

High-Fidelity, Near-Field Microwave Gates in a
Cryogenic Surface Trap

Candidate: Marius A. Weber

Supervisor: Prof David Lucas

30th of November, 2022

Abstract

We present a novel dynamical decoupling strategy for near field microwave gradient driven, Mølmer-Sørensen style, two-ion quantum logic gates, which suppresses errors from both fluctuations in the qubit frequency and imperfection in the decoupling drive itself. Using a microwave-integrated surface-trap which is operated cryogenically at ~ 25 K and a magnetically insensitive $^{43}\text{Ca}^+$ qubit at 288 G, we demonstrate a 331 μs two-ion quantum logic gates, with $4.9(11) \times 10^{-3}$ logic error probability. This is below the 1% error threshold required for quantum error correction and represents a $\sim 10\times$ gate time reduction when compared to previously demonstrated near field gradient driven microwave gates below the 1% error probability threshold. Additionally, two faster gates were demonstrated without the use of dynamical decoupling. Respectively, these two gates had gate operation durations of 216.8 μs & 153.8 μs and measured gate error probabilities of $8.5(20) \times 10^{-3}$ & $9.8(21) \times 10^{-3}$. Further, we develop a method for rapid calculation of ion transport operations. We successfully demonstrate ion transport as well as crystal splitting and merging operations within two different ion traps using the waveforms calculated by this ion transport toolbox.

Acknowledgements

My first thanks go to my supervisor Prof. David Lucas for giving me the opportunity to play with atoms, lasers, and soldering irons in his lab. David has been a steady source of enthusiasm and encouragement. His experience has been invaluable when he poked his head into the lab. I am further grateful for his careful proof reading of this thesis.

This work was made possible through the combined and steady effort of many members of the Oxford Ion Trap group. From within the microwave trap team: Dr. Ryan Hanley has accompanied me throughout the vast majority of this work and his optimism, experience, and near unlimited patience have kept me sane over the years. Dr. Clemens Löschnauer's meticulous support kept me on the right track on many a late evening. Dr. Thomas Harty first introduced me to ion trapping and provided a wealth of experience developing electronics. Dr. Jochen Wolf designed large swathes of the apparatus and provided a level of organisation to aspire to. Dr. Mario Gely brought fresh enthusiasm to the team in the hour of need and provided valuable help proof reading this thesis. Finally, Aaron Leu and Molly Smith bring some much needed new life to the microwave trap team.

I am thankful to Dr. David Nadlinger who is always up for a discussion, provided a wealth of insight, and listened to my frustrations. Dr. Raghavendra Srinivas was always ready to share his experience and helped proof read this thesis. Peter Drmota supported me when operating the HOA2 trap and shared his experience simulating quantum dynamics in *Julia*. Dr. Chris Ballance always found time to share his rich knowledge of physics. There are many more members of the Oxford Ion Trap group which I have been fortunate to spend time with, thank you for being a fantastic team. It would not have been the same without you. The many department support staff who smoothed the road are also deserving of my gratitude.

I would also like to thank Prof. Zoran Hadzibabic and Prof. Stafford Withington for their excellent undergraduate courses which inspired me to set out on this journey. I send a special thank you to Elisabeth Wannenwetsch, Andrea Heering, and Jörg Heinrichs. These every day heroes not only believed in me when few others did, but their support is a large part of getting me where I am today.

I owe a debt of gratitude to my friends who listened, distracted, and supported me throughout the years, this road would not have been possible without you.

Danke an meine Eltern, Jutta und Alfons, und meinen Bruder Julian für ihre stete Zuversicht und Unterstützung. Danke Jutta für deine unermüdliche Ermutigung meinen eigenen Weg zu gehen und Alfons für die vielen technischen Spielereien seit frühestem Alter.

Finally, I thank and acknowledge the funding provided by ARO and EPSRC via NQIT and QCS Hub which enabled my work.

Contents

Contents	v
1 Introduction	1
1.1 Quantum information processing	1
1.2 Trapped ions for quantum information processing	2
1.3 Scaling a trapped ion quantum processor	4
1.3.1 Ion transport	5
1.3.2 Microwave driven quantum logic with ions	6
1.3.3 Cryogenic operation	13
1.4 Structure of thesis	13
2 Operating an Ion-Trap with $^{43}\text{Ca}^+$ at 288 G	15
2.1 The trapping potential	15
2.1.1 Ponderomotive force	15
2.1.2 Static electric fields	16
2.2 Photoionisation loading	16
2.3 Structure of $^{43}\text{Ca}^+$	17
2.3.1 Fine structure	17
2.3.2 Hyperfine structure	17
2.3.3 Magnetic field insensitive transitions	18

2.4	Laser cooling	21
2.4.1	Doppler cooling	21
2.4.2	Ground state cooling	24
2.5	Qubit state preparation	24
2.5.1	Microwave enhanced state preparation	26
2.5.2	Microwave transfer to the clock qubit	27
2.6	Qubit readout	28
2.6.1	Shelving procedure	28
2.6.2	Microwave transfer from clock qubit	30
2.6.3	Threshold readout	33
3	Two-Qubit Gate Theory	35
3.1	Microwave spin-dependent force	36
3.1.1	Hamiltonian of trapped ion microwave interaction	36
3.1.2	Effective Lamb-Dicke parameter	39
3.2	The Mølmer-Sørensen gate	39
3.2.1	The ideal Mølmer-Sørensen gate	40
3.2.2	Gate fidelity	44
3.2.3	Common error sources	46
3.2.3.1	Motional frequency miscalibrations	46
3.2.3.2	Qubit frequency error	47
3.2.3.3	Kerr coupling to spectator modes	48
3.2.3.4	Off-resonant excitation of hyperfine transitions	51
3.3	Error suppression techniques	51
3.3.1	Off-resonant excitation	52
3.3.2	Sideband Walsh modulation	54
3.3.3	Dynamical decoupling of qubit detuning errors	56

3.3.4	Walsh modulated dynamical decoupling of qubit detuning errors	58
4	Apparatus	64
4.1	The ion trap assembly	64
4.1.1	Partial nulling of the microwave field	65
4.1.2	Cryogenic Pillbox	67
4.1.3	DC electrode filters	68
4.2	Vacuum Assembly	68
4.3	Imaging system	69
4.4	Cryogenic equipment	70
4.5	Magnetic Field coils	70
4.5.1	Active stabilisation of electromagnet current	72
4.5.2	Ion servo	76
4.6	Laser Systems	76
4.6.1	Beam polarisations	80
4.7	Control System	82
4.7.1	Realtime controlled hardware	83
4.7.1.1	DIO-BNC	83
4.7.1.2	Fastino v1.2	84
4.7.1.3	Phaser (baseband)	84
4.7.1.4	Sampler	87
4.7.1.5	Urukul (AD9910)	87
4.7.1.6	SU-Servo	87
4.7.2	Reference Clock	88
4.8	Amplitude stable RF signal chain	88
4.9	Microwave signal chain	90
4.9.1	Microwave duty cycle management	97

4.10	Ground loop control	99
4.10.1	Noisy Kasli switch mode power supply	100
4.10.2	Grounding configuration rules	100
4.10.3	Heating rate improvement	101
5	Time Dependent Trap Control	103
5.1	Computational approach & goals	104
5.1.1	Aims of the computational package	106
5.1.2	Controlled potential parameters	106
5.2	The trap model	107
5.3	Calculating electrode voltages	109
5.3.1	Static cost-function	110
5.3.2	Computation speed optimisation	112
5.4	Evolving potential wells	112
5.4.1	Dynamic cost-function	113
5.4.2	Well trajectories	114
5.5	Ion separation	117
5.5.1	Ion-separation trajectory	117
5.5.2	Separation algorithm	120
5.6	Characterisation	125
5.6.1	Separation and transport reliability	127
5.6.1.1	HOA2	128
5.6.1.2	Microwave integrated trap	128
5.6.2	Radial mode heating during transport	130
6	Two-Qubit Gate Implementation	133
6.1	Gate sequence	133
6.1.1	Cycle alignment	134

6.1.2	Laser cooling	135
6.1.3	State preparation	136
6.1.4	Readout	136
6.1.5	Clock qubit operations	138
6.1.6	Housekeeping operations	139
6.1.7	Duty-cycle management	140
6.2	Calibration procedure	141
6.2.1	Duty-cycle management	141
6.2.2	Pulse shaping	144
6.2.3	Transfer pulse calibration	144
6.2.4	Analysis pulse calibration	144
6.2.5	Microwave field characterisation and AC-Zeeman shift	145
6.2.6	Sideband amplitude balancing	147
6.2.7	Dynamical decoupling drive phase	150
6.2.8	Fine gate duration adjustment & sideband detuning	151
7	Two-Qubit Gate Characterisation	154
7.1	Gate measurements	154
7.1.1	Parity analysis	155
7.1.2	SPAM correction	157
7.1.3	Dynamical decoupling strength	159
7.1.4	Gate input power	160
7.2	Error characterisation	162
7.2.1	Gate duration calibration error	162
7.2.2	Mode frequency stability	162
7.2.3	Heating rate	164
7.2.4	Coupling to spectator modes	165
7.2.4.1	Frequency addressing of motional mode	165

7.2.4.2	Kerr coupling	167
7.2.5	Motional dephasing	168
7.2.6	Microwave field	168
7.2.6.1	Microwave field characterisation error	169
7.2.6.2	Slow microwave field drifts	170
7.2.6.3	Fast microwave field drifts	172
7.2.6.4	Broadband microwave noise	173
7.2.7	Off-resonant qubit excitation by sidebands	173
7.2.8	Qubit spectator transitions	173
7.2.9	Raman laser photon scatter	176
7.2.10	Error summary	176
8	Conclusions	177
	Bibliography	181
A	Quantum-Jump Enhanced SPAM Optimisation	198
A.1	3D _{3/2} to 4P _{3/2} quantum jump enhanced spectroscopy	198
A.2	Optical state-preparation optimisation	200
B	Microwave Field and RF Resonance Post “Experiment Upset”	204
B.1	Spontaneous RF resonance jumps	204
B.2	Mode frequency variation	206
B.3	Shift of microwave null	206
B.4	Shift of microwave resonance	210
C	Fastino Output Noise Characterisation & Kasli Ground Loops	211
C.1	Measurement setup	211
C.2	Initial performance	212
C.3	Fixes to noise sources which are part of the Fastino design . . .	213

C.4 Kasli switch-mode noise (with common mode current chokes in Fastino outputs)	214
C.4.1 Modifications that didn't help reduce spurs	217
C.5 Kasli switch-mode noise (without common mode current chokes in Fastino outputs)	217
C.6 Summary	217
D Accounting for Trap Filters in Potential Well Evolution	219
D.1 Compensating DAC voltages for trap filters	219
D.1.1 On Trap Electrode Filters	219
D.1.2 Inverting the Filter	220
D.1.3 Notes on Fourier Transform	220
E Rules of Thumb for Potential-Well Solver Parameters	221

Chapter 1

Introduction

1.1 Quantum information processing

Quantum mechanical systems can efficiently, that is in polynomial time, solve problems that are believed to be intractable for a conventional computer [1, 2]. One of these problems is simulation of quantum mechanics itself, with many potential applications in physics, chemistry and medicine. Beyond simulating quantum mechanics, the list of algorithms achieving improved scaling on quantum computers is steadily growing. Among the most famous are the Grover search and Shor factorisation algorithms [3, 4]. Efficient quantum algorithms promise large dividends when run on a sufficiently large and reliable quantum computer. However, efforts to build such a machine face stark challenges formalised in the DiVincenzo criteria for scalable quantum computers [5]. The five core requirements identified by DiVincenzo may be summarised as follows:

1. A scalable physical system with well characterized qubits (the quantum mechanical analogue to a bit)
2. The ability to initialize the state of the qubits to a simple fiducial state,

such as $|000\dots\rangle$

3. Long relevant decoherence times, much longer than the gate operation time
4. A “universal” set of quantum gates (an arbitrary gate can be constructed from single qubit rotations and a suitable two qubit gate)
5. A qubit-specific measurement capability

In face of these challenges, many platforms are being explored as the basis of future quantum computers. Among the leading candidates are photonic [6], superconducting [7], nitrogen vacancy [8], Rydberg atom [9], and ion-trap [10] based systems.

1.2 Trapped ions for quantum information processing

Starting in the mid 90s, there has been a growing interest in using trapped ions for quantum information processing (QIP) [11, 12]. Ion traps leverage the net electric charge of ions to confine them in free space. Ion trap quantum information processors typically employ a Paul trap which is capable of loading a linear chain of ions in a potential well. These trapped ions typically have similar charge and mass. The natural repulsion of similarly charged ions ensures good mutual isolation of the internal quantum states of multiple ions held within the trap. Meanwhile the suspension of an ion-crystal in vacuum provides isolation of ion internal states from the environment.

Trapped ion qubits are typically encoded in the superposition of internal eigenstates of the ion. For ions of the same isotope, these internal eigenstates are identical, avoiding per qubit calibration requirements. The eigenstates

encoding the qubit are usually chosen to be in stable or meta stable electronic states that are separated by an electric-dipole forbidden transition. Common choices for the qubit transitions are microwave domain magnetic dipole [13] and optical electric quadrupole [14] transitions. On such transitions the decay timescales are sufficiently long to not constitute significant logic errors. Further, when operating on suitable transitions, long coherence times on the scale of minutes [15, 16] to hours [17] have been demonstrated with trapped ion qubits.

Logic operations on trapped ion qubit transitions are commonly driven using laser beams which can be focused to enable ion addressing. Using trapped ions, a “universal” set of logic operations have been demonstrated, with typical gate timescales of $\sim 100 \mu\text{s}$ [18, 19, 20, 21, 22], although gates as fast as $\sim 1 \mu\text{s}$ have been demonstrated [23]. Trapped ion qubit coherence times therefore far exceed the required gate operation times making trapped ions suitable for both memory and logic operations. As noted by DiVincenzo [5], use of these gates in large scale quantum processors is likely to require the use of quantum error correction protocols [24, 25]. The most error resilient quantum error correction scheme is the surface code, which requires a threshold logic operation error probability of 1% [26]. Using trapped ion quantum processors, all required logic operations have been demonstrated with error probabilities below this threshold. However, implementing error correction with near 1% gate error probabilities incurs a large overhead in the number of physical qubits and gates required to encode logical operations. Due to this error correction overhead, practically useful error-corrected processors are unlikely to be built in the near-term. Consequently, “noisy intermediate scale quantum” (NISQ) processors are a more realistic near-term use of trapped-ion quantum information processors [2]. In the absence of error correction, such NISQ processors

are critically dependent on the reliability of their logic operations, which is a strong suit of the ion trap based quantum processors, when compared to other quantum processor technologies [10]. Due to their leading logic operation reliability, trapped ion based quantum processors are well suited to satisfy the third and fourth DiVincenzo criteria.

Looking beyond logic operations, ion-trap based quantum processors have demonstrated the highest fidelity qubit state preparation and readout capabilities of any QIP platform [10] by employing optical pumping, and shelving techniques [27]. To date, state preparation and measurement error probabilities as low as $(9.0 \pm 1.3) \times 10^{-5}$ have been demonstrated [18, 28, 29]. These optical state preparation and measurement operations can address individual ions using spacial selectivity, such as focused laser beams. Consequently trapped ion quantum computing can be considered leading on the second through to fifth DiVincenzo criteria. However, scalability remains a major challenge for trapped ion quantum information processors.

1.3 Scaling a trapped ion quantum processor

To date, ion trap quantum processors have been demonstrated with tens of coherently controlled qubits [30, 31]. These state of the art devices are yet too small to attain practical superiority over conventional computers, also known as ‘quantum supremacy’ [32, 33]. Scaling up ion trap based quantum processors to practically relevant processors faces a range of challenges. For a quantum processor to be considered scalable, it should be able to increase its basic computational elements, such as qubits and applied logic gates, without suffering an incommensurate increase in required resources. Such resources include properties such as operation speed, device cost, footprint and required energy [10].

Some of the challenges to scaling an ion-trap based quantum processor include crosstalk when addressing specific qubits, and the overhead of laser cooling all trapped ions [10]. However, in the following we will focus on the potential role of ion transport and microwave driven logic operations in tackling some of the scaling challenges faced by trapped ion quantum processors.

1.3.1 Ion transport

Naively one might attain a larger scale ion trap quantum processor by adding increasing numbers of ions to a sufficiently deep potential well. However, the leading multi-qubit gate schemes for trapped ion based processors require spectral addressing of individual normal modes of the collective ion motion. Confining more ions into a single potential well results in spectral crowding of these normal modes, slowing down gate operations, lest logic errors be induced by the spectating modes. Further, the coupling strength of the logic gate interaction decreases with increasing number of confined ions, requiring increasing drive powers to retain the same operation speed. To address this issue the use of a ‘quantum charge coupled device’ (QCCD) architecture has been proposed [12, 34, 35]. This ion trap architecture confines a limited number of ions in any given potential well, but scales to larger qubit numbers by employing multiple wells. The different potential wells are located in a range of specialised trap regions which may have dedicated uses such as qubit memory or inter-qubit logic operations.

As multi-qubit logic gates for trapped ions commonly rely on the collective ion motion as a ‘quantum bus’ coupling the state of different ions, splitting the logic encoding ions across multiple wells breaks the all to all qubit connectivity available when employing a single potential well. Consequently, in the QCCD architecture ions must be re-arranged between potential wells to facilitate

required logic operations. These re-arrangements should include the capability to merge and separate ions into and out of a shared potential well, as well as moving and reordering ions across the QCCD.

As the number of trapped ions increases ion-loss also becomes a larger concern. Considering a device using 1000 ions with a 4 hour lifetime of each trapped ion we expect to lose an ion every ~ 10 s. In a quantum processor this represents the destruction of a qubit. For calculations exceeding this timescale, it becomes critical to detect and rectify such ion loss errors. It is likely that ion transport to replace lost ions will feature in large scale ion trap processors, whether or not they adopt the QCCD architecture.

Though the QCCD architecture alleviates significant obstacles to scaling trapped ion processors, it is not without its own scaling challenges. Key limiting factors are expected to be the limited optical access available for laser delivery, as well as the large number and complexity of electronic control signals for transport operations. Due to the complexity of control signals required to perform transport operations, QCCD devices might might reasonably run into significant challenges when using ~ 1000 ions. None the less, the QCCD architecture is a promising avenue for scaling trapped ion based quantum processors as a device containing ~ 100 logic ions is expected to be suitable for NISQ processors with ‘quantum supremacy’ for practical tasks [2]. A possible mitigating strategy to overcome the scaling challenge associated with the large number of control signals is integration of the signal generation into the trap structure, which is an active field of research [36, 37].

1.3.2 Microwave driven quantum logic with ions

To avoid the optical access limitations encountered by the QCCD architecture, it is desirable to integrate the delivery of fields driving logic operations

into the trap structure. One such approach is the use of chip integrated optics to deliver laser light to targeted ions in the trap [10, 38]. However, when using qubits encoded in a microwave-domain magnetic dipole transition, it is also possible to drive the required quantum logic operations using microwave magnetic fields near the qubit transition frequency. Using integrated microwave instead of integrated optical signals for driving quantum logic operations comes with a range of advantages. Firstly, the optical Raman transition approach when driving microwave domain qubits using lasers suffers from fundamental photon scattering induced gate errors which are avoided by the use of microwave fields [20, 39, 40]. Secondly, the amplitude and phase of microwave signals can be controlled more precisely and stably than for optical signals. Thirdly, compared to laser optics, microwave electronics is a more mature technology with readily available miniaturised high-performance components. Due to well established micro-fabrication techniques, integration of microwave electrodes into ion traps is significantly simpler than fabricating suitable traps with integrated optics. Using such microwave integrated ion traps, both single and multi qubit gates with error probabilities below the 1% error correction threshold have been demonstrated. Microwave driven single qubit operations have been demonstrated with average error probability of $1.0(3) \times 10^{-6}$, which is the lowest single qubit operation error demonstrated to date in any quantum computation platform [18]. Further, “laser free” two qubit gate operations have been demonstrated with error probabilities of 0.3% and below [21, 41, 42]. None the less, two qubit logic operations remain a weak point of microwave driven quantum logic, as the ~ 1 ms time required for these high fidelity two-qubit operations is much slower than the ~ 100 μ s gate durations typically achieved with laser driven quantum logic operations.

The speed disparity between microwave and laser driven two qubit logic

operations arises from the need to generate a spin dependent force when generating ion-ion interactions via the collective ion motion. For magnetic dipole qubits, spin dependent forces are generated by magnetic field gradients. For optically generated spin dependent forces, the free space wavelength of optical light provides sufficient field gradient to generate practical spin dependent forces. Due to the $\sim 10^7$ times longer wavelength of microwave radiation, a similar far-field approach is not practical for microwave radiation [12]. Fortunately, a variety of techniques for generating stronger spin dependent forces than available with free space microwave radiation have been established and are discussed below. An overview of demonstrated two-qubit gate durations and error probabilities is shown in figure 1.1.

Static magnetic fields One way of enhancing the microwave mediated spin dependent force is to place the ions in an inhomogeneous static magnetic field [43, 44, 45]. The energy eigenstates of ions placed in a static magnetic field experience the Zeeman effect, which modifies the eigenstate energies depending on the experienced magnetic field. Consequently, in a spatially varying magnetic field the eigenstate energies also vary spatially which corresponds to a force. When magnetic dipole transitions are driven in such a varying field, the ion is displaced, depending on its original spin state, a “spin dependent force”. This technique is also known as “MAGnetic Gradient induced Coupling” (MAGIC) and has the benefit of additionally enabling spectral addressing of individual ions via the spatially varying Zeeman shift caused by the inhomogeneous magnetic field.

The enhanced spin dependent force and consequent faster gate speed achieved by the MAGIC method depends on attaining large static magnetic field gradients and using transitions with a large magnetic field sensitivity. The required

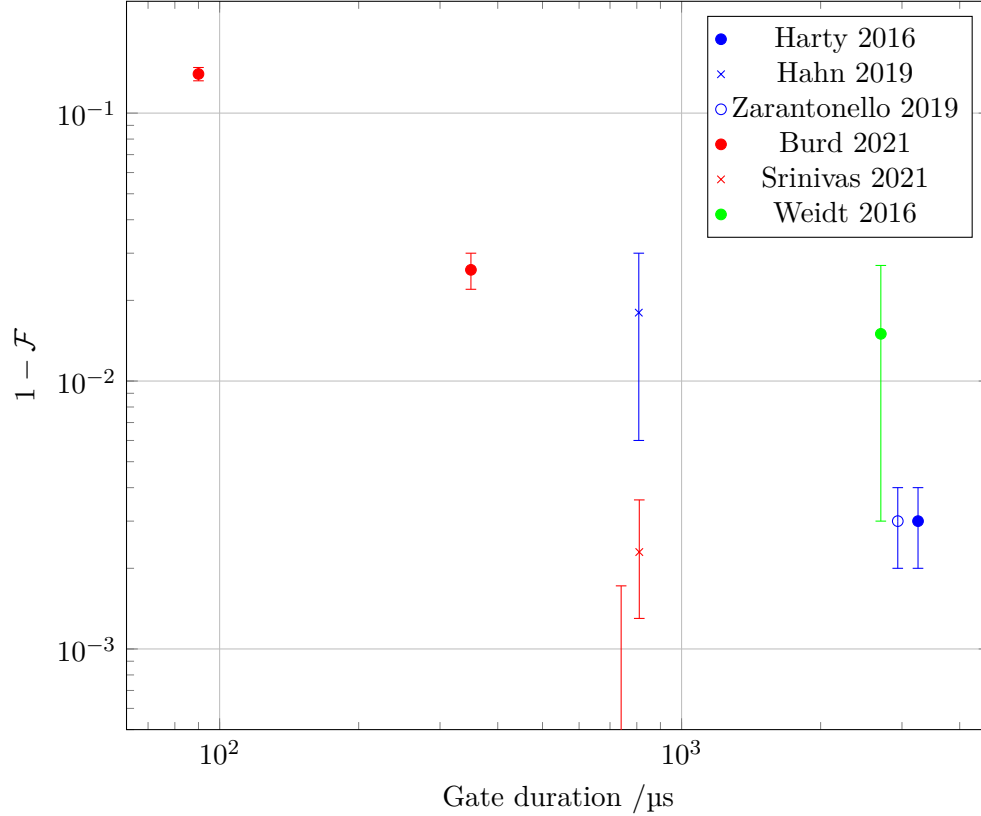


Figure 1.1: **Laser free two-ion gates** - Selection of laser free two ion gate demonstrations using the near field microwave gradient, near-motional-frequency magnetic field gradient and MAGIC gate mechanisms. Demonstrations using near field microwave gradients are shown in blue. Results using the MAGIC and near-motional-frequency magnetic field gradient methods are shown in green and red respectively. The fastest near field microwave gradient driven gate demonstrations which are below the 1% error probability required for error correction have gate durations $t_g = 2.938$ ms. The fastest demonstrated error correctable laser free gate was performed using the near-motional-frequency magnetic field gradient method with a gate duration $t_g = 0.74$ ms.

magnetic field gradient in this approach may be generated using DC current electromagnets [46] or permanent magnets [45, 47], allowing the production of magnetic field gradients without additional energy dissipation. However, the MAGIC approach suffers a drawback in the form of its required magnetic field sensitivity of the qubit transition. The qubit magnetic field sensitivity allows environmental magnetic field noise to couple to the qubit state, leading to decoherence. During two-qubit operations, decoherence from ambient magnetic field noise may be mitigated using atomic dressing techniques [48]. However, the qubits remain unprotected during memory operations and single qubit rotations. To mitigate this qubit memory limitation, dynamical decoupling [49] or a second qubit transition which is not magnetic field sensitive [18] need to be employed. A further drawback of the MAGIC approach is an unfavourable scaling of gate speed with higher motional frequencies, which are desirable to suppress sensitivity of the ion motion to electric field noise.

Near field microwaves Another way of enhancing the spin dependent force generated by a given microwave field strength is by placing the ion in the near-field of a microwave current carrying electrode. In the near-field regime, the magnetic field gradient is not constrained by the microwave radiation wavelength, but rather by the geometry of the microwave current distribution, which may be engineered by shaping the electrode structure. Considering the magnetic field a distance d from a current carrying wire, we note that the magnetic field and magnetic field gradients for a given current scale as $\propto 1/d$ and $\propto 1/d^2$ respectively. Consequently moving the ion closer to the electrode structure improves both the available magnetic field gradient and the magnetic field gradient to amplitude ratio.

Two qubit gates using the near field gradient technique were first proposed

in 2008 [50] followed by experimental demonstration in 2011 [13]. Since then, near field gradient two qubit gates have rapidly advanced in speed and gate fidelity. In preceding works this technique was used to generate two-qubit gates with ~ 3 ms gate durations and 0.3(1)% gate error probabilities which pass the 1% error correction threshold [41, 42]. These more recent experiments further improve the microwave magnetic field to gradient ratio by engineering electrode structures which produce a magnetic field zero crossing near the ion position. This increased microwave gradient to amplitude ratio helps suppress the driving of undesired qubit dynamics during gradient driven two-qubit operations.

A key advantage of using near field microwave radiation for driving two qubit gates is that, as the spin dependent force does not depend on a spatially varying Zeeman shift, a magnetic field insensitive transition may be chosen for the qubit. This allows for significantly longer qubit coherence times which can be as long as minutes [15, 16] to hours [17]. However, the use of near field microwave gradients also suffers from disadvantages associated with the high frequency of the oscillating gradient.

The two qubit gate operation time of near field microwave driven gates is determined by the magnitude of the magnetic field gradient. In a given ion and electrode arrangement, increasing the gate speed requires an increase in the microwave electrode current. However, due to the limited skin depth of microwave currents, resistive losses in the trap electrodes and wire bonds increase with the microwave frequency. Further, the currents in the microwave electrode induce currents in neighbouring electrodes, which oppose the generated magnetic field and field gradient. Both the image currents and resistive losses act to reduce the attained magnetic field gradient and consequently two-qubit gate speed for a given microwave input power and scale with the

frequency of the magnetic field gradient.

Near-motional-frequency magnetic field gradient The most recently proposed method of enhancing the microwave generated spin dependent force, employs a magnetic field gradient that oscillates at a frequency close to the resonance of the collective trapped ion motion [51, 52]. This scheme represents somewhat of a hybrid approach of the two schemes, where a spin dependent force is coupled via the qubits magnetic field sensitivity, but the magnetic field gradient is generated in the near field of current carrying electrodes. Recent, demonstration of this gate scheme attained gate error probabilities $< 0.17\%$ with a gate operation time of 0.74 ms [21]. Together with the near field microwave gradient results demonstrated by Harty [41] and Zarantonello [42], these results represent the state of the art in “laser free” two ion logic gates.

As the frequency of the strong magnetic field gradient in the near-motional-frequency scheme is typically several orders of magnitude below the microwave frequencies associated with the qubit transition, this scheme mitigates the input power disadvantages of the near field microwave gradient scheme. As the physics of the near-motional-frequency scheme are most similar to the static gradient approach, it also shares the same drawback of requiring mitigation of the reduced memory performance of magnetic field sensitive qubits from ambient magnetic field noise. However, the speed of near-motional-frequency gates scales more favourably with the frequency of the trapped ion collective motion.

1.3.3 Cryogenic operation

To improve on the slow ~ 1 ms operation times associated with oscillating-gradient “laser free” two qubit gates, it is appealing to leverage the favourable scaling of gate speed with reduced ion-electrode distance d . However, moving the trapped ions closer to the trap electrodes results in reduced depth of the confining potential well. This reduced potential depth makes it likelier for ions to be knocked out of the ion trap by background gas collisions. Further, as the ion approaches the trap surface it is subjected to increased electric field noise from the trap electrodes. This noise is not generally understood, but is believed to be associated with surface patch potentials on the electrodes [53, 54, 55, 56]. As the laser-free two qubit gate schemes engineer qubit interactions through the collective ion motion, the increased electric field noise from operating closer to the electrodes may result in an intolerable increase in two qubit gate errors. Operating the ion trap and surrounding pillbox at cryogenic temperatures helps offset both the lifetime and electric field noise challenges by freezing out background gas atoms and empirically observed reduced field noise [54, 57]. Further, cooling the ion trap allows wire bonds and electrode structures to tolerate larger currents without damage¹.

1.4 Structure of thesis

In this work we develop and characterise a toolbox for on demand calculation of ion transport operations. Further, we improve on the speed of near field microwave driven laser-free two qubit gates. With an eye on these aims, chapter 2 presents a overview of basic ion trap operation, which is followed by

¹This is due to a combination of reduced resistive losses, increased thermal headroom to damaging temperatures, and increased thermal conductivity aiding extraction of dissipated energy.

a presentation of near field microwave gradient driven two qubit gate physics in chapter 3. The experimental apparatus used for the majority of this work is detailed in chapter 4. Within chapter 5 the operation and characterisation of a toolbox providing on-demand calculation of ion trap transport operations is presented. We follow these ion transport results with implementation details of the two qubit gate characterisation and tune up procedure in chapter 6. Attained two-qubit gate speeds, fidelities and likely error sources are presented in chapter 7. Finally, conclusions are presented in chapter 8.

Chapter 2

Operating an Ion-Trap with $^{43}\text{Ca}^+$ at 288 G

In the following a brief theory overview for trapping ^{43}Ca at 288 G is presented.

2.1 The trapping potential

A key aspect for manipulating trapped-ion qubits is spatially confining, also known as trapping, the ions in free space. As ions carry a net charge it is convenient to do this using electric fields in a linear Paul trap [58, 59].

2.1.1 Ponderomotive force

The free space Maxwell equations constrain the in-vacuum electric potential V via the notable result $\nabla^2 V = 0$. This indicates that it is not possible to create a free space electric potential minimum using static fields. Linear Paul traps circumvent this issue by using radio-frequency electric quadrupole fields. The ponderomotive force from this field gives rise to a pseudo-potential. The pseudo-potential from an oscillating electric quadrupole has a minimum at

the centre of the quadrupole field. In a linear Paul trap the pseudopotential provides confinement in the two directions perpendicular to the quadrupole axis. The pseudo-potential shares the translational symmetry of the underlying electric quadrupole field. Consequently, the pseudo-potential provides confinement in the two ‘radial’ directions, but does not provide confinement along the quadrupole axis.

2.1.2 Static electric fields

The pseudo-potential generated through radio frequency fields establishes two-dimensional confinement of the ions. The effective potential experienced by the ions can be further shaped by applying static electric fields. These may complement the pseudopotential by establish confinement along the quadrupole axis. In the case of segmented electrode Paul traps, static electric fields may also be used to tune many aspects of the potential experienced by the ion (see chapter 5).

2.2 Photoionisation loading

We load ions into the trap potential using photoionisation of neutral calcium vapour. Neutral calcium vapour is generated by heating a ^{43}Ca enriched calcium sample in a resistively heated oven [60]. Isotope selectivity is achieved using a standard two-stage ionisation process [61]. In our case we use a 423 nm laser driving the $4s^2\ ^1S_0 \longleftrightarrow 4s4p\ ^1P_1$ transition. As isotope shifts on this transition are much larger than the transition linewidth, this process is isotope selective. Excited atoms are then ionised using a 378 nm incoherent beam produced by a diode. When the 423 nm and 378 nm beams are overlapped with the trapping potential described in 2.1, photoionisation of suitably slow atoms results in trapped ions.

2.3 Structure of $^{43}\text{Ca}^+$

We now turn our attention to the stored ion. The fine and hyperfine structure of the ion is essential to our laser cooling, state-preparation, readout, and quantum logic operations. In our case, we use $^{43}\text{Ca}^+$ ions at a magnetic flux density of 288 G.

2.3.1 Fine structure

The $^{43}\text{Ca}^+$ ion has a single valence electron, giving rise to an alkali metal like fine-structure. For our purposes, we are interested in the $4S_{1/2}$ ground level as well as the two lowest lying P and D states. The associated transition level structure is shown in figure 2.1.

2.3.2 Hyperfine structure

^{43}Ca has a nuclear spin of $I = 7/2$ giving rise to a multitude of hyperfine states within each level of the $^{43}\text{Ca}^+$ fine structure. In the presence of magnetic fields the eigenstates and energies are modified by the Zeeman effect. The experiments in this thesis are carried out at a field of 288 G. This puts us outside the low field regime where F is no longer a good quantum number. For convenience, we will none-the-less refer to states using the notation $|F, m\rangle$. This should be understood as referring to the state that is attained when starting at low magnetic field with $|F, m\rangle$, followed by adiabatically increasing the magnetic field. In the case of the $4S_{1/2}$ ground level the states are given analytically via the Breit-Rabi formula. However, other levels require numerical evaluation. Figure 2.2 shows an overview of the hyperfine structure of the relevant levels.

For this work, the hyperfine structure of the $4S_{1/2}$ ground level is of partic-




Figure 2.1: $^{43}\text{Ca}^+$ **level structure** - The diagram shows the fine structure levels used within the experiment. Laser driven, optical transitions are indicated by red and blue arrows. The $4S_{1/2}$ hyperfine structure is enlarged and labelled with the transition frequencies of the $|F = 4, m = +4\rangle$ to $|F = 3, m = +3\rangle$ ‘stretch’ transition and the $|F = 4, m = +1\rangle$ to $|F = 3, m = +1\rangle$ ‘clock’ qubit transition. Figure reproduced from Weber et al. [62].

ular interest. Figure 2.3 highlights this structure and shows transitions used within this work.

2.3.3 Magnetic field insensitive transitions

One consideration when choosing an ion and transition for quantum logic is the coherence time of states encoded in that transition. For hyperfine, magnetic-dipole transitions this coherence time is often limited by magnetic field noise [64]. The noisy magnetic field creates uncertainty off the transition frequency via the Zeeman effect. Fortunately, not all transition frequencies are equally sensitive to magnetic field changes. Cases where the transition frequency is at a stationary point with respect to magnetic field changes are known




Figure 2.2: **Hyperfine structure of $^{43}\text{Ca}^+$ in a 288 G magnetic field** - Energy level diagram showing the hyperfine state energies of the five $^{43}\text{Ca}^+$ energy levels used in this work. The average energy of each level is indicated by a long solid line. Hyperfine states are indicated by shorter lines which are colour coded according to the F quantum number of the state attained when adiabatically decreasing the magnetic field to the low field regime. The lateral position of energy levels is arranged according to their total electronic angular momentum quantum number J . Within the energy levels hyperfine states are arranged according to their m quantum number. The vertical position of the energy levels and hyperfine states is to scale with their relative energies, where the displacement of the hyperfine states from the average level energy is scaled up 30 000 fold for improved visibility. This figure is reproduced from Wolf [63].




Figure 2.3: $4S_{1/2}$ hyperfine structure at 288 G (to scale). The F quantum number is indicated via the level colour, the m quantum numbers are represented via the horizontal level position and are labelled at the bottom. The transition between clock qubit states is indicated in blue and labelled ‘q’. Optical pumping prepares the ion in the $|F = 4, m = +4\rangle$ state. The ion can be transferred to the $|F = 3, m = +1\rangle$ clock qubit state by applying successive microwave π -pulses on the transitions labelled ‘a’, ‘b’, and ‘c’. Optical state readout also occurs from the $|F = 4, m = +4\rangle$ state. Applying pulses to transitions ‘a’, ‘b’, and ‘c’ in opposite order transfers the $|F = 3, m = +1\rangle$ state amplitude back to the $|F = 4, m = +4\rangle$ state. During readout, the qubit may be in any superposition of the $|F = 3, m = +1\rangle$ and $|F = 4, m = +1\rangle$. Consequently, when addressing transition ‘c’ during the readout transfer, the spectating transition ‘d’ (dotted) may be off-resonantly excited by microwaves which are spectrally addressed to transition ‘c’. The dashed transitions ‘e’ and ‘f’ are used to assist with the optical state-preparation process.

as clock transitions. In the case of $^{43}\text{Ca}^+$ at 288 G the $|F = 3, m = +1\rangle$ to $|F = 4, m = +1\rangle$ transition is a clock transition. This transition is henceforth referred to as the “clock qubit”.

2.4 Laser cooling

Ions within ion traps commonly require regular cooling. In our case, the reason for this is two-fold. Firstly, the ion-motion is gradually excited due to environmental influences. Excessive ion-motion leads to undesired loss of the ion from the trap. Secondly, the Mølmer-Sørensen style two-qubit logic operations implemented in this thesis become more error sensitive with increased ion-motion. We therefore employ laser-cooling techniques to reduce the ion-motion as far as practical and well below the standard Doppler cooling limit. Here, we provide a brief overview of the employed cooling procedures. As is evident from section 2.3.2, $^{43}\text{Ca}^+$ at 288 G has a rich hyperfine structure, complicating standard laser cooling techniques. A more in depth treatment and characterisation of the cooling procedures we use is presented by Löschnauer in [65].

2.4.1 Doppler cooling

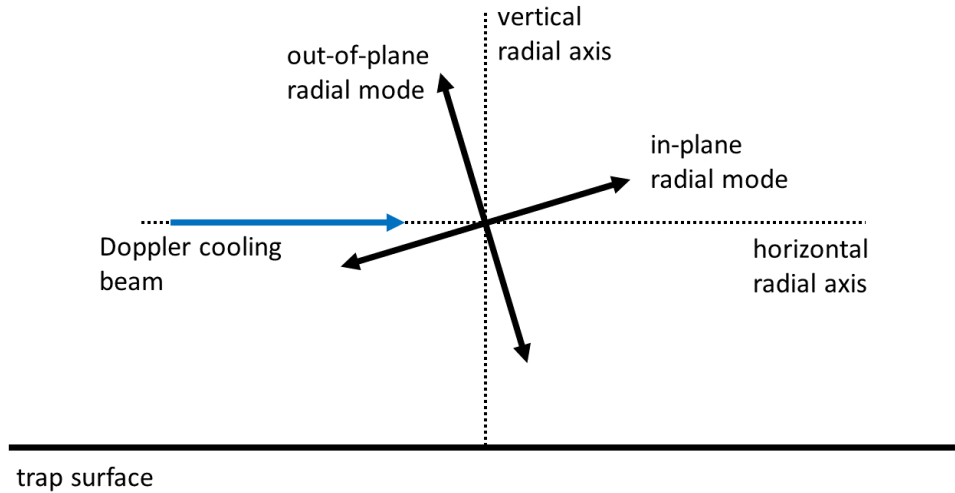
Retention of ions within the trap potential is facilitated by Doppler cooling. Doppler cooling is a standard laser cooling technique that leverages the Doppler shift induced by a particles motion. As an ion moves towards a light source, the incident light is blue-shifted. Doppler cooling uses a red detuned laser beam, for which such a blue shift results in increased photon scattering as the laser is shifted towards resonance with an optical transition. As the ion is shifted towards resonance it feels an increasing force from the larger photon scattering rate near resonance. The resulting force-field is non-conservative

and dissipates ion motion. For trapped ions, the motion is decomposed into normal modes of the collective ion motion. To cool all modes of the ion motion, we therefore want the propagation direction of the cooling laser beam to have a projection onto all oscillation directions of the normal modes, as shown in figure 2.4.

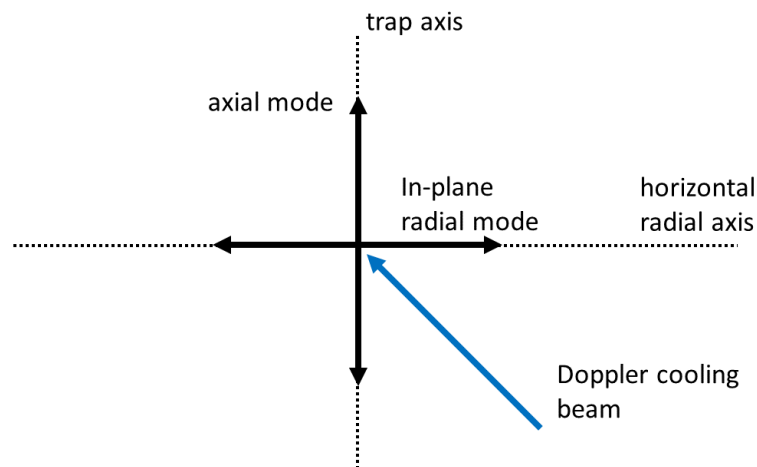
We perform Doppler cooling using the 397 nm $4S_{1/2}$ to $4P_{1/2}$ transition. The propagation direction 397 nm Doppler cooling laser beam is chosen to cool all motional modes (see figure 2.4). As the $4P_{1/2}$ states can decay to the $3D_{3/2}$ level, a 866 nm re-pumping laser beam is used. To effectively repump all $3D_{3/2}$ hyperfine states, this beam is power broadened with approximately 1500 saturation intensities. Finally, a 854 nm deshelving beam is used to ensure that ions shelved in the $3D_{5/2}$ level can also be cooled.

As explored in section 2.3.2, the ion state may be in one of many hyperfine states, within the ground level these transitions are spread over a ~ 3 GHz frequency range. The majority of this spread is caused by the splitting between the $F = 3$ and $F = 4$ hyperfine manifolds (see figure 2.1). We therefore employ an Electro-Optic-Modulator (EOM) with the 397 nm laser to address both the $F = 3$ and $F = 4$ manifolds. Further, we simplify the cooling dynamics by choosing an elliptically polarised 397 nm Doppler cooling beam, such that the σ^- polarisation component of the beam is minimised. This choice of polarisation biases the ion state to be in high m quantum number states, reducing the number of states the ion is likely to occupy.

The effectiveness of laser cooling depends on a combination of the laser detuning from the ion transition and the ion temperature. In the limit of hot ion motion, Doppler broadening makes a larger laser detuning desirable for effective cooling. To aid the cooling effectiveness of the initially hot ions produced from the photo-ionisation process, and to help recool ions after background



(a)



(b)

Figure 2.4: **Normal mode orientations** - Diagram showing the orientation of the normal modes with respect to the trap structure and 397 nm Doppler cooling beam. a) Vertical cross-section showing the 15° tilt of the radial normal modes. Due to the tilt, the motion of both modes have a projection onto the horizontal Doppler cooling beam direction. b) Vertical cross-section showing the 45° angle of the Doppler cooling beam to the trap and mode axes.

gas collisions, we therefore employ both a 397 nm ‘Doppler beam’ and a lower power 397 nm ‘detuned Doppler beam’ which is -400 MHz detuned from the ‘Doppler beam’.

2.4.2 Ground state cooling

The temperature achieved with the Doppler cooling procedure described in section 2.4.1 is too high to facilitate high fidelity two qubit logic operations. A standard technique for cooling trapped ions to their motional ground state is resolved, first-order Raman sideband cooling. However, ground state cooling using resolved Raman sideband cooling also requires a suitably low starting temperature. The standard Doppler cooling procedure outlined above does not achieve sufficiently low ion temperatures to allow resolved Raman sideband cooling. To enhance the effectiveness of the laser cooling, we adopt a dark resonance cooling scheme where we tune the 397 nm and 866 nm laser frequencies to coherently eliminate some of the $4P_{1/2}$ population. [65, 66] The dark resonance produced by the coherent $4P_{1/2}$ elimination, has significantly sharper dependence of the 397 nm fluorescence on the laser frequency than standard Doppler cooling [67]. The increased fluorescence gradient and reduced photon scatter, result in a reduced laser cooling limit allowing us to achieve motional occupation numbers of ~ 2 quanta on the single ion inplane radial mode. This cooling step is then followed by the standard technique of resolved Raman sideband cooling to the motional ground state [12, 65].

2.5 Qubit state preparation

To initialise our logical qubits, we need to be able to prepare a cold ion in a known internal state. We implement this via optical pumping with a σ^+ polarised 397 nm laser beam. To prepare the ion no matter what initial hy-




Figure 2.5: **Optical state-preparation** - Diagram showing the energy levels and optical processes used for optical pumping to the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state. The states in the $4S_{1/2}$ electronic ground level are grouped into two energy bands with $F = 3$ and $F = 4$ quantum number respectively. Two σ^+ polarised 397 nm frequencies are used to excite both groups of $4S_{1/2}$ states to the $4P_{1/2}$ level. Spontaneous emission from the $4P_{1/2}$ returns us to the ground state with statistically higher m quantum number or is repumped using the 866 nm laser. Figure reproduced from Harty [68]

perfine state it started in, an EOM is used to produce spectral components which separately address the $F = 3$ and $F = 4$ manifolds within the $4S_{1/2}$ level (see figure 2.5). σ^+ polarised light propagating along the magnetic field axis only drives $\Delta m = +1$ transitions from the $4S_{1/2}$ to $4P_{1/2}$ level. The excited $4P_{1/2}$ state spontaneously decays with $\Delta m = 0, \pm 1$, giving a stochastic increase in m . Once the ion is in the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state this process self-terminates. The ion can no longer be excited by the σ^+ polarised 397 nm beam as there is no $4P_{1/2}$, $m = +5$ state. Finally, as the optical pumping process depends on excitation to the $4P_{1/2}$ level, it is possible for leakage to the $3D_{3/2}$ states to occur. This leakage is mitigated using the same 866 nm re-pumping beam as in the Doppler cooling scheme [65, 68].

2.5.1 Microwave enhanced state preparation

The fidelity of preparing the desired state with the procedure outlined above is limited by the polarisation purity of the 397 nm laser beam used for optical pumping. Any σ^- or π polarisation impurities will scatter the ion out of the desired $4S_{1/2}, |F = 4, m = +4\rangle$ state, setting up an imperfect equilibrium. Appendix A details experimental methods for optimising the polarisation of the 397 nm optical pumping beam. However, it remains experimentally challenging to produce high purity σ^+ polarised light. We therefore follow the purely optical state preparation detailed above with a microwave enhanced state preparation procedure. This reduces our sensitivity to the optical polarisation purity [68].

As our desired state is within the $F = 4$ manifold, we can note that transitions from this state are several GHz detuned from the transitions driven from the $F = 3$ manifold. Disabling the EOM which produces the frequency splitting in our 397 nm beam, allows frequency selective optical pumping of the $F = 3$ level (see figure 2.6a). After a few scattering cycles the ion is likely to have decayed to the $F = 4$ manifold and be in a higher m state. We may now use microwave transfer pulses to map the undesired $F = 4$ states into the $F = 3$ manifold. Cycling through the $F = 3$ addressed optical pumping and microwave mapping allows for high fidelity state preparation with reduced sensitivity to the 397 nm beam polarisation.

In practice, after the initial optical state preparation it is unlikely for the ion to occupy states with $m < 2$. We therefore only map the $|F = 4, m = +3\rangle$ and $|F = 4, m = +2\rangle$ states to the $F = 3$ manifold during the microwave enhancement step (see figure 2.6b). Equation 2.1 summarises the state preparation pulses sequence with n microwave enhancement cycles.




Figure 2.6: **Microwave enhanced state preparation** - a: Diagram showing the optically addressed energy levels during microwave enhanced state preparation. By disabling the EOM, the σ^+ polarised 397 nm beam only has a single frequency component which is spectrally addressed to the $4S_{1/2}$, $F = 3$ states. The $F = 3$ states are optically pumped excitation to and decay from the $4P_{1/2}$ level. Meanwhile, the $F = 4$ states, are not interacted with. b: diagram indicating the $4S_{1/2}$ hyperfine transitions used for microwave enhanced state preparation. Microwave π pulses are driven on the $|F = 4, m = +3\rangle \rightarrow |F = 3, m = +3\rangle$ and $|F = 4, m = +2\rangle \rightarrow |F = 3, m = +2\rangle$ transitions. Figures reproduced from Harty [68].

$$397 \text{ nm } \sigma^+ \text{ (EOM ON)}, \left\{ 397 \text{ nm } \sigma^+ \text{ (EOM OFF)}, \mu W \pi\text{-pulses} \right\}^n, 397 \text{ nm } \sigma^+ \text{ (EOM OFF)} \quad (2.1)$$

2.5.2 Microwave transfer to the clock qubit

The state preparation procedure discussed so far prepares the ion in the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state. However, we wish to perform our logic operations on the clock transition between the $|F = 3, m = +1\rangle$ and $|F = 4, m = +1\rangle$ states (see section 2.3.3). To initialise the ion within our chosen qubit, we consecutively apply calibrated microwave π -pulses on the transitions ‘a’, ‘b’,

and ‘c’ (see figure 2.3). After these microwave transfer pulses, the ion state is initialised in the $4S_{1/2}, |F = 3, m = +1\rangle$ state.

2.6 Qubit readout

To assess our implemented quantum logic operations we also wish to readout the qubit state. We achieve this via optical shelving and state dependent fluorescence detection.

State dependent fluorescence detection is done by repeatedly exciting one of the qubit states to a short lived electronic state. As this excited state spontaneously decays, the emitted photons may be detected via a camera or photo multiplier tube. As this scattering only occurs if the ion was in a specific qubit state, this process converts the qubit state into the presence or absence of these scattered photons.

For Ca^+ qubits it is common to achieve this state selective fluorescence by incoherently mapping one of the qubit states to the long lived $3D_{5/2}$ level, also referred to as the ‘shelf’. The other qubit state is left within the $4S_{1/2}$ ground state. Similarly to the Doppler cooling procedure, a 397 nm beam is used to excite ions within the $4S_{1/2}$ level to the short-lived $4P_{1/2}$ energy level. If the ion started within the electronic ground state, 397 nm emission from spontaneous decay of the $4P_{1/2}$ level may then be used to infer the original qubit state. However, as discussed in section 2.4.1 the $4P_{1/2}$ level can also decay to the lower lying $3D_{3/2}$ level. A 866 nm repumping beam is used to prevent the fluorescence cycle being terminated by this decay [69, 70, 71].

2.6.1 Shelving procedure

As outlined above, part of the qubit readout procedure involves incoherently mapping one of the qubit states into the $3D_{5/2}$ level. We do this by excit-

ing the 393 nm E1 transition from the $4S_{1/2}$ electronic ground state to the $4P_{3/2}$ level (see figure 2.1). Unlike the $4P_{1/2}$ level, which is used in Doppler cooling, the $4P_{3/2}$ level can also decay to the long lived $3D_{5/2}$ states. There are three main challenges with this approach. Firstly, to be suitable for qubit readout, the shelving procedure must be state selective as shelving the wrong qubit state results in readout errors. Fortunately, the ground level hyperfine structure of $^{43}\text{Ca}^+$ at 288 G is spread over a much larger frequency range than the natural linewidth of the 393 nm $4S_{1/2} \rightarrow 4P_{3/2}$ transition. We can therefore use frequency selectivity to only shelve the desired qubit state. Secondly, though we desire the ion to decay from the $4P_{3/2}$ to the $3D_{5/2}$ state, the $4P_{3/2}$ branching ratios indicate that the majority of $4P_{3/2}$ decays lead back to the $4S_{1/2}$ ground level. Lest that the aforementioned frequency selectivity terminate the shelving process, spontaneous decay to the ground state must return us to the same hyperfine state from which the shelving process was initiated. We achieve this by addressing the $4S_{1/2}$, $|F = 4, m = +4\rangle \rightarrow 4P_{3/2}$, $|F = 5, m = +5\rangle$ transition with a σ^+ polarised beam (see figure 2.7). Finally, there is also an undesired 850 nm $4P_{3/2} \rightarrow 3D_{3/2}$ decay. Comparing branching ratios between the desired $4P_{3/2} \rightarrow 3D_{5/2}$ and undesired $4P_{3/2} \rightarrow 3D_{3/2}$ decay, we note that approximately 10% of the decays to a D-state are to the $D_{3/2}$ level. Unlike the $D_{5/2}$ level, the $D_{3/2}$ level is part of the fluorescence detection cycle. This decay therefore results in a $\sim 10\%$ shelving error probability if left unmitigated.

High fidelity shelving of $^{43}\text{Ca}^+$ at lower magnetic fields is commonly implemented using a pulsed 850 nm repumping scheme [68, 69, 70]. Though we are at approximately twice the previously demonstrated magnetic field, we adopt the same approach, which is outlined below.

As the $P_{5/2}$ state that we populate is the $|F = 5, m = +5\rangle$ state, there are

only three possible decays to the $3D_{3/2}$ level. The $3D_{3/2}$ states these result in are the $|F = 5, m = +5\rangle$, $|F = 5, m = +4\rangle$, and $|F = 4, m = +4\rangle$ states. We repump the undesired $3D_{3/2}$ states to the $4P_{3/2}$, $|F = 5, m = +5\rangle$ state, using 850 nm laser light which is polarisation, and frequency addressed to each of the transitions. To avoid creating dark resonances with the 393 nm laser, pulses of the repumping beams and 393 nm laser are alternated. Further, to avoid driving undesired transitions from the $3D_{3/2}$, $|F = 4, m = +4\rangle$ and $|F = 5, m = +4\rangle$ states, the beams addressing these transitions are σ^+ polarised and pulsed on before the third transition is repumped. The π -transition from the $|F = 5, m = +5\rangle$ is then driven using an elliptically polarised beam, minimising its σ^- polarisation content. Repeating this 393 nm - 850 nm cycle as shown in equation 2.2, allows for shelving fidelities on the 10^{-4} level.

$$\left\{ 393 \text{ nm } \sigma^+, 850 \text{ nm } \sigma^+, 850 \text{ nm } \pi \& \sigma^+ \right\}^n, 393 \text{ nm } \sigma^+ \quad (2.2)$$

In the above equation pulses are labelled via their wavelength and polarisation, n represents the number of repumping cycles. Appendix A details experimental methods for optimising the polarisations and frequencies of the 850 nm repumping beams.

2.6.2 Microwave transfer from clock qubit

In our discussion of the shelving procedure (section 2.6.1), it became apparent that our chosen qubit states $4S_{1/2}$, $|F = 3, m = +1\rangle$, and $4S_{1/2}$, $|F = 4, m = +1\rangle$ are not suitable for state-selective optical shelving. Instead, state selective optical shelving is performed from the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state. Similarly to our state preparation procedure, we therefore coherently transfer the $4S_{1/2}$, $|F = 3, m = +1\rangle$ state to the $|F = 4, m = +4\rangle$ state with a microwave pulse sequence. Naively, one might simply apply a sequence of π -pulses on the

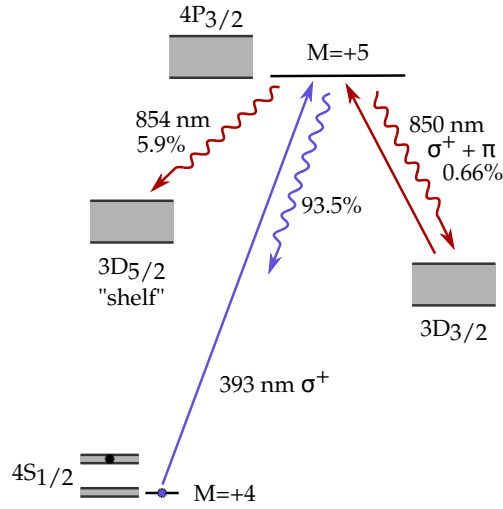


Figure 2.7: **Readout transitions** - Energy level diagram showing the optical shelving procedure. A spectrally addressed 393 nm σ^+ beam selectively excites the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state to the $4P_{3/2}$, $|F = 5, m = +5\rangle$ state. The $4P_{3/2}$, $|F = 5, m = +5\rangle$ state may spontaneously decay to the initial $4S_{1/2}$, $|F = 4, m = +4\rangle$ state, the desired $3D_{5/2}$ ‘shelf’ level or the spectating $3D_{3/2}$ level. Decays back to the $4S_{1/2}$ level are re-excited by the 393 nm σ^+ beam. Undesired decays to the spectating $3D_{3/2}$ level are re-injected into the shelving cycle using 850 nm pulses which are spectrally addressed to repump all possible $3D_{3/2}$ hyperfine states back to the shelving cycle via excitation to the $4P_{3/2}$, $|F = 5, m = +5\rangle$ state. Figure reproduced from Harty [68].

transitions ‘c’, ‘b’, and ‘a’ (see figure 2.3). However, this approach has two significant shortcomings [68].

Firstly, our selective shelving scheme relies on frequency addressing of the desired 393 nm transition. However, the rich hyperfine structure of $^{43}\text{Ca}^+$ gives rise to a multitude of spectator transitions. These transitions allow shelving of the ‘bright’ $|F = 4, m = +1\rangle$ qubit state, which should not be shelved. The closest 393 nm spectator transition detuning from the ‘bright’ $|F = 4, m = +1\rangle$ clock qubit state is only 48 MHz. This detuning is comparable to the 23 MHz natural linewidth of the 393 nm $4S_{1/2} \rightarrow 4P_{3/2}$ transition and results in poor shelving selectivity between the ‘bright’ $|F = 4, m = +1\rangle$ and ‘dark’ $|F = 4, m = +4\rangle$ hyperfine states. Spectator transitions for states

within the $F = 3$ manifold are generally much further detuned. We therefore follow the microwave transfer of the $|F = 3, m = +1\rangle$ qubit state to the optically shelved $|F = 4, m = +4\rangle$ state with a microwave π -pulse coherently mapping the $|F = 4, m = +1\rangle$ state onto the $|F = 3, m = +1\rangle$ state. This results in a much larger nearest spectator transition detuning of 2.8 GHz

Secondly, the σ polarised transitions within the $4S_{1/2}$ level come in near degenerate pairs. Applying microwave pulses to one of the paired σ polarised transitions drives significant off-resonant transitions on the near degenerate transition. Consequently, using a π pulse to transfer the $|F = 3, m = +1\rangle$ state via transition ‘c’ also drives significant dynamics from the ‘bright’ state via the 84 kHz detuned transition ‘d’. After applying the similarly near degenerate ‘b’ pulse this results in some of the ‘bright’ $|F = 4, m = +1\rangle$ state being mapped into a superposition of states which are readily shelved by the 393 nm laser pulse. To avoid this process the microwave transfer pulse on the ‘c’ transition is split into a Ramsey sequence consisting of $\pi/2$ pulses separated by a time delay. The $\pi/2$ pulses are applied with equal phase resonantly to the ‘c’ transition, resulting in a complete population transfer. On the ‘d’ transition, the two pulses drive off-resonant Rabi oscillation. The resulting excitation is suppressed by empirically tuning the inter-pulse delay such that the two off-resonant oscillations coherently cancel.

These considerations indicate that the naive sequence employing a sequence of π pulses on the transitions ‘c’, ‘b’, and ‘a’ is not suitable for achieving high shelving selectivity. We therefore employ the more suitable sequence of transfer pulses on the transitions ‘c’, ‘q’, ‘b’, ‘a’ (see figure 2.3). Crucially, the ‘c’ transfer pulse is a Ramsey sequence tuned to suppress off-resonant excitation of the ‘bright’ $|F = 4, m = +1\rangle$ qubit state.

2.6.3 Threshold readout

Having established a procedure for mapping states of the $^{43}\text{Ca}^+$, 288 G clock qubit into a fluorescing and dark state, we also require a reliable way of detecting the presence or absence of fluorescence. Fortunately, a variety of fluorescence detection techniques using PMTs or CCD cameras are available for effectively detecting the 397 nm fluorescence of Ca^+ ions [71]. Many of the more advanced detection techniques seek to decrease the number of photons required to make an accurate determination on the presence or absence of fluorescence. Crucially, though the $3\text{D}_{5/2}$ shelf level is ‘long lived’, its 1.2 s lifetime imposes a limit on the dark state detection fidelity based on the time taken to detect sufficient photons to make a determination on the presence or absence of fluorescence. As we will discuss in section 4.3, we employ a high numerical aperture imaging system. The high photon collection efficiency of this imaging system allows us to achieve sufficient detection fidelity with a PMT and simple thresholding technique.

In the thresholding fluorescence detection technique, ions are determined to be bright or dark based on the total number of photons detected in a fixed time interval. The exact time interval and thresholds are calculated numerically under the assumption of a Poisson distributed background detection rate and a non-poissonian count rate associated with a dark ion that arises from accounting for the finite $3\text{D}_{5/2}$ shelf lifetime. A detailed derivation of this formula for a single ion is presented by Langer [72]. Readout of multiple ions may then be done by spatially separating and individually querying the fluorescence of each ion using the ion transport techniques presented in chapter 5.

Due to control system limitations, the work in this thesis uses simpler two-ion thresholding technique, unless otherwise stated. When using the two

ion thresholding technique, two photon count thresholds are determined for a fixed detection time. If the count rate is below both thresholds, both ions are determined to be dark. If the count rate is above both thresholds both ions are determined to be bright. If the photon count is between the thresholds one of the ions is determined to be dark and the other bright. In the latter case it is not determined which of the two ions is bright or dark. However, as we will see in section 3.2.2, “which ion” information is not required to characterise the global microwave operations presented in this work.

As the two-ion thresholding error was not a limiting factor, a sensible detection time was determined by treating the “two bright” event as an ion with half the $3D_{5/2}$ shelf lifetime with a poissonian background corresponding to the expected count rate for a single fluorescing ion. As ion fluorescence produces many more counts than background scatter, the two-ion thresholding error is dominated by the distinguishability of the single and two bright ion events, justifying this approximation. As we are now treating two-ion thresholding like a single ion thresholding problem, we simply use the method presented by Langer [72] to determine a suitable detection time and our upper detection threshold. To determine the lower detection threshold, we again employ the expressions found by Langer [72] to find an optimum threshold count distinguishing a single and two dark ions.

Chapter 3

Two-Qubit Gate Theory

We now turn our attention to the qubit-qubit interactions necessary for implementing multi-qubit quantum logic gates. The two qubit gates commonly used for trapped ions exploit a spin-dependent force acting on the ions. In our case this spin-dependent force is provided by the amplitude gradient of our microwave field. As the name suggests spin-dependent forces exert a force on the ion that depends on the internal spin state of the ion. Consequently, the ion motion driven by spin-dependent forces is entangled with the ions internal state.

For multi-ion crystals, the coulomb repulsion results in the motion of all ions being coupled. The motional state of the ions must therefore be thought of as an excitation of the collective ion motion, which is quantised in the normal modes incorporating all ions within the crystal. As normal modes are shared between ions, entanglement of this collective motion with the ion internal states via a spin-dependent force may be exploited to distribute entanglement between the otherwise negligibly interacting internal ion states [11, 73].

3.1 Microwave spin-dependent force

For magnetic dipole transitions, spin-dependent forces are generated by magnetic field gradients. When using plane wave radiation, the magnetic field gradient is inversely proportional to the wavelength. Consequently, the spin-dependent force produced by ~ 10 cm microwave radiation is 10^5 times weaker than the spin-dependent force produced by 1000 nm laser light. The far-field radiation approach employed with laser driven gates is therefore not practical when using microwave radiation [12]. Fortunately, a variety of techniques for generating stronger spin-dependent forces have been established. These include the application of static magnetic field gradients [43], and oscillating magnetic field gradients near the ion motional frequency [21]. However, we adopt the near-field microwave amplitude gradient approach initially presented by Ospelkaus [50]. In the near-field regime the ion is placed in proximity to current carrying electrodes. If the ion-electrode distance is kept small compared to the microwave wavelength, the magnetic field gradients are determined by the geometry of the current carrying electrodes, rather than the far-field microwave radiation wavelength. Consequently, much larger microwave magnetic field gradients can be engineered while maintaining practical microwave powers.

3.1.1 Hamiltonian of trapped ion microwave interaction

We are interested in the dynamics that near-field microwave radiation drives on the ion-internal and collective ion-motional states. For simplicity, we neglect spectator transitions within the hyperfine structure of the ion. Further, we consider the case where the detuning Δ of the applied microwaves from the qubit transition frequency ω_{qubit} is near the frequency of a single normal mode of the ion crystal ω_{motion} . For an N ion crystal, we therefore simplify our model

to N qubits that are coupled by a single harmonic oscillator, representing the frequency selected normal mode. Within the rotating wave approximation, the dynamics of ion n in the presence of a $\cos((\omega_{qubit} + \Delta)t + \phi)$ time-varying microwave field are described by the interaction picture Hamiltonian presented in equation 3.1 [50, 74].

$$\begin{aligned} \tilde{H}_{I,n}(t) = & \frac{1}{2}\hbar\Omega_{Q,n}e^{-i(\Delta t - \phi)}\hat{\sigma}_{+,n} + \frac{1}{2}\hbar\Omega_{z,n}e^{-i((\omega_{qubit} + \Delta)t - \phi)}\hat{\sigma}_{z,n} \\ & + \frac{1}{2}\hbar\Omega'_{Q,n}e^{-i(\Delta t - \phi)}\hat{\sigma}_{+,n} \left(\hat{a}e^{-it\omega_{motion}} + \hat{a}^\dagger e^{it\omega_{motion}} \right) \\ & + \frac{1}{2}\hbar\Omega'_{z,n}e^{-i((\omega_{qubit} + \Delta)t - \phi)}\hat{\sigma}_{z,n} \left(\hat{a}e^{-it\omega_{motion}} + \hat{a}^\dagger e^{it\omega_{motion}} \right) \\ & + \text{h.c.} \end{aligned} \quad (3.1)$$

In this equation, qubit dynamics are encoded via the $\hat{\sigma}_{+,n}$ and $\hat{\sigma}_{z,n}$ operators. Respectively, these are the spin raising and Pauli operators of the n^{th} qubit. Motional interactions of the normal mode are presented using the creation operator \hat{a} . Finally, the Ω terms encode the strength of these interaction terms.

The first interaction term represents Rabi oscillations driven by the microwave magnetic field at the n^{th} ion location. The interaction strength of this term is given by equation 3.2.

$$\Omega_{Q,n} = -\frac{1}{\hbar}\bar{\mu}_Q B_{Q,n} \quad (3.2)$$

where $B_{Q,n}$ indicates the amplitude of the magnetic field component matching the qubit transition polarisation evaluated at the n^{th} ion location, meanwhile $\bar{\mu}_Q$ represents the magnetic dipole matrix element of the qubit transition. The second interaction term represents a time-varying modulation of the qubit transition frequency, caused by the instantaneous magnetic field experienced by the ion. Similarly the fourth term accounts for additional field variation

associated with excitation of the normal mode. As discussed in section 2.3.3, we are using a clock qubit, where the qubit transition frequency is stationary with respect to the external magnetic field. In our case, the first order frequency dependence represented by the second and fourth term is therefore absent (see equation 3.3) [74]. Consequently, we disregard these terms from here on.

$$\Omega_{z,n} = \Omega'_{z,n} = 0 \quad (3.3)$$

Finally, we note that the third term represents a coupling between the qubit states and the collective ion motion and will be leveraged in section 3.2 to implement multi-qubit logic operations. The associated coupling strength is presented in equation 3.4.

$$\Omega'_{Q,n} = -\frac{1}{\hbar} \bar{\mu}_Q q_0 b_n \vec{q} \cdot \left(\vec{\nabla} B_Q \right)_n \quad (3.4)$$

As before, $\bar{\mu}_Q$ represents the magnetic dipole matrix element of the qubit transition. The notation $\left(\vec{\nabla} B_Q \right)_n$ indicates “grad of the amplitude of the spatially-varying, qubit-polarised magnetic field component”, evaluated at the position of the n^{th} ion. The vector \vec{q} is the unit vector in the oscillation direction of the selected normal mode. Different ions participate with different amplitudes and phases in a given normal mode, this is conveniently represented by the normalised eigenvectors \vec{b} where each element encodes the relative phase and participation amplitude of an ion. b_n is the element of the eigenvector of our selected normal mode corresponding to the n^{th} ion. Finally, q_0 is the spatial extent of the ground state wave function of the ion associated with the normal mode. q_0 is determined solely by the normal mode frequency and ion mass as shown in equation 3.5 [74].

$$q_0 = \sqrt{\frac{\hbar}{2m\omega_{motion}}} \quad (3.5)$$

3.1.2 Effective Lamb-Dicke parameter

From our discussion of the microwave interaction Hamiltonian, the microwave magnetic field amplitude experienced by the ion drives Rabi oscillations on the qubit frequency. Meanwhile, spin-dependent motional interactions are driven by the magnetic field amplitude gradient. Proceeding analogously to optically driven qubit operations, we define the effective Lamb-Dicke parameter for the microwave interaction as the ratio between these interaction strengths as shown in equation 3.6.

$$\tilde{\eta} \equiv \left| \frac{\Omega'_{Q,n}}{\Omega_{Q,n}} \right| = \left| \frac{q_0 b_n \vec{q} \cdot (\vec{\nabla} B_Q)_n}{B_{Q,n}} \right| \quad (3.6)$$

Typical values of $\tilde{\eta}$ achieved when using near field microwaves are of the order 10^{-3} . This is roughly 5 orders of magnitude larger than for far field microwave radiation. However, it is worth noting that these effective microwave Lamb-Dicke parameters are much smaller than $\eta \approx 0.1$ Lamb-Dicke parameters which are readily achievable for optically driven interactions. Consequently, when comparing microwave and optically generated spin-dependent forces, we note that much stronger microwave interactions must be applied to achieve comparable spin-dependent forces.

3.2 The Mølmer-Sørensen gate

For two-qubit logic operations we implement a Mølmer-Sørensen style geometric phase gate (MS gate) [73]. Unlike earlier multi-qubit gate approaches [11], geometric phase gates operate by driving closed phase-space trajectories on the collective motion of the ion crystal. As the employed driving force is spin-dependent, the driven trajectories depend on the initial qubit state. Thus, these qubit states have different phase-space trajectories which accu-

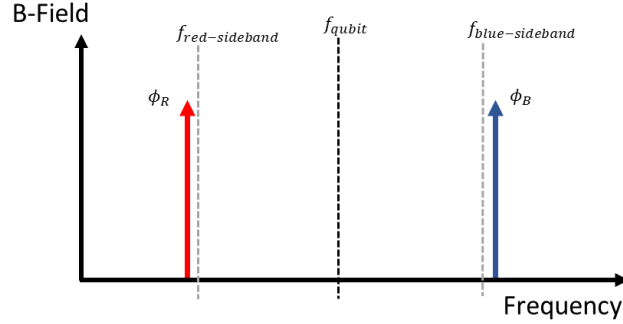


Figure 3.1: **MS gate tones** - The MS gate employs two microwave tones. The red and blue detuned microwave sideband tones are chosen as for a normal MS gate. As the MS gate employs two microwave tones, it has two drive phase degrees of freedom. The phases of the red and blue sideband drives (ϕ_R and ϕ_B) determine the average and difference phases ($\bar{\phi}$ and $\Delta\phi$) of the MS gate.

ulate different geometric Berry phases [75]. As the area enclosed in the phase space trajectories is insensitive to the initial motional state, geometric phase type gates are widely used to implement high fidelity multi ion logic operations [76].

In the following, we briefly review key aspects of the MS gate. More detailed treatments are presented by Mølmer [73] and Sepiol [74].

3.2.1 The ideal Mølmer-Sørensen gate

The standard MS gate consists of two equal amplitude microwave tones. Respectively, the two microwave tones are detuned by $\pm\delta$ from the blue and red motional sidebands of the qubit transition, see figure 3.1. From our previous discussion of the microwave interaction Hamiltonian in section 3.1.1, and discarding terms oscillating at $2\omega_{motion}$, we can write down the Hamiltonian

for motional interaction of the n^{th} ion in the presence of such a red and blue sideband tone, as shown in equation 3.8 [74].

$$\begin{aligned}\tilde{H}_{MS,n}(t) &= \tilde{H}_{B,n}(t) + \tilde{H}_{R,n}(t) \\ &= \frac{1}{2} \hbar \Omega'_{Q,n} \left(\hat{\sigma}_{+,n} e^{i\bar{\phi}} + \hat{\sigma}_{-,n} e^{-i\bar{\phi}} \right) \left(\hat{a} e^{-i(\delta t - \frac{\Delta\phi}{2})} + \hat{a}^\dagger e^{i(\delta t - \frac{\Delta\phi}{2})} \right) \quad (3.7)\end{aligned}$$

$$= \frac{1}{2} \hbar \Omega'_{Q,n} \hat{\sigma}_{\bar{\phi},n} \left(\hat{a} e^{-i(\delta t - \frac{\Delta\phi}{2})} + \hat{a}^\dagger e^{i(\delta t - \frac{\Delta\phi}{2})} \right) \quad (3.8)$$

In the above, we have rewritten the phases of the red (ϕ_R) and blue (ϕ_B) sideband tones as:

$$\Delta\phi \equiv \phi_B - \phi_R \quad (3.9)$$

$$\bar{\phi} \equiv \frac{1}{2} (\phi_B + \phi_R) \quad (3.10)$$

For convenience, we also defined $\hat{\sigma}_{\bar{\phi},n}$ in terms of the x and y Pauli matrices of the n^{th} qubit.

$$\hat{\sigma}_{\bar{\phi},n} \equiv \hat{\sigma}_{x,n} \cos(\bar{\phi}) - \hat{\sigma}_{y,n} \sin(\bar{\phi}) \quad (3.11)$$

In the case of two trapped ions experiencing identical magnetic field gradients, we have $\Omega'_Q = |\Omega'_{Q,n}|$ and our the MS interaction Hamiltonian may be expressed in terms of a collective spin operator $\hat{S}_{\bar{\phi}}$.

$$\tilde{H}_{MS}(t) = \hbar \Omega'_Q \hat{S}_{\bar{\phi}} \left(\hat{a} e^{-i(\delta t - \frac{\Delta\phi}{2})} + \hat{a}^\dagger e^{i(\delta t - \frac{\Delta\phi}{2})} \right) \quad (3.12)$$

$$\hat{S}_{\bar{\phi}} = \frac{1}{2} \sum_{n=1}^2 \text{sign}(\Omega'_{Q,n}) \hat{\sigma}_{\bar{\phi},n} \quad (3.13)$$

Inspecting equation 3.12, it is apparent that the $+1$ and -1 eigenstates of the $\hat{S}_{\bar{\phi}}$ operator undergo equal and opposite phase space trajectories. Further,




Figure 3.2: **Phase space trajectories** - Schematic indicating the phase space trajectories taken by different eigenstates of $\hat{S}_{\bar{\phi}}$ for gates performed on the radial rocking mode. When the driving force is detuned from the motional resonance, the $|+-\rangle$ and $|-\rangle$ state trajectories form a loop returning to the origin. In doing so the $|+-\rangle$ and $|-\rangle$ states, acquire geometric phase proportional to the enclosed area. Figure reproduced from Sepiol [74].

the two zero eigenstates of the operator remain stationary in phase space. The phase space trajectories driven on the $+1$ and -1 eigenstates depend on the sideband detuning δ , the entanglement driving strength Ω'_Q and the sideband difference phase $\Delta\phi$, see figure 3.2.

As noted by Mølmer [73], a maximally entangling two qubit gate can be implemented by choosing a gate interaction duration t_g and sideband detuning δ_g which satisfies relations 3.14 and 3.15.

$$t_g = \frac{2\pi K}{\delta_g} = \frac{\pi\sqrt{K}}{\Omega'_Q} \quad (3.14)$$

$$\delta_g = 2\Omega'_Q\sqrt{K} \quad (3.15)$$

In the above expression K is a positive integer representing the chosen number of closed loops within the phase space trajectories of the $+1$ and -1 eigenstates over the course of the gate interaction.

For reasons presented in sections 3.2.3 and 4.1, the two-qubit gate demonstration within this thesis were carried out on the radial rocking mode. In the case of the rocking mode, we have $\hat{S}_{\bar{\phi}} = \frac{1}{2} (\hat{\sigma}_{\bar{\phi},1} - \hat{\sigma}_{\bar{\phi},2})$. Working in the $\hat{\sigma}_{\bar{\phi},n}$ eigenbasis, the truth table for such a gate is as follows [74]:

$$|++\rangle \mapsto |++\rangle \quad (3.16)$$

$$|+-\rangle \mapsto -i |+-\rangle \quad (3.17)$$

$$|-+\rangle \mapsto -i |+-\rangle \quad (3.18)$$

$$|--\rangle \mapsto |--\rangle \quad (3.19)$$

This truth table implements a two-qubit phase gate in the $\hat{\sigma}_{\bar{\phi},n}$ eigenbasis. Combined with single qubit rotations, this gate makes up a ‘universal’ set of gates, satisfying the 4th DiVincenzo criterion [5].

For convenience, we also present the truth table within the logical basis, which is shown below [74]:

$$|11\rangle \mapsto \frac{1}{\sqrt{2}} e^{-i\pi/4} (|11\rangle + i e^{-2i\bar{\phi}} |00\rangle) \quad (3.20)$$

$$|10\rangle \mapsto \frac{1}{\sqrt{2}} e^{-i\pi/4} (|10\rangle + i |01\rangle) \quad (3.21)$$

$$|01\rangle \mapsto \frac{1}{\sqrt{2}} e^{i\pi/4} (|10\rangle - i |01\rangle) \quad (3.22)$$

$$|00\rangle \mapsto \frac{1}{\sqrt{2}} e^{i\pi/4} (e^{2i\bar{\phi}} |11\rangle - i |00\rangle) \quad (3.23)$$

Notably, this gate maps separable two-qubit states into the maximally entangled Bell states.

3.2.2 Gate fidelity

In the real world, a variety of sources can lead to errors when implementing the MS gate we outlined. Such errors are a discrepancy to the desired truth table. This discrepancy is commonly characterised via the fidelity \mathcal{F} of implementing a desired pure state $|\psi\rangle$.

$$\mathcal{F} = \langle \psi | \hat{\rho} | \psi \rangle \quad (3.24)$$

In equation 3.24 $\hat{\rho}$ indicates the density matrix representing the experimentally produced state. From the fidelity, the error probability ϵ of not producing the state $|\psi\rangle$ is determined as $\epsilon \equiv 1 - \mathcal{F}$. Notably, the error probability need not be the same across all target states $|\psi\rangle$. Therefore, to fully characterise the error of an implemented gate, process tomography must be used [77]. Alternatively, interleaved randomised benchmarking also allows the average error $\bar{\epsilon}$ to be determined [78, 79, 80]. However, both these procedures require the use of single qubit addressing, which is still being developed within our experimental apparatus [81]. In the absence of single qubit addressing, symmetric subspace randomised benchmarking may be used to attain average errors over a subspace of the truth table [82, 83]. None-the-less, in this work, we adopt the parity analysis technique. The parity analysis technique, provides an estimate of the gate fidelity by measuring the fidelity of preparing the expected Bell state, shown in equation 3.23.

In parity analysis the fidelity of the produced bell state is attained by combining multiple distinct measurements after applying the MS-gate to two qubits prepared in the $|00\rangle$ state. Firstly, the ‘population’ measurement simply performs a measurement of the qubit state after application of the MS gate. Secondly, the ‘parity’ of the produced state is measured. To measure the parity, the original MS-gate is followed by a tunable phase $(\pi/2)_\theta$ -pulse. This

additional pulse is referred to as the ‘analysis’ pulse. To establish the parity, qubit measurements with multiple analysis pulse phases must be performed. A detailed treatment of the Bell state fidelity derived parity analysis is presented by Sepiol [74]. Here, we only make a few qualitative observations and present the resulting Bell state fidelity expression.

Firstly, by inspecting equation 3.23 we see that the population measurement should yield no events where the ions are observed to be in different eigenstates. The probability of such events $P_{10} + P_{01}$ corresponds to Bell state errors.

Secondly, the measurement outcomes after further application of the analysis pulse depends on the relative phase of the analysis pulse to the average MS gate phase $\bar{\phi}$. For a perfect MS gate and analysis pulse, the probability of measuring different outcomes between the two qubits varies as:

$$P_{a,10}(\phi_a) + P_{a,01}(\phi_a) = \cos^2(\phi_a + \gamma) \quad (3.25)$$

where ϕ_a is the phase of the analysis pulse and γ is some phase offset. The notation $P_{a,qq}(\phi_a)$ indicates the probability of measuring the state $|qq\rangle$ after application of an analysis pulse with phase ϕ_a . In practice, γ depends on both the average phase $\bar{\phi}$ of the MS gate and Zeeman shifts experienced by the ion. Crucially, for a perfect MS gate and $\pi/2$ analysis pulse, $P_{a,10} + P_{a,01}$ has full contrast between zero and unity. Loss of this contrast is indicative of Bell state errors. The measurement outcomes determine the so-called ‘parity’ $\mathcal{P}(\phi_a)$ which is defined in equation 3.26

$$\mathcal{P}(\phi_a) \equiv 1 - 2(P_{a,10}(\phi_a) + P_{a,01}(\phi_a)) \quad (3.26)$$

Equation 3.27 shows the expression for the fidelity derived by Sepiol [74]. Conveniently, this fidelity may be rewritten in terms of the probability of

measuring different qubit states in the measurements carried out with and without an analysis pulse (see equation 3.29).

$$\mathcal{F} = \frac{1}{2}(P_{11} + P_{00}) + \frac{1}{4}(\max\{\mathcal{P}\} - \min\{\mathcal{P}\}) \quad (3.27)$$

$$= \frac{1}{2}(1 - (P_{10} + P_{01})) + \frac{1}{2}(\max\{P_{a,10} + P_{a,01}\} - \min\{P_{a,10} + P_{a,01}\}) \quad (3.28)$$

$$= \frac{1}{2}(1 + \max\{P_{a,10} + P_{a,01}\} - \min\{P_{a,10} + P_{a,01}\} - (P_{10} + P_{01})) \quad (3.29)$$

From equation 3.29 it becomes apparent that measurement statistics of the measured gate fidelity can be described by drawing from 3 independent binomial distributions which correspond to the probability of measuring a single bright ion in the population measurement and each of the two analysis phase extrema.

3.2.3 Common error sources

A variety of miscalibrations and environmental factors may degrade the fidelity of a MS style two qubit gate. With an eye on implementing error mitigation strategies, we discuss a selection of error sources below. A more complete presentation of gate errors is given in section 7.2.

3.2.3.1 Motional frequency miscalibrations

The ideal MS gate seeks to only drive dynamics caused by the geometric phase accumulated via the motional trajectories driven by the spin-dependent force. When we discussed the ideal MS gate in section 3.2.1, we noted that implementing the entangling operation requires not only accumulation of a specific geometric phase. Rather, we must accumulate the correct geometric phase and disentangle the qubit state from the collective ion motion. For a

given motional interaction strength Ω'_Q , these conditions required us to use a specific gate duration t_g (equation 3.14), and detuning of the microwave sideband tones from the motional sideband resonances of the qubit transition δ_g (equation 3.15). If equations 3.14 & 3.15 are not satisfied, the gate fidelity is reduced according to the actual sideband detuning δ and gate duration t , see equation 3.30 [65].

$$\mathcal{F}(t, \delta) = \frac{1}{8} \left(3 + e^{-16|\alpha(t, r)|^2(\bar{n}+1/2)} + 4e^{-4|\alpha(t, r)|^2(\bar{n}+1/2)} \sin [4\Phi(t, r)] \right) \quad (3.30)$$

$$\alpha(t, \delta) = \frac{\Omega'_Q}{2\delta} e^{i\Delta\phi} (e^{i\delta t} - 1) \quad (3.31)$$

$$\Phi(t, \delta) = \left(\frac{\Omega'_Q}{2\delta} \right)^2 (\delta t - \sin[\delta t]) \quad (3.32)$$

For an ideal gate we desire $\alpha(t, \delta) = 0$ and $\Phi(t, \delta) = \pi(1/8 + (K-1)/2)$. Notably, the sensitivity to deviations from $\alpha(t, \delta) = 0$ increases with larger motional occupation number. On the other hand, the sensitivity to motional detuning errors for a given interaction strength Ω'_Q may be reduced by choosing a slower, higher loop number K MS gate as this results in a larger ideal gate detuning δ_g , as per equation 3.15.

Experimentally, calibration accuracies of order $\Delta t \approx 1 \mu\text{s}$, and $\Delta\delta \approx 100 \text{ Hz}$ are reasonable. For a single loop, $t_g = 300 \mu\text{s}$ MS gate ($\Omega'_Q = 2\pi \times 2.5 \text{ kHz}$), such calibration inaccuracies are expected to result in a gate error probability of $\sim 0.5\%$.

3.2.3.2 Qubit frequency error

For an ideal MS gate, the microwave tones are symmetrically detuned from the qubit frequency, see figure 3.1. Mischaracterisation of the qubit frequency

therefore corresponds to detuning of the average microwave tone frequency from the qubit resonance frequency. Typical sources of such errors might be external magnetic field miscalibrations, AC-Stark shifts and AC-Zeeman shifts. As we are employing a clock qubit and do not use lasers to drive our logical interactions, our qubit frequency shifts are dominated by AC-Zeeman shifts.

Due to the limited magnetic field gradient produced by the trap electrode structure, driving two qubit gates also results in the ions experiencing significant microwave amplitude. Considering a single loop $t_g = 230 \mu\text{s}$ gate the AC-Zeeman shift arising from each sideband tone is expected to be approximately 1 MHz. Fortunately, cancellation between the Zeeman shifts produced by the two amplitude balanced microwave tones of a MS gate reduces the overall AC-Zeeman shift to $\sim 50 \text{ kHz}$. Both of these AC-Zeeman shifts are significantly larger than the qubit detuning tolerance of a $t_g = 230 \mu\text{s}$ MS gate. In fact, to obtain a $< 1\%$ gate error probability a $230 \mu\text{s}$ single loop gate only tolerates $< 200 \text{ Hz}$ Zeeman shift mischaracterisation. This corresponds to a common mode fractional sideband amplitude drift of 2×10^{-3} which is comparable to typical gain stabilities of microwave amplifiers. Limits on differential amplitude drifts between the two microwave tones are even tighter with a 1×10^{-4} fractional drift resulting in a 1% gate error for a single loop $t_g = 230 \mu\text{s}$ MS gate. Consequently, both differential and common mode amplitude stability of the microwave field is of significant concern for performing high fidelity two-qubit gates with near field microwaves.

3.2.3.3 Kerr coupling to spectator modes

During our discussion of the microwave motional interactions we neglected any interaction associated with spectator motional modes. One such interaction is

the Kerr non-linearity which represents a direct interaction between different normal modes of the ion crystal. In the following we focus on anharmonicity intrinsic to the Coulomb interaction between the ions within the crystal. In doing so, we neglect any additional anharmonicity which may be caused by the trap fields. As we perform our logic operations on the two-ion in-plane radial rocking mode, we need only consider interactions with the axial stretch mode and out-of-plane rocking mode as interactions with centre of mass modes are prevented by symmetry [84, 85]. From the treatment given by Nie [85], we can identify that the motional state of the spectating stretch and rocking modes cause a shift of the motional frequency of the inplane rocking mode. The form of the energy change ΔE associated with this interaction is shown in equation 3.33

$$\Delta E = \chi_s n_s n_g \quad (3.33)$$

Where we indicate the motional quanta of the addressed motional mode by n_g and the motional occupation number of the spectating mode by n_s . Finally, the χ_s indicates a constant characterising the strength of the non-linear interaction. Following the treatment of Nie [85], we can identify the formula for the Kerr effect due to the stretch and rocking modes as equations 3.34 & 3.35 respectively.

$$\chi_{stretch} = -\sqrt{3} \left[1 + \frac{3\omega_z^2}{2(4\omega_g^2 - 3\omega_z^2)} \right] \frac{\omega_z^{7/3}}{\omega_g} \left(\frac{2\hbar}{\alpha^2 M c^2} \right)^{1/3} \quad (3.34)$$

$$\chi_{rocking} = \frac{3}{4} \frac{\omega_z^{10/3}}{\omega_g \omega_{oop-rocking}} \left(\frac{2\hbar}{\alpha^2 M c^2} \right)^{1/3} \quad (3.35)$$

Here we take ω_z , $\omega_{oop-gate}$, and ω_g to indicate the mode frequency of the axial centre of mass, out-of-plane rocking mode and the desired gate rocking mode respectively. Further, the constants α , \hbar , c , and M indicate the fine

structure constant, reduced Planck constant, vacuum speed of light, and the single ion mass, respectively.

Assuming the spectator modes are in a thermal state with average motional occupation number \bar{n}_s , this Kerr effect imparts not only an average mode frequency shift, but also causes an apparent superposition of target mode frequencies according to the variance of the spectator state motional occupation number $\bar{n}_s (\bar{n}_s + 1)$. As the average frequency change is calibrated for as part of the gate tune up procedure presented in section 6.2, gate errors are only caused by this spread of mode frequencies. Propagating this superposition of mode frequencies through the lowest order expansion of the gate fidelity given in equation 3.30, we find that gate error resulting from the spectator mode s may be approximated by equation 3.36 [65].

$$\epsilon_{kerr} \simeq \frac{2\bar{n}_g + 1}{8K} \bar{n}_s (\bar{n}_s + 1) (\chi_s t_g)^2 \quad (3.36)$$

Inspecting the axial stretch and radial rocking mode interaction strength in equations 3.34 and 3.35, we note that the Kerr error scales unfavourably with the axial confinement strength ω_z . Respectively, the stretch and rocking mode induced Kerr sensitivities χ scale as $\omega_z^{7/3}$ and $\omega_z^{10/3}$. If the axial stretch and out of plane rocking modes are not ground state cooled, their temperatures are determined by the Doppler cooling temperature limit. Consequently, the occupation number of the axial stretch mode scales as $\bar{n}_{\text{stretch}} \propto \omega_z^{-1}$. The overall Kerr error induced by the axial stretch and radial rocking mode scale as $\omega_z^{2+2/3}$ and $\omega_z^{6+2/3}$ respectively. Notably, the Kerr effect induced error does not improve with gate loop number as $t_g^2 \propto K$.

Using typical mode frequencies of $\omega_z = 2\pi \times 1$ MHz, $\omega_g = 2\pi \times 5.6$ MHz, and $\omega_{\text{oop-rocking}} = 2\pi \times 6.0$ MHz combined with a 15° mode radial mode tilt angle to the horizontal we can estimate the Doppler limit for each relevant

mode to be $\bar{n}_{\text{stretch}} \approx 5$ and $\bar{n}_{\text{oop-rocking}} \approx 24$. For a single loop, $t_g = 300 \mu\text{s}$ MS gate this would result in stretch and radial mode induced Kerr errors of $\epsilon_{\text{stretch}} \approx 1 \times 10^{-4}$, and $\epsilon_{\text{oop-rocking}} \approx 6 \times 10^{-5}$. However, the Kerr error, especially that associated with the out of plane the rocking mode, rapidly gains significance as the axial confinement strength is increased. Increased axial confinement may be desirable to allow increased spectral separation of the in-plane radial centre of mass and rocking modes, see section 7.2.4.1.

3.2.3.4 Off-resonant excitation of hyperfine transitions

The ideal qubit we have considered so far was a two level system where only a single spin transition is possible. However, as we noted in section 2.3.2, $^{43}\text{Ca}^+$ has a rich hyperfine structure, giving rise to a multitude of spectator spin transitions. Of these, four σ polarised transitions are directly connected to the qubit states. At a magnetic field of 288 G, these transitions are ~ 100 MHz detuned from the resonance frequency of the qubit transition. When performing square microwave pulses of suitable amplitude and frequency to drive $\sim 300 \mu\text{s}$ MS gates, there is a $\sim 1\%$ chance of off-resonantly exciting these spectator transitions, resulting in qubit leakage errors.

Besides exciting spectator spin transitions, the π component of the microwave field driving the MS gate drives Rabi oscillation on the qubit transition. In principle, these Rabi oscillations commute with the MS gate Hamiltonian and need not necessarily be considered an error, but implement a deterministic qubit rotation which may be calibrated.

3.3 Error suppression techniques

Significant two-qubit gate errors are likely to arise from four main sources, off-resonant driving of spectator hyperfine transitions, off-resonant transitions

Error suppression technique	Suppressed error sources
Pulse shaping	Off-resonant excitation of spectator transitions & qubit
Sideband Walsh modulation	Mode frequency fluctuation
Resonant decoupling tone	Qubit frequency fluctuation
Walsh modulated resonant decoupling tone	Qubit frequency fluctuation & imperfect decoupling pulse

Table 3.1: **Error suppression techniques** - We discuss the error suppression techniques listed above. The motivating errors source behind these techniques is indicated in the right-hand column.

on the qubit transition, motional frequency uncertainty, and finally qubit frequency shifts. In the following, we discuss techniques suppressing each of these error sources. An overview of the suppression techniques for each error is presented in table 3.1. In particular, off-resonant excitation and qubit frequency errors are exacerbated by the strong drive fields needed to compensate for the comparatively low Lamb-Dicke parameter available for “laser-free” two qubit gates, see section 3.1.2.

3.3.1 Off-resonant excitation

When we discussed the microwave interaction in section 3.1.1, we noted that microwave tones, besides driving the desired motional interaction, also drive Rabi oscillations. When considering the ideal MS gate we neglected this term. Reintroducing it now, and recalling that the gate tones have equal power and opposite detuning, we write the MS Hamiltonian of the n^{th} ion as follows.

$$\begin{aligned}
 \tilde{H}'_{MS,n}(t) &= \tilde{H}'_{B,n}(t) + \tilde{H}'_{R,n}(t) \\
 &= \frac{1}{2}\hbar\Omega_{Q,n} \left(e^{i(\delta+\omega_{\text{motion}})t-i\frac{\Delta\phi}{2}} + e^{-i(\delta+\omega_{\text{motion}})t+i\frac{\Delta\phi}{2}} \right) \left(\hat{\sigma}_{+,n}e^{i\bar{\phi}} + \hat{\sigma}_{-,n}e^{-i\bar{\phi}} \right) \\
 &\quad + \frac{1}{2}\hbar\Omega'_{Q,n} \left(\hat{\sigma}_{+,n}e^{i\bar{\phi}} + \hat{\sigma}_{-,n}e^{-i\bar{\phi}} \right) \left(\hat{a}e^{-i(\delta t-\frac{\Delta\phi}{2})} + \hat{a}^\dagger e^{i(\delta t-\frac{\Delta\phi}{2})} \right)
 \end{aligned}
 \tag{3.37}$$

Above, we have factored out the time dynamics of the qubit rotations driven by the two microwave tones. Assuming the microwave amplitude and gradient are in phase, the qubit subspace operator of the two terms is identical. As the first term in equation 3.37 acts as an identity operator on the motional state, the two terms commute. Off-resonant carrier excitation can therefore, at worst, generate a qubit rotation error. Such errors may be fixed by book keeping and suitable compensation with single qubit operations. None-the-less, this necessitates further calibrations, which is undesirable. We therefore turn our attention to ways of suppressing the resulting qubit rotation.

Inspecting the time dependence of the qubit rotation term we may regroup the complex exponentials into a sinusoid as shown below.

$$\tilde{H}''_{MS,n}(t) = \hbar\Omega_{Q,n} \cos\left((\delta + \omega_{motion})t - \frac{\Delta\phi}{2}\right) \hat{\sigma}_{\bar{\phi},n} \quad (3.38)$$

Equation 3.38 is equivalent to a single amplitude modulated microwave tone which is resonant with the qubit transition. The sinusoidal amplitude modulation results in a π phase shift of the driving field in $\Delta t = \frac{\pi}{(\delta + \omega_{motion})}$ time intervals. As the tone is resonant, the driven dynamics on both sides of a phase flip cancel by symmetry. In the absence of qubit frequency errors the largest possible driven rotation is that driven over a single interval Δt . Further, for a given gate duration, the amplitude modulation envelope may be shifted by tuning the difference phase $\Delta\phi$. So far, this parameter is unconstrained and we may choose a value which coherently eliminates the residual qubit rotation.

An alternate way of suppressing the off-resonant qubit excitation is through pulse shaping. This may be represented in equation 3.38 through a time dependent $\Omega_{Q,n}$. Ramping into and out of the sideband tones slowly such that $t_{\text{ramp}} \gg \Delta t$ adiabatically suppresses the final qubit rotation. This pulse shap-

ing approach has the benefit of not requiring additional book-keeping and also suppresses excitation of other hyperfine spectator transitions. In our case, a pulse edge time on the few μs time scale provides sufficient suppression and does not significantly extend the gate duration. We therefore adopt the pulse shaping approach for controlling off-resonant excitation.

3.3.2 Sideband Walsh modulation

Discussing the ideal MS gate in section 3.2.1, we noted that a MS gate may be performed with an arbitrarily large number of phase space loops K . For a given gate interaction strength Ω'_Q , increasing the number of loops, increases the gate duration required to perform an MS gate. However, as noted in section 3.2.3.1, the resulting MS gate becomes more resilient to mode frequency mischaracterisation. This resilience can be further enhanced using the standard technique of sideband Walsh modulation [86], see figure 3.3.

Sideband Walsh modulation exploits the fact that a multi-loop gate returns to its initial phase-space position part way through the two-qubit gate. By changing the difference phase $\Delta\phi$ of the microwave tones at times where such a loop closure is expected, the trajectory of subsequent phase space loops is changed without upsetting the ideal gate operation as the same geometric phase is accumulated. As noted by Hayes et al. [86], changing the gate difference phase between the two values $\Delta\phi = \Delta\phi_0, \Delta\phi_0 + \pi$ according to Walsh functions of order $2^l - 1$, $l \in \mathbb{N}_0$ suppresses residual entanglement of the qubits with the motion and improves the scaling of the gate error ϵ for small mode frequency mischaracterisation $\Delta\delta \ll \Omega'_Q$ as shown in equation 3.39.

$$\epsilon \propto \Delta\delta^{2^{l+1}} \quad (3.39)$$

Notably sideband Walsh modulation of order $2^l - 1$ relies on the motional loop number of the implemented MS gate being an integer multiple of 2^l .




Figure 3.3: **Walsh sequences** - Walsh functions of interest for dynamical decoupling up to order 31. The y-axis shows the phase contribution associated with the Walsh modulation. The x-axis shows the elapsed gate operation time, normalised to the total gate duration. Different Walsh functions are offset by 2π for visual clarity.




Figure 3.4: **DDMS gate tones** - The DDMS gate employs three microwave tones. The red and blue detuned microwave sideband tones are chosen as for a normal MS gate. However, a third microwave tone is applied at the qubit frequency half way between the two microwave sideband tones. Typically the power of the ‘carrier’ decoupling tone is much weaker than that of the sideband tones. As the DDMS gate employs three microwave tones, it has three drive phase degrees of freedom. Beyond the average and difference phases ($\bar{\phi}$ and $\Delta\phi$) the phase of the carrier tone is tuned to match $\bar{\phi}$.

Consequently, for a given interaction strength Ω'_Q , use of sideband Walsh modulation of order $2^l - 1$ requires a factor $\sqrt{2^l}$ gate duration increase over a single loop gate of equal interaction strength.

3.3.3 Dynamical decoupling of qubit detuning errors

The motional interaction generated when using near field microwave gradients rotates the qubits around an axis in the x-y plane, as indicated by the definition of $\hat{\sigma}_{\bar{\phi},n}$ in equation 3.11. Due to the low Lamb-Dicke parameter, the large microwave fields needed to drive the two-qubit gate interaction generate significant AC-Zeeman shifts. Unfortunately, the $\hat{\sigma}_{\bar{\phi},n}$ gate dynamics do not commute with $\hat{\sigma}_z$ AC-Zeeman shifts. As the generated AC-Zeeman shifts are

significantly larger than the motional interaction strength $\Omega'_{Q,n}$, fractionally small AC-Zeeman shift variations will result in significant gate errors. As seen in section 3.2.3, without mitigation measures the AC-Zeeman shifts are expected to cause significant gate errors in excess of 1%.

Fortunately, the $\hat{\sigma}_{\bar{\phi},n}$ two-qubit gate interaction commutes with $\hat{\sigma}_{\bar{\phi},n}$ qubit rotations. This feature of the gate interaction allows us to apply a third microwave tone during the two-qubit gate interaction. Provided the third microwave tone is resonant with the qubit transition and of appropriate phase to drive $\hat{\sigma}_{\bar{\phi},n}$ Rabi oscillations, the gate dynamics are left unchanged, see figure 3.4. Such a three tone gate is referred to as a dynamically decoupled Mølmer-Sørensen (DDMS) gate. As demonstrated by Sepiol and Harty et al. [74, 41], DDMS gates are significantly less sensitive to qubit detuning errors, such as are caused by AC-Zeeman shifts [87, 88]. Fundamentally, this effect relies on the qubit Rabi frequency associated with the carrier tone being significantly larger than the two-qubit gate frequency. The stronger the carrier drive, the greater the associated suppression of qubit detuning errors. However, this represents a trade off between sensitivity to AC-Zeeman shift fluctuations and sensitivity to variations in the Rabi frequency of the carrier tone.

Previous work by Sepiol managed the driven qubit rotations by splitting the overall DDMS gate into two equal sections. The qubit rotation driven in each of these sections was calibrated to be a $m \times 2\pi$ pulse where m is an integer. Further, a calibrated π pulse was employed to coherently cancel Rabi frequency fluctuations common to both gate halves. This pulse was placed between both gate halves and required specific calibration of its amplitude to account for duty-cycle effects induced by the high power sideband tones. Further, the phase of this pulse is critical and requires calibration due to AC-Zeeman shifts. Beside the need for additional calibrations, this approach has

the downside that such “echo” pulses may only be placed when the motional gate dynamics have completed a closed trajectory in phase-space. As the gate interaction time t_g increases with loop number K as $t_g \propto \sqrt{K}$, this requirement means that higher order echo sequences come with increasingly large gate speed penalties.

3.3.4 Walsh modulated dynamical decoupling of qubit detuning errors

Due to the large AC-Zeeman shifts associated with faster laser-free gates, the use of a third, “dynamical decoupling” tone is desirable. However, we wish to avoid the drawbacks associated with the previous approach. With this goal in mind, we consider the Hamiltonian of a DDMS gate interaction shown in equation 3.40. Note that we have dropped the off-resonant Rabi oscillation term associated with the sideband tones and do not include the motional term of the dynamical decoupling tone as it is far detuned from the motional sideband transitions.

$$\begin{aligned} \tilde{H}_{DDMS,n}(t) = & \frac{1}{2} \hbar \Omega_{Q,n}^{\text{DD}} \left(\hat{\sigma}_{+,n} e^{i\phi_{\text{DD}}} + \hat{\sigma}_{-,n} e^{-i\phi_{\text{DD}}} \right) \\ & + \frac{1}{2} \hbar \Omega'_{Q,n} \left(\hat{\sigma}_{+,n} e^{i\bar{\phi}} + \hat{\sigma}_{-,n} e^{-i\bar{\phi}} \right) \left(\hat{a} e^{-i(\delta t - \frac{\Delta\phi}{2})} + \hat{a}^\dagger e^{i(\delta t - \frac{\Delta\phi}{2})} \right) \end{aligned} \quad (3.40)$$

When writing the Hamiltonian above, we have used $\Omega_{Q,n}^{\text{DD}}$ to indicate the frequency of the resonant Rabi oscillations driven by the dynamical decoupling tone. Similarly, ϕ_{DD} indicates the phase of applied dynamical decoupling tone, see figure 3.4.

Inspecting the DDMS Hamiltonian shown in equation 3.40, we note that the dynamical decoupling and MS terms commute if

$$\phi_{\text{DD}} = \bar{\phi} + n\pi, n \in \mathbb{Z} \quad (3.41)$$

We are free to choose any ϕ_{DD} with integer n . Further, we are free to change our choice of n during the gate operation without upsetting the entanglement dynamics driven by the sideband tones. The latter opens the option to echo out Rabi oscillation driven by the decoupling tone. As these phase flip based echo sequences are independent from the gate dynamics, higher order echo sequences may be implemented without imposing any requirements on the phase-space trajectory of the gate. This avoids the gate duration penalty incurred by previous schemes.

In principle arbitrary phase flip sequences of the dynamical decoupling tone may be chosen. In the following we choose to break the dynamical decoupling tone into 2^l equal segments of which the phase ϕ_{DD} may be chosen independently. The equal decoupling tone division is suitable for Walsh modulation of order up to W_{2^l-1} , see figures 3.3 and 3.5 [89]. Higher order Walsh modulated dynamical decoupling sequences will echo both the resonant Rabi oscillation driven inside each segment and polynomial Rabi frequency changes across the gate sequence, making them particularly desirable. This dynamical decoupling strategy is hence forth referred to as “Walsh Modulated Decoupling of Mølmer-Sørensen” (WMDMS) gates. As higher order echo sequences provide increasing immunity to decoupling Rabi frequency fluctuations, we no longer need to calibrate the driven Rabi oscillations. Further, significantly larger decoupling drive powers may be used without incurring significant qubit rotation errors.

In our considerations above, we noted several advantages of WMDMS gates over more traditional DDMS gates when the gate operation is resonant with the qubit transition. We now turn our attention to the performance of the




Figure 3.5: **WMDMS gate pulse sequence** - The diagram shows a WMDMS gate pulse sequence using two, Walsh1-1 modulated sideband pulses and eight, Walsh-7 modulated carrier pulses. The carrier phase, and average sideband phase are indicated in the plots below the pulse sequences. These phases are set to either zero or π for each of the pulses. Phase shifts arising from pulse shaping of, and dead times between sideband pulses are not included in the indicated phase.

WMDMS gate in the presence of detuning errors. We will need to consider both the protection of the motional interaction against detuning error and the continued effectiveness of the qubit rotation echo.

In the presence of detuning errors the dynamical decoupling term no longer commutes with the motional dynamics (as is necessary for the decoupling to increase resilience to qubit detuning errors). Intuitively, the effectiveness of the dynamical decoupling may be thought of as spin-locking to the eigenstates

of the decoupling interaction. The effectiveness of spin-locking relies on the associated Rabi oscillation exploring the entirety of their loop on the Bloch sphere. We therefore expect the decoupling effectiveness of the WMDMS gate to break down as the phase modulation interval approaches the decoupling drive Rabi frequency. Consequently, WMDMS gates require a sufficiently strong decoupling tone that the qubit rotation driven inside each decoupling drive segments is larger than 2π , see figure 3.6.

Perfect self cancellation of the driven qubit rotation via a flip of the drive phase is only possible for resonant Rabi oscillations. When visualised on the Bloch sphere, off-resonant Rabi oscillation lead to a tilted rotation axis. In the worst case scenario, a phase flip is performed when the qubit is at the extreme point of its Bloch sphere trajectory from the resonant trajectory. This offset in the qubit state is not remedied by the phase flip and will constructively build up over phase flips. Fortunately, this error may be kept small by ensuring that the decoupling drive strength $\Omega_{Q,n}^{\text{DD}}$ is sufficiently strong compared to the qubit detuning δ_q and number of phase flips j .

$$\Omega_{Q,n}^{\text{DD}} \gg (1 + j) \delta_q \quad (3.42)$$

The corresponding impact on the gate fidelity is demonstrated in figure 3.7. In practice we can readily satisfy the conditions for good qubit detuning error suppression.




Figure 3.6: **WMDMS qubit frequency error** - Simulated WMDMS gate resilience to qubit frequency errors for a single-loop gate with $\Omega_{Q,n}^{\text{DD}} t_g = 2\pi \times 50$. The probability of not producing the desired entangled state is indicated on the y-axis. The x-axis indicates the qubit frequency error normalised to the gate duration. The plot shows the qubit detuning induced error of a standard MS gate and the reduced error for a range of WMDMS modulation orders. Notably, the resilience of the WMDMS gates decreases drastically when the modulation order exceeds the number of qubit rotations $\frac{\Omega_{Q,n}^{\text{DD}} t_g}{2\pi}$.

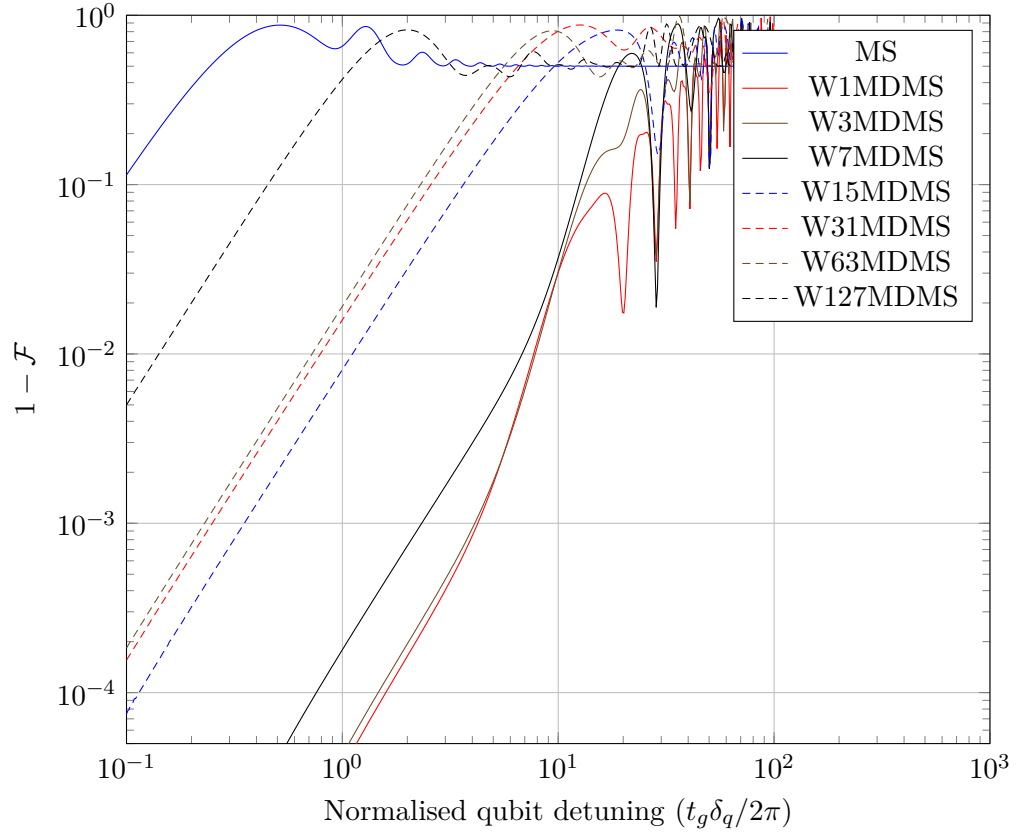


Figure 3.7: **WMDMS qubit detuning limit** - Simulated WMDMS gate resilience to qubit frequency errors for a single-loop gate with $\Omega_{Q,n}^{\text{DD}} t_g = 2\pi \times 100$. The probability of not producing the desired entangled state is indicated on the y-axis. The x-axis indicates the qubit frequency error normalised to the gate duration. The plot shows the qubit detuning induced error of a standard MS gate and the reduced error for a range of WMDMS modulation orders. Similarly to the behaviour in figure 3.6, the resilience of the WMDMS gates decreases drastically when the modulation order exceeds the number of qubit rotations $\frac{\Omega_{Q,n}^{\text{DD}} t_g}{2\pi}$. Further, the detuning resilience rapidly deteriorates as the detuning error approaches $\delta_q \approx 10^1 \times 2\pi/t_g$.

Chapter 4

Apparatus

In this chapter we explore the apparatus which, unless otherwise specified, was used in the experimental demonstrations presented in this work. Large parts of the apparatus were described in preceding work. Consequently the aim of this section is to provide an overview of the key apparatus. Detailed presentations of the apparatus are found in references [63, 65, 62]. Where the apparatus has been updated since completion of the preceding work, a more detailed presentation is made.

4.1 The ion trap assembly

The ion trap used in this work is a surface Paul trap with segmented DC electrodes and an integrated microwave electrode. Figure 4.1 shows a rendered image of the surface trap, which consists of 22 DC electrodes around a central RF electrode and $\lambda/4$ microwave resonator. The trap is a single layer design, fabricated from gold deposited on a sapphire substrate. The choice of a sapphire substrate provides both low dielectric losses, and high thermal conductivity, avoiding heat build-up from resistive losses. A detailed presentation of the trap design, fabrication, and characterisation is presented by Weber et




Figure 4.1: **Ion Trap** - a: render of the surface Paul trap. Electrical connections for DC, radio frequency and microwave signals are indicated by labels. b: Enlarged view of the centre of the surface trap electrodes. Electrodes are colour coded as either DC, radio frequency, microwave or electrical ground. Figures reproduced from Weber et al. [62].

al. [62] and is not reproduced here.

4.1.1 Partial nulling of the microwave field

A key feature of the trap is its integrated microwave resonator. The resonator consists of a single electrode, the length of which is chosen such that it acts as a $\lambda/4$ resonator. The incoupling end of the electrode is capacitively coupled to the microwave feed line, the other end of the resonator electrode is connected to the trap ground. The electrode is routed such that it makes a ‘U’ shape as indicated in figure 4.2 (a). Such an arrangement causes the formation of a zero crossing of the in-plane microwave magnetic field component as outlined in figure 4.2 (b). In our apparatus the in-plane magnetic field zero crossing corresponds to a null of the π polarised magnetic field component. Notably, the null is only present for the π polarised microwave magnetic field component and there is no cancellation of the σ polarised microwave field. Characterisation measurements of the microwave magnetic field were performed by using a trapped ion as a magnetic field probe and are presented in table 4.1. Some of the results presented in this thesis were measured after an experimental upset




Figure 4.2: **Partial microwave field nulling** - a: Schematic view of the radio frequency and microwave electrodes. Respectively, radio frequency, microwave and electrically grounded electrodes are indicated in blue orange and grey. The position of the pseudo potential minimum is indicated by the red dashed line. Finally, B_0 indicates the direction of the externally applied static magnetic field. Notably, the microwave current in the ‘U’ shaped microwave electrode flows in opposite directions in the two legs of the ‘U’. b: A simplified cross section of the magnetic field generated by current flowing in the microwave electrode. The figure has been reproduced from Weber et al. [62].

Pillbox Temperature	gradient ($\frac{\partial B_\pi}{\partial x}$)	amplitude at RF null (B_π)
300 K	$20.4 \pm 0.2 \text{ T m}^{-1}$	$92.2 \pm 0.4 \text{ } \mu\text{T}$
77 K	$30.7 \pm 0.2 \text{ T m}^{-1}$	$166.8 \pm 0.5 \text{ } \mu\text{T}$
21 K	$39.4 \pm 0.2 \text{ T m}^{-1}$	$223.4 \pm 0.9 \text{ } \mu\text{T}$

Table 4.1: **Microwave field** - Measured field gradient, and amplitude at the RF null position of the microwave π field when operating at a selection of trap temperatures with 1 W microwave input power. The microwave underwent substantial change after an experiment upset (see appendix B).

that changed the microwave field produced by the trap. The upset is discussed in appendix B. Design details and in-depth analysis of the microwave field are presented elsewhere [63, 62].




Figure 4.3: **Vacuum Assembly** - Rendered image of key in-vacuum components. (i) cryostat head; (ii) thermal strap; (iii) microwave coax; (iv) DC electrode connections; (v) CF100 base flange; (vi) 51 pin micro-D feed through; (vii) temperature sensor; (viii) pillbox lid; (ix) surface-trap; (x) atomic ovens. Figure reproduced from Weber et al. [62].

4.1.2 Cryogenic Pillbox

The surface trap is eutectically bonded to a cryogenically cooled, copper pillbox. Besides providing thermal contact to the cryostat and cutouts for optical access to the trap, the pillbox provides mounting points for in-vacuum printed circuit boards (PCBs), see figure 4.3. There are five such PCBs, the first of these hosts a resistive temperature detector (RTD) and is mounted to the bottom of the pillbox. The remaining four PCBs are wirebonded to the trap electrodes and provide electrical connection from the in-vacuum wiring harness to the trap electrodes. The microwave and radio frequency electrodes

each have a dedicated PCB for their signal delivery. The remaining two PCBs provide connection to the ‘DC’-electrodes of the trap. A detailed design breakdown of the pillbox and PCBs are presented by Wolf [63] and Weber et al. [62].

4.1.3 DC electrode filters

The surface trap electrodes are connected to adjacent printed circuit boards (PCBs) via wirebonds. These PCBs are mounted to the cryogenic pillbox adjacent to the surface trap. Besides allowing connection to the wiring harness, which delivers the electrical signals that should be applied to the electrodes, the PCBs connecting to the ‘DC’ electrodes (see figure 4.1) feature an RC low-pass filter¹ with a ~ 160 kHz filter cut-off frequency. This filter is designed to limit heating of the ion motion caused by voltage noise on these electrodes. After the measurements presented in section 5.6 were completed, a subset of these filters suffered loss of the RC filter capacitor. The resulting increased noise dominated the ion heating process. Consequently, the two-qubit gate demonstration was performed with an additional external second-order RC low pass filter² with poles at ~ 200 kHz and ~ 300 kHz. This configuration achieved the same ion temperatures and heating rates as prior to these apparatus failures.

4.2 Vacuum Assembly

The ion trap and pillbox are contained within an octagonal ultra high vacuum chamber featuring six side viewports for laser access and a top viewport for imaging. The 7th octagon face connects to a nipple with ion pumps, and

¹The filter uses 1 k Ω resistors [SRT Restech, CHR 0603-HT 4K + 1K 2%] and 1 nF capacitors [SRT Microceramique, 0603A102JBCBN (100V)]

²The filter uses 20 Ω resistors and 100 nF capacitors

ion gauges. Finally, the opposite face from the pumps connects to a nipple allowing connection to a flow cryostat. The cold head of this cryostat nipple is connected to the pillbox via a flexible copper strap, providing good thermal transfer while retaining mechanical isolation from the cryostat. Besides the pillbox, the inside of the vacuum chamber hosts two resistively heated atomic ovens, one containing calcium and one strontium metal [60]. The ovens are placed adjacent to, but thermally isolated from, the pillbox such that their atomic beams propagate through the trapping region. The calcium contained within this oven is isotopically enriched with ^{43}Ca . The strontium oven was not used in this work. Electrical signals are passed to an in-vacuum wiring harness through a 51 pin micro-D connector. Full details of the vacuum chamber design are presented by Wolf [63].

4.3 Imaging system

The ion trap system collects light emitted by the trapped ions via a $\text{NA} = 0.6$ composite primary lens, which is located above the vacuum chamber. Shortly after the composite lens, an aperture is built into the imaging assembly that may be used as an alignment aid. During normal operation this aperture is fully open. Adjustable slits are placed in the image plane of the primary lens, enabling spacial filtering of the detected light. Next, the imaging light is relayed using an aspheric³ and achromatic lens⁴. As we wish to detect the ion state using a photo-multiplier-tube (PMT), while maintaining a live image of the ions for diagnostic purposes, we place a 90:10 beam splitter inside the optical relay. The majority of collected photons are directed to the PMT⁵ for state detection. Meanwhile, the lower intensity arm is directed to an

³Thorlabs, AL1225G-A

⁴Thorlabs, AC254-100-A

⁵Hamamatsu, H10682-210 MOD

EM-CCD camera⁶ which is used for diagnostic purposes. Background light detection from ambient light is suppressed by using a notch filter⁷ in front of the PMT and a short-pass filter⁸ in front of the camera. Figure 4.4 shows a schematic view of the imaging system arrangement.

4.4 Cryogenic equipment

The ion trap may be operated both at room and cryogenic temperatures. To enable cryogenic operation, a helium flow cryostat⁹ is used. The cryostat is supplied with liquid helium (LHe) from a dewar. Control of the helium flow rate is achieved with a mass flow controller¹⁰ placed in the room temperature exhaust of the cryostat system. Full details of the cryogenic design are presented by Wolf [63].

4.5 Magnetic Field coils

The static magnetic field setting the quantisation axis of the ions is provided by three orthogonal electromagnets. Each electromagnet consists of two coils in a Helmholtz configuration around the trap centre. The strongest of these, the ‘main’ electromagnet, produces a ~ 288 G magnetic field and is oriented to provide a horizontal magnetic field which is perpendicular to the trap axis. The orientation of this magnetic field coincides with the desired quantisation axis of the trapped ions. The other two ‘trim’ electromagnets are used to compensate stray magnetic fields, such as the earth magnetic field. Design details of these coils are presented by Wolf [63]. However, it is worth noting that the

⁶Andor, iXonUltra897

⁷Semrock, 6FF01-400/12-25

⁸Semrock, FF01-468/SP-25

⁹Janis, ST-400 UHV Supertran.

¹⁰M J Wilson Group Ltd, 4850ABB2T1B1K3C00A

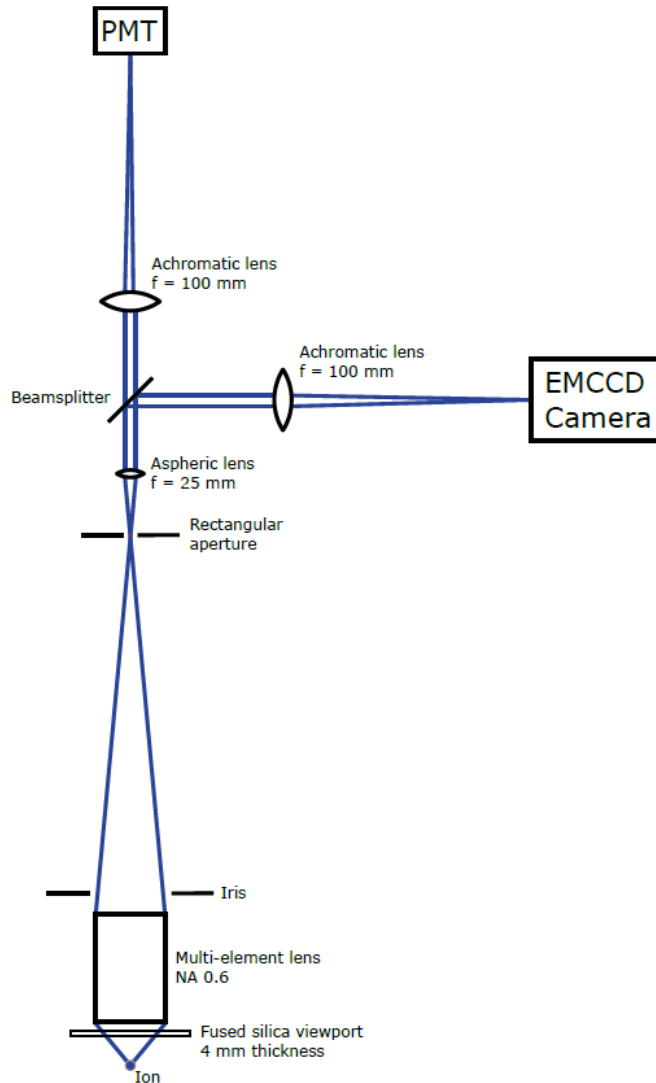


Figure 4.4: **Imaging system** - The schematic shows the imaging optics used for detection of 397 nm fluorescence emitted by the trapped ions. Mirrors used to redirect the beam path within the imaging system have been omitted. Image reproduced from Löschnauer [65].

magnetic field demanded from the ‘main’ magnet is near three orders of magnitude stronger than that demanded of the ‘trim’ magnets. Consequently, the trim coils are suitably powered using a bench top power supply¹¹. The design of the water cooled ‘main’ magnet is such, that a magnet current of ~ 148 A is needed to produce the required magnetic field. This current is supplied by an Agilent 6671A power supply.

4.5.1 Active stabilisation of electromagnet current

As outlined in section 2.3.2, the energy structure of $^{43}\text{Ca}^+$ is sensitive to the applied magnetic field. Though we choose to perform logic operations on a first order magnetic field insensitive ‘clock’ transition, magnetic field fluctuations may still introduce logic errors via two mechanisms. Firstly, the state preparation and readout procedures require transfer pulses to be carried out on magnetic field sensitive transitions. Secondly, the microwave fields used to drive the two-qubit logic operations produce magnetically sensitive dressed states. Consequently magnetic field fluctuations result in qubit decoherence errors during two-qubit logic operations. Due to the large field demanded from the ‘main’ electromagnet, the magnetic field noise experienced by the trapped ions is dominated by current noise of the ‘main’ electromagnet. We choose to suppress this electromagnet current noise using a closed feedback loop, which shunts current past the electromagnet coils.

The control loop measures the current passing through the ‘main’ magnet using a LEM, IT-605 S current transducer¹². This provides a suitable current sensing range and has a 300 kHz small signal bandwidth, which is sufficiently high to not limit the attainable feedback loop bandwidth. Ana-

¹¹TTi, QL355TP

¹²To better utilise the sensing range of the transducer we use two passes of the ‘main’ magnet cable through the transducer




Figure 4.5: Magnetic field stabilisation setup - Circuit diagram showing the configuration used to implement active stabilisation of the electromagnet current. The Agilent 6671A power supply provides a current $I_0 \approx 148$ A and has output capacitance $C_{out} = 44$ mF. Due to the wire and coil resistance we have $V_0 \approx 4.1$ V. After the power supply, the current path I_{coil} leads to the electromagnet and current transducer. The other current path I_{fb} passes through a bipolar transistor controlled by the feedback signal.

logue preprocessing, and digitisation of the sense signal are performed using a modified open-source mezzanine board [90] and the open source ‘Stabilizer’¹³ CPU-based feedback controller [91]. The analogue frontend implemented by the mezzanine board has a low temperature coefficient resistor¹⁴, converting the current transducer current into a differential voltage signal. The mezzanine board has a small signal, analogue bandwidth of 100 kHz and performs analogue subtraction of a tunable voltage offset from the differential voltage signal followed by $\sim 2000\times$ amplification of the resulting offset subtracted error signal. This pre-processed analogue error signal is digitised using the 16-bit ADC¹⁵ built into Stabilizer. After digitisation, the stabilizer microprocessor implements a 500 MS/s digital control loop which produces an analogue

¹³We employed a Stabilizer v1.0

¹⁴Vishay Precision Group, Y160710R0000D9R

¹⁵LTC2310-16

control signal using the inbuilt stabilizer DAC¹⁶. This analogue control signal is post-processed on the mezzanine board where it then controls the shunted current which bypasses the ‘main’ electro magnet via a power transistor¹⁷. As the power supply output capacitance $C_{\text{out}} = 44 \text{ mF}$ and self inductance of the electromagnet $L_{\text{out}} \approx 5 \text{ mH}$ form a resonance at $\sim 10 \text{ Hz}$, feedback at higher frequencies requires appreciable output range to overcome the $\propto f^{-2}$ response suppression past this LC resonance. Consequently the mezzanine board was modified to extend the shunt current feedback range to 520 mA.

Modifications of the mezzanine board were performed for two reasons. Firstly, the ceramic capacitor filtering the analogue frontend subtracted signal¹⁸ was replaced with a $22 \mu\text{F}$ film capacitor¹⁹ to eliminate piezoelectric and pyroelectric noise originating from the ceramic capacitor. Secondly, the feedback current range was increased by increasing the analogue gain from the DAC signal to the transistor current. This gain increase was achieved by replacing resistors R10 and R16 with 1k5 and 4k7 resistors.

Due to the low resonance frequency between the power supply and electromagnet, a naive PID loop does not provide satisfactory loop bandwidth while retaining sufficient phase margin. Further, the LEM, IT-605 S current transducer has intrinsic 15.84 kHz and $3 \times 15.84 \text{ kHz}$ spurs which a PID loop would imprint onto the electromagnet current. Consequently a more suitable digital control loop was employed to maintain satisfactory phase margin and mitigate the effects of sensing spurs and high frequency ringing. This control loop consisted of five cascaded bi-quad filters, allowing considerable tunability to attain a suitable controller response. Figure 4.6 shows the measured²⁰ cur-

¹⁶AD5542A

¹⁷STMicroelectronics, TIP29A

¹⁸C33 of the mezzanine board

¹⁹Rubycon, 16MU226MD35750

²⁰Measurements were performed using a Rohde & Schwarz, UPV audio analyser

rent noise spectrum improvement attained by choosing a control loop tuned via the following procedure:

1. Place a notch filter with quality factor $q = 0.1$ at the frequency matching the most significant sensor spur at 15.84 kHz
2. Place a notch filter with quality factor $q = 0.5$ matching the next most significant frequency sensor spur at 47.52 kHz
3. Place a 4th order, critically damped low-pass filter at 40 kHz to suppress high frequency ringing.
4. Use a PI loop with corner frequency of 5 Hz and integrator gain limiting pole at 1 mHz. The corner frequency is chosen to avoid ringing due to the LC resonance of the electromagnet and power supply. Meanwhile, the integrator gain limit is chosen to limit hysteresis of the control loop on major configuration changes.
5. Place a lead compensator around the target loop bandwidth. Here we chose with a zero at 7 kHz and pole at 9 kHz.
6. Increase overall loop gain until the control loop starts ringing and then step back gain by 50%

Conversion of low-pass and notch filters to biquad filter coefficients was performed using the formula presented by Leibbrandt [92]. The desired PI and lead compensator poles and zeros were converted to biquad coefficients using the Z-transform.

Characterising the coherence time between the magnetically sensitive $4S_{1/2}$, $|F = 4, m_f = +4\rangle$ and $|F = 3, m_f = +3\rangle$ states employing this control loop

gives a two order of magnitude coherence time improvement, from $\sim 50 \mu\text{s}$ to $> 1 \text{ms}$ ²¹.

4.5.2 Ion servo

The active feedback loop described above stabilises the magnetic field experienced by the ion. However we also wish the static magnetic field experienced by the ion to have a specific magnitude such that the $|F = 4, m_f = +1\rangle$ to $|F = 3, m_f = +1\rangle$ transition is a clock transition, see section 2.3.3. We sense the magnetic field strength experienced by the ion using Rabi spectroscopy of the magnetically sensitive $|F = 4, m_f = +4\rangle$ to $|F = 3, m_f = +3\rangle$ transition. As detuning of the measured resonance frequency from the expected transition frequency corresponds to a magnetic field error, we adjust the setpoint of the active feedback loop such that the measured and expected transition frequency match within error. Typically, the $|F = 4, m_f = +4\rangle$ to $|F = 3, m_f = +3\rangle$ transition frequency is measured to within $\pm 0.6 \text{kHz}$, corresponding to a static detuning error of $\mp 0.3 \text{mG}$.

4.6 Laser Systems

As outlined in chapter 2, lasers are used for a range of functions within the experiment. These laser systems are identical to the configuration presented by Löschnauer [65], the following therefore provides an overview of laser systems.

All light sources used in this work fall into one of two frequency bands. A red band ranging from 850 nm to 866 nm, and a blue band from 378 nm to 423 nm. With the exception of a 378 nm free running diode²², all light sources are external cavity diode lasers²³. Where absolute frequency stability

²¹Observed coherence time varies depending on environmental conditions

²²Toptica, iBeam smart

²³Toptica, DL pro HP

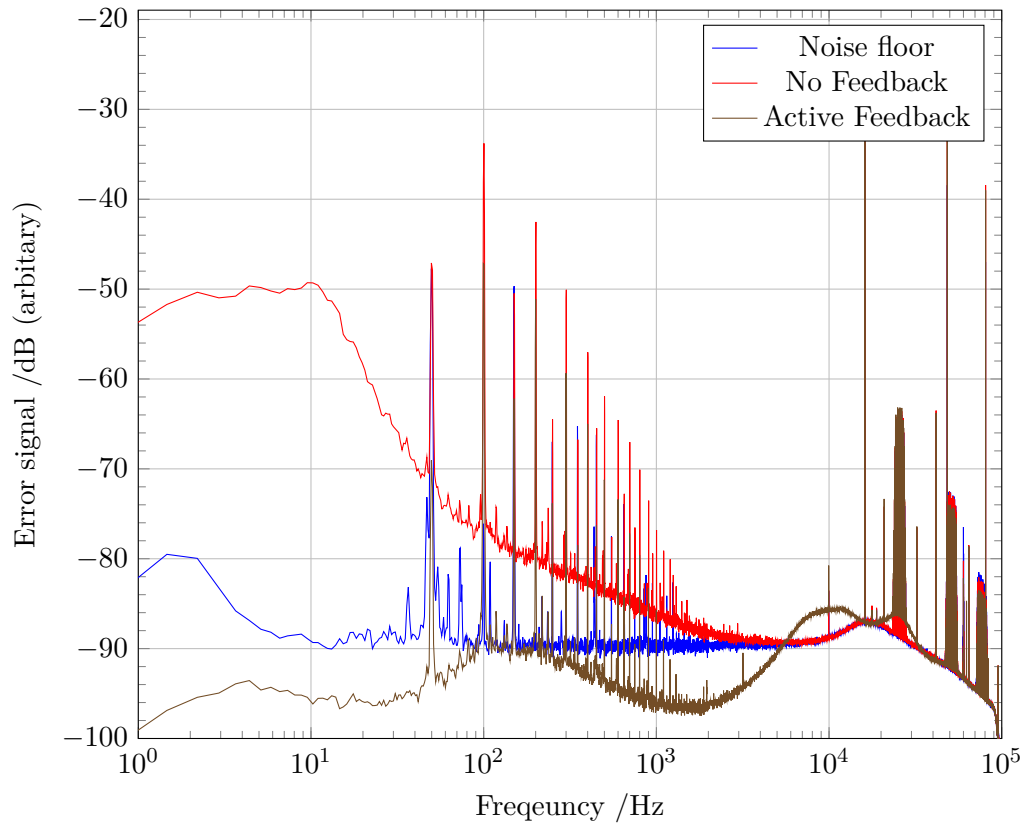


Figure 4.6: **Magnetic field stabilisation** - Measured coil current error signal spectrum. The noise spectrum in the absence of active current stabilisation is indicated in red. Notably, the coil current has mains noise spurs at multiples of the 50 Hz mains electricity cycle. The most significant of these are at 100 Hz and harmonics thereof. Enabling active feedback reduces the error signal noise spectrum up to the 5 kHz feedback bandwidth. The corresponding spectrum is shown in green. Finally, the error signal noise floor is indicated in blue. At frequencies up to the 5 kHz feedback bandwidth, and in the absence of active stabilisation, the error signal noise is above the noise floor. Reductions in the error signal below this noise floor do not result in improved electromagnet current stability. The noise floor also shows significant spurs at 15.84 kHz and harmonics. These are generated by the LEM, IT-605 S current transducer. To avoid imprinting these on the coil current, the control loop uses digital notch filters blocking the 15.84 kHz and 3×15.84 kHz spurs.

and laser linewidth are less critical, laser systems are locked to a wavemeter²⁴ with frequency adjustment intervals of approximately 30 s. Lasers requiring few MHz frequency stability, are locked to low drift reference cavities²⁵ via the standard Pound-Drever-Hall (PDH) method [93]. An overview of the employed light sources is presented in table 4.2.

Except for the photo-ionisation light sources, the light emitted by the lasers, is coupled into single mode polarisation maintaining fibres and delivered to rack mounted breadboards. These contain beam splitters, free space acousto-optic modulators (AOMs) and, in the case of one of the 397 nm lasers, an electro-optic modulator (EOM). The configuration for each laser is outlined in table 4.2. The EOM and AOMs are used to frequency shift the laser light, meanwhile AOMs are also used to control the light intensity. The beams produced on these breadboards are delivered to the optical table supporting the trap assembly via polarisation maintaining single mode fibres. As their beam paths are later shared and they have sufficiently similar wavelength the 393 nm σ^+ and 397 nm σ^+ beams are coupled into the same optical fibre using a free space 10:90 beam splitter²⁶. Similarly, the 850 nm π and 854 nm beams are merged with a fibre splitter²⁷.

Except for the photoionisation light, beams delivered to the trap table pass through a polarising beam splitter (PBS)^{28,29}, converting polarisation drifts to intensity fluctuations. The light then passes through an optical pickoff directing some of the light to a photodiode, which is used for beam intensity measurement. Figure 4.7 shows all beam paths to the ion and the beams propagating via each of these paths. Notably, the photo-ionisation, blue Dop-

²⁴HighFinesse GMBH, WS7 (PID2-MC8)

²⁵Stable Laser Systems (SLS), custom product or NPL, custom product

²⁶Thorlabs, BSN04

²⁷Thorlabs, PN850R5A1

²⁸blue lasers: Altechna custom PBS optimised for 393 nm to 423 nm

²⁹red lasers: Thorlabs, PBS122

λ	Type	Frequency reference	Optical switching & frequency tuning of delivered beam	Uses
378 nm	Free running diode	None	Shutter	Photo-ionisation
393 nm	Toptica, DL pro HP	SLS cavity	single pass AOM	Optical shelving readout
397 nm	Toptica, DL pro HP	SLS cavity	EOM, double pass AOM, followed by three beam paths with single pass AOMs	Laser cooling, fluorescence detection & optical state preparation
397 nm	Toptica, DL pro HP	Wavemeter	two beam paths wit a single pass AOM each	Raman sideband cooling, Raman thermometry
423 nm	Toptica, DL pro HP	Wavemeter	Shutter	Photo-ionisation
850 nm	Toptica, DL pro HP	NPL cavity	three beam paths with single pass AOMs	Optical shelving readout
854 nm	Toptica, DL pro HP	Wavemeter	single pass AOM	Optical state preparation
866 nm	Toptica, DL pro HP	SLS cavity	Double pass & single pass AOM	Laser cooling, fluorescence detection & optical state preparation

Table 4.2: **Light sources** - This table indicates the wavelength (λ), and apparatus details of the employed light sources. Additionally, the uses of each light source within this work are listed.

Wavelength	Beam waist radius	Saturation power (P_0)
393 nm	16 μm	396 nW
397 nm	16 μm	377 nW
423 nm	16 μm	493 nW
850 nm	28 μm	121 nW
854 nm	28 μm	118 nW
866 nm	28 μm	113 nW

Table 4.3: **Beam spot sizes** - Measured spot sizes, and the associated transition saturation power of the beams delivered to the trapped ions are tabulated above. Table reproduced from Löschnauer [65].

pler cooling, 850 nm π , and 854 nm beams are all delivered to the ion via the same path, while being delivered to the table from three different fibres. The photo-ionisation and blue Doppler beams are combined using a PBS³⁰. The blue and red beams are then combined on a dichroic mirror³¹. Similarly, the vertically polarised Raman, and 866 nm beams are overlapped on a dichroic mirror³².

Each beam path is focused onto the ion using a 200 mm focal length lens. The resulting measured spot sized based on wavelength and beam width are shown in table 4.3.

4.6.1 Beam polarisations

Besides the photoionisation, 854 nm, and 866 nm beams, the light delivered to the ion should have a specific polarisation. After passing through the PBS placed before the photodiode pickoff, the light in each beam is linearly polarised in the horizontal plane.

In the case of the 397 nm Doppler cooling beam, we desire to eliminate the σ^- polarisation component of the beam, which is achieved by placing

³⁰Altechna custom PBS optimised for 393 nm to 423 nm

³¹Thorlabs, DMPL567

³²Thorlabs, DMPL805




Figure 4.7: **Trap optical beams** - Diagram indicating the direction of optical beams relative to the ion trap and applied magnetic field. Optical setup schematic. The 397 nm Doppler cooling laser propagates at 45° to the trap axis, enabling cooling of both radial and axial modes of the collective ion motion. The 866 nm re-pumping laser is set up counter-propagating to the 397 nm Doppler cooling beam, facilitating better dark resonance cooling. The photoionisation and 854 nm beams are set up co-propagating with the 397 nm Doppler beam to reduce the number of beam alignments needed to trap ions. As the 850 nm and 854 nm beams are similar in wavelength it is convenient for 850 nm π beam to co propagate with the 854 nm beam. The σ^+ -polarised beams are required to propagate parallel to the applied magnetic field. Due to their similar wavelengths the 393 nm and 397 nm σ^+ beams are setup co-propagating through the same beam path. The 850 nm σ^+ beam is delivered to the ion from the opposite direction allowing the blue and red lasers to both use optics optimised for their respective wavelength. The propagation directions of the 397 nm Raman beams are chosen such that the Raman transition recoil projects onto the radial modes of the ion motion.

blue optimised $\lambda/2$ ³³ and $\lambda/4$ ³⁴ first order wave-plates immediately before the final focusing lens. As we wish to fulfil the same condition for the 850 nm π beam, we place a further red optimised $\lambda/2$ ³⁵ and $\lambda/4$ ³⁶ first order wave-plates before the dichroic combining the blue and red beams on this beam path. The polarisation of the case of the vertically polarised Raman beam, the horizontal polarisation is rotated using a $\lambda/2$ wave-plate³⁷.

Looking to the circularly polarised beams, we employ a Glen-Taylor polariser³⁸ to achieve high polarisation purity. The polariser is located immediately before a $\lambda/4$ wave-plate³⁹ and the focusing lens of the red and blue σ polarised beam paths. As the polarisation of the 397 nm σ^+ beam is critical for state preparation fidelity, the wave-plate in this path is mounted on a tip-tilt mount that is used to optimise the circular polarisation for a wave-plate for the 397 nm wavelength.

4.7 Control System

Realtime control of the apparatus is provided using the ARTIQ (Advanced Real-Time Infrastructure for Quantum physics) control system, using the Sinara open-hardware ecosystem [94, 91]. The ARTIQ realtime control is programmed using a python derived high level programming language. Further, ARTIQ allows control of non-realtime lab equipment through remote procedure calls using the python language. Beyond allowing for programming of experiment procedures, ARTIQ provides a dataset structure accessible to experiments, as well as realtime graphing of acquired data. To improve the

³³Thorlabs, WPH05M-405

³⁴Thorlabs, WPQ05M-405

³⁵Thorlabs, WPH05M-850

³⁶Thorlabs, WPQ05M-850

³⁷Altechna, custom product

³⁸Thorlabs, GT10-A and GT10-B for blue and red beams respectively

³⁹Thorlabs, WPQ05M-405 and WPQ05M-850 for blue and red beams respectively

expressiveness of written experiments, the ‘ndscan’ ARTIQ extension is used for the majority of the presented experiments [95].

4.7.1 Realtime controlled hardware

The system offers nanosecond control of event timing controlled via a ‘master’ Field Programmable Gate Array (FPGA)⁴⁰ which executes the experiment program. The FPGA is connected to Eurocard Extension Modules (EEMs) via ribbon cables. These EEMs provide modular hardware capabilities which are controlled in realtime. Beyond the ‘master’ FPGA and its connected EEMs, a ‘satellite’ FPGA⁴¹ is used. This is connected to the ‘master’ FPGA via SFP fibre optic link. This arrangement offers use of an increased number of EEMs and retains nanosecond realtime control. In the following we detail the EEM modules used with the ‘master’ and ‘satellite’ FPGAs. All of these extension modules are part of the open source Sinara hardware ecosystem [91].

4.7.1.1 DIO-BNC

This EEM module provides two banks of four configurable digital inputs/outputs. We employ the board with one bank configured as digital inputs. These inputs are used as edge counters for the clicks produced by the imaging system PMT (see section 4.3) and as an optional experiment trigger which is indexed to the 50 Hz mains cycle. However, outside of characterising the magnetic field experienced by the ion (see section 4.5.1) this mains cycle trigger was not used in the experiments presented within this work. The second DIO board bank of four channels is configured as digital outputs providing realtime controlled TTL signals. One of these TTL outputs is used to provide realtime control of the EOM used with the 397 nm cooling and state preparation beam via an

⁴⁰Sinara-HW, Kasli v1.0 (speed grade -3)

⁴¹Sinara-HW, Kasli v1.1 (speed grade -3)

RF switch⁴². A second output channel is used as a trigger for an optional oscilloscope.

4.7.1.2 Fastino v1.2

This module is a 32 channel, 16-bit DAC⁴³ with a ± 10 V output range. The multichannel DAC features an analogue bandwidth of ~ 0.4 MHz and a sample update rate of 2.55 MSs^{-1} per channel, avoiding contention between channel updates. Further, the design places particular emphasis on low output noise in the ~ 0.1 MHz to ~ 5 MHz frequency range, corresponding to the range commonly coinciding with resonances of the collective ion motion. We use Fastino to control the voltage of the ‘DC’ trap electrodes, see section 4.1.

As part of this work, prototype testing of Fastino v1.0 was performed and Fastino revisions v1.1 and v1.2 were developed, providing a > 30 dB digital to analogue crosstalk reduction over Fastino v1.0. Fastino broadband noise performance data is presented in figure 4.8. As established during prototype testing, Fastino outputs are filtered using a 6.5 mH common mode current choke⁴⁴ to suppress common mode current noise originating from the Kasli switch mode power supply.

4.7.1.3 Phaser (baseband)

This is an RF ‘Arbitrary Waveform Generator’ (AWG) module providing two IQ pair analogue outputs in the ± 500 MHz range. The two IQ pair signals are generated using a 16 bit parallel DAC⁴⁵ operated at 1 GSs^{-1} . Each of the two IQ pair channels supports up to 5 independent tones in a ± 12.5 MHz band within the ± 500 MHz phaser output range. These tones may be independ-

⁴²Mini-Circuits, ZASWA2-50DR-FT+

⁴³AD5542ABCPZ

⁴⁴Würth Elektronik, SL2

⁴⁵DAC34H84

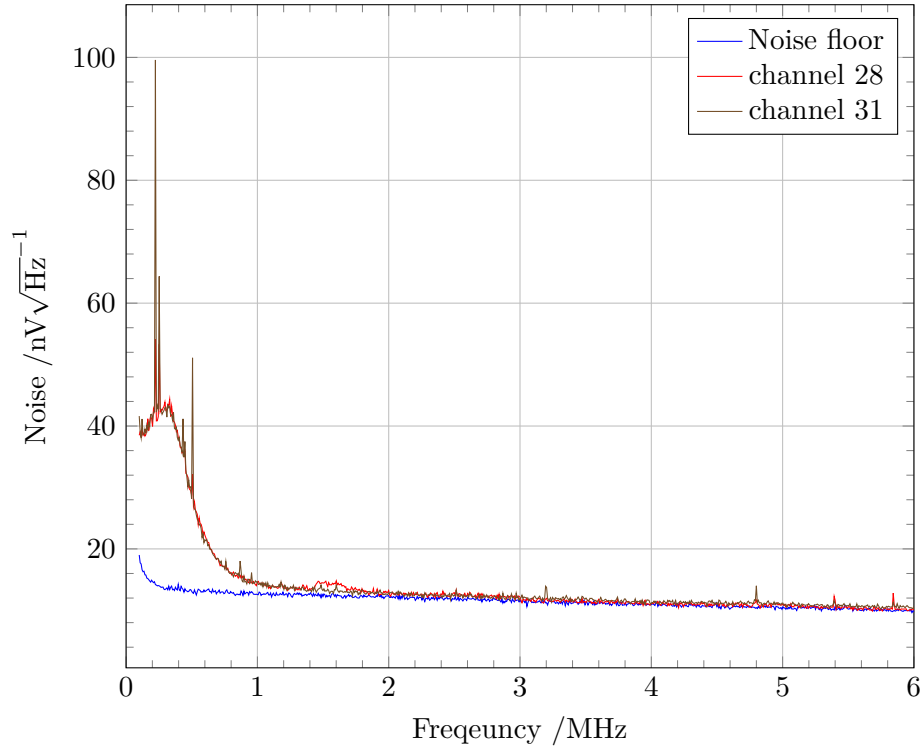


Figure 4.8: **Fastino noise** - Measured noise spectrum of Fastino outputs in-situ within the apparatus with no digital updates. The measurement was performed using a Agilent E4405B spectrum analyser with a 100 kHz DC-block set to a 1 kHz measurement bandwidth. For clarity, noise measurements for only two of the 32 output channels are shown. These are channel 28 in red and channel 31 in green. The noise floor of the spectrum analyser is indicated in blue. At frequencies beyond 200 kHz, the measured broadband noise drops off towards the measurement noise floor due to the bandwidth of the Fastino output filters. Some isolated noise spurs are measured within the output channels. Crucially, the noise spurs in the 0.5 MHz to 6 MHz band associated with the normal modes of the collective ion motion are weak and it is straightforward to engineer mode trap potentials such that collective ion motion resonances do not coincide with these noise spurs.

ently amplitude, phase, and frequency shaped with a 25 MSs^{-1} update rate. We use Phaser to generate the microwave signals driving our quantum logic operations.

Communication between the controlling FPGA and the Phaser EEM follows a 320 ns digital communication cycle. Due to implementation details of the digital sample generation, the output latency and pulse area produced when programming a pulse depend on phase alignment with this 320 ns communication cycle. Consequently, to achieve deterministic pulse sequences, the sequence playback must be aligned with this 320 ns cycle.

Phaser employs a multi-stage digital signal up-conversion process from the 25 MSs^{-1} user shaped signal to its 1000 MSs^{-1} DAC output. The bandwidth of the upsampling process is such that amplitude, phase, and frequency changes written at 25 MSs^{-1} are smoothly interpolated to the 1000 MSs^{-1} DAC output. However, amplitude, phase, and frequency changes written at lower update rates will result in undesirable stepped and/or ringing output signals. Consequently pulse shaping operations on Phaser are written with 25 MSs^{-1} per output tone. This is event rate is roughly an order of magnitude larger than the highest event rates Kasli is capable of sustaining⁴⁶. As the realtime event buffer in our ARTIQ version is limited to 512 realtime events, we are forced to leave significant dead times between shaped phaser pulses to avoid overwhelming Kasli. Finally, to attain the full $\pm 500 \text{ MHz}$ IQ signal output frequency range, the numerically controlled oscillator (NCO) integrated into the DAC34H84 must be switched between different frequencies. This switch must be programmed via the DAC SPI interface. The switching time of this NCO appears to be non-deterministic, forcing the use of a $\sim 20 \mu\text{s}$

⁴⁶Sustained event rates depend on software implementation & ARTIQ version. When using direct memory access to play back pre-recorded sequences the highest attained sustained event rate was $\sim 3 \text{ MSs}^{-1}$.

dead time after switching the DAC integrated NCO frequency.

4.7.1.4 Sampler

This extension module provides an 8 channel ADC⁴⁷ with a sample rate of 1.5 MSs^{-1} . Through a variable gain instrumentation amplifier, each channel has a separately programmable sample range which is configurable in factors of ten between $\pm 10 \text{ V}$ and $\pm 0.01 \text{ V}$. We operate with the $\pm 10 \text{ V}$ voltage range which has a -6 dB bandwidth of 200 kHz .

4.7.1.5 Urukul (AD9910)

A 4 channel 1 GSs^{-1} ‘Direct Digital Synthesis’ (DDS) RF source with a -3 dB bandwidth from $< 1 \text{ MHz}$ to $> 400 \text{ MHz}$. The generated RF signal has a frequency resolution of 0.23 Hz . Further, the output amplitude and phase are tunable with 14 and 16 bit resolution respectively. The output power of each channel is individually scaled using a digital step attenuator. Finally, all four channels may be synchronised with respect to each other and the ‘master’ FPGA clock, enabling deterministic phase control.

4.7.1.6 SU-Servo

Unlike the devices listed so far SU-Servo is a combination of multiple EEM modules, providing modified functionality. SU-Servo consists of a Sampler EEM and multiple Urukul EEMs. When operating as SU-Servo, the Sampler sense-voltage controls the Urukul output power via a PI feedback loop. However, as currently implemented, SU-servo does not offer deterministic phase control.

⁴⁷LTC2320-16

In our apparatus, we use SU-Servo to stabilise the laser intensities delivered to the ion by using the sense-voltages provided from the trap table photodiodes (see section 4.6). Via SU-Servo, the photo diode voltages determine the RF-power applied to the AOMs controlling the delivered laser power, enabling closed loop feedback.

4.7.2 Reference Clock

All frequency synthesisers and logic clocks of the realtime control system are derived from a rubidium atomic frequency standard⁴⁸. This standard provides a 10 MHz reference clock. The 125 MHz logic clock for the control system is derived from the 10 MHz clock using a quartz oscillator⁴⁹.

4.8 Amplitude stable RF signal chain

The ~ 65 MHz RF field generating the trap pseudopotential (see section 2.1.1) is generated using an Urukul channel. As we wish to perform MS gates involving the radial ion motion, we require the RF field amplitude experienced by the ion to be sufficiently stable to limit motional drifts to ~ 100 Hz. As we operate with radial mode frequencies of ~ 5.6 MHz, this $\sim 2 \times 10^{-5}$ fractional frequency stability requires a $\sim 2 \times 10^{-5}$ fractional RF amplitude stability. Within the lab environment, Urukul only produces a $\sim 10^{-3}$ fractional amplitude stability, which is well short of the required amplitude stability. To passively improve the amplitude stability, while retaining frequency tunability, we pass the Urukul generated signal through the open-source ‘Squareatron’ amplitude stabilisation board [96]. This PCB passively amplitude stabilises the RF signal by passing it through a low phase noise logic converter⁵⁰. The

⁴⁸SRS FS725

⁴⁹Wenzel Associates, 500-32063

⁵⁰LTC6957IDD-3#PBF

logic converter converts the sinusoidal RF input into a square wave between the board ground and voltage rail. The voltage rail is provided by a low noise linear regulator⁵¹, whose setpoint is regulated using a precision voltage reference⁵². Both these components are supplied from a low noise 5 V linear regulator⁵³. Through the conversion to a square wave, the superior stability of the voltage rail is transferred onto the RF amplitude. The sinusoidal RF signal is recovered using a bandpass filter⁵⁴. To avoid ambient temperature fluctuations and air currents inducing output power fluctuations, the ‘Squareatron’ PCB is enclosed in a made-to-measure metal pillbox enclosing the board. This pillbox is temperature stabilised to within < 1 mK using a Peltier heating/cooling element. The Peltier element is controlled using the open-source Sinara-HW ‘Thermostat’ v2.0 temperature controller using a 10k thermistor to sense the Squareatron temperature. To better stabilise the temperature boundary conditions of the board, the metal pillbox is enclosed in insulating foam which is placed inside a closed metal box. Suitably temperature stabilised, Squareatron produces a low amplitude drift +10 dBm output RF signal with < 10 ppm amplitude stability [97].

As the +10 dBm Squareatron RF power is insufficient to produce the desired pseudopotential, the Squareatron output is amplified⁵⁵. The Squareatron output signal is delivered to the amplifier using low loss coax cable⁵⁶ which is insulated using pipe insulation. The amplifier and its power supply⁵⁷ are similarly temperature stabilised in a made to measure pillbox as outline for Squareatron. Once amplified the RF signal again passes via two insulated

⁵¹LT3042EMSE#PBF

⁵²LTC6655BHMS8-3.3#PBF

⁵³LTC6957IDD-3#PBF

⁵⁴Mini-Circuits, BPF-C45

⁵⁵API, BXMP1042

⁵⁶Mini-circuits, 86-6SM+

⁵⁷Wenzel, LNVR-24-15-5-1

low-loss coax cables⁵⁸ to an RF resonator, boosting the signal voltage, to the ~ 80 V required for the desired trap pseudopotential. The post resonator signal is passed into the vacuum chamber and delivered to the trap RF electrode via the wiring harness and in-vacuum RF PCB, see section 4.1.2. The attained mode frequency stability is demonstrated in figure 4.9. The attained RF amplitude stability is limited by the temperature variation of the RF resonator. Due to mechanical constraints of the vacuum assembly this is not temperature stabilised and only insulated using a nitrogen filled bag.

Using this configuration, the root means square (RMS) deviation of the ~ 5.6 MHz single-ion, in-plane, radial normal mode over a 20 min was measured to be 82 ± 16 Hz.

Some data presented in this work was taken after a major upset of the apparatus. This permanently degraded the RF amplitude stability, details of this upset are presented in appendix B.

4.9 Microwave signal chain

We now detail the setup of the microwave signal chain. The in-vacuum components of this chain are presented in detail by Wolf [63]. Here we will simply note that the microwave signal is delivered to the microwave electrode via a wirebonded PCB routing the signal through the pillbox. Outside the pillbox, the signal is routed via a rigid coax cable to a SMA vacuum feed through, providing electrical connection to the exterior of the vacuum system.

Outside the vacuum chamber, we employ two distinct microwave signal chains, shown in figure 4.10. In the following, the two microwave signal chains are referred to via their intended use as the ‘gate’ and ‘generic’ microwave drives. The ‘generic’ signal chain is suitable for producing tones at frequencies

⁵⁸Mini-circuits, 86-6SM+




Figure 4.9: **Mode frequency drift** - Measured drift of the in-plane, radial, single-ion normal-mode-frequency. The measurement was performed by injecting an RF signal onto one of the ‘DC’ trap electrodes and observing the resulting motional excitation using via resolved Raman sideband spectroscopy. The two presented measurements show the mode frequency drift when ‘Squareatron’ and the RF amplifier are temperature stabilised. The blue measurement series indicates the drift in the absence of additional thermal insulation of the RF resonator board. The improved stability when insulating the RF resonator with a nitrogen filled bag is indicated in red. Notably, insulating the RF resonator removes ~ 600 Hz mode frequency variations on < 10 min timescales, which are associated with the laboratory air conditioning and magnetic field coil water chiller. Comparing the 20 min RMS frequency deviation of the two configurations, we find an improvement from 195 ± 30 Hz to 82 ± 16 Hz from insulating the RF resonator.

anywhere in the 2.4 GHz to 3.3 GHz frequency range, covering all microwave transitions driven in this work. The ‘gate’ chain provides a more linear and higher amplitude signal which is band limited to 3.123 ± 0.010 GHz.

Both the ‘generic’ and ‘gate’ signal chains take an intermediate frequency (IF) signal input, which is produced by one of the two Phaser IQ-pair outputs. To keep the paths of the two components of a given IQ-signal matched, the ‘generic’ drive signals are routed via the two front panel SMA outputs of Phaser. Meanwhile, the ‘gate’ drive signal is routed via the U.FL PCB headers⁵⁹. The two IF IQ-signals are up-converted to the microwave domain using IQ-mixers⁶⁰. The ‘generic’ and ‘gate’ drives respectively use a local oscillator (LO) frequency of 2.85 GHz and 3.0 GHz. These are produced by tabletop microwave synthesisers⁶¹. Synthesiser output powers are set to produce +13 dBm, which saturates the IQ-mixer LO input, while limiting LO leakage.

In each chain, the microwave domain IQ-mixer output is passed through an isolator followed by reflective filters. For the generic drive these filters consist of a 3.8 GHz low pass⁶² and 1.32 GHz high pass filter⁶³. The ‘gate’ drive uses an isolator⁶⁴, and 3.123 ± 0.010 GHz cavity filter⁶⁵. In part of the work the ‘gate’ drive used an additional 3.3 GHz low pass⁶⁶ and 3.0 GHz high pass⁶⁷ filters. The filtered microwave signal is then amplified. For the ‘generic’ drive a Mini-Circuits ZHL-16W-43+ 16 W low noise amplifier is used. The choice

⁵⁹This connection is made using a Female SMA to Female U.FL Coaxial Cable (RS-RPO, 794-2894)

⁶⁰Eclipse, IQ2040MP4

⁶¹The ‘generic’ drive synth is Keysight MXG N5181B, the ‘gate’ drive synth is a Rohde & Schwarz SMA100A

⁶²Mini-Circuits, VLF-3800+

⁶³Mini-Circuits, VHF-1320+

⁶⁴ATM, ATc3-6 circulator with third port 50Ω terminated

⁶⁵ETL Systems, custom product

⁶⁶Mini-Circuits, VLP-41

⁶⁷Mini-Circuits, VHP-26

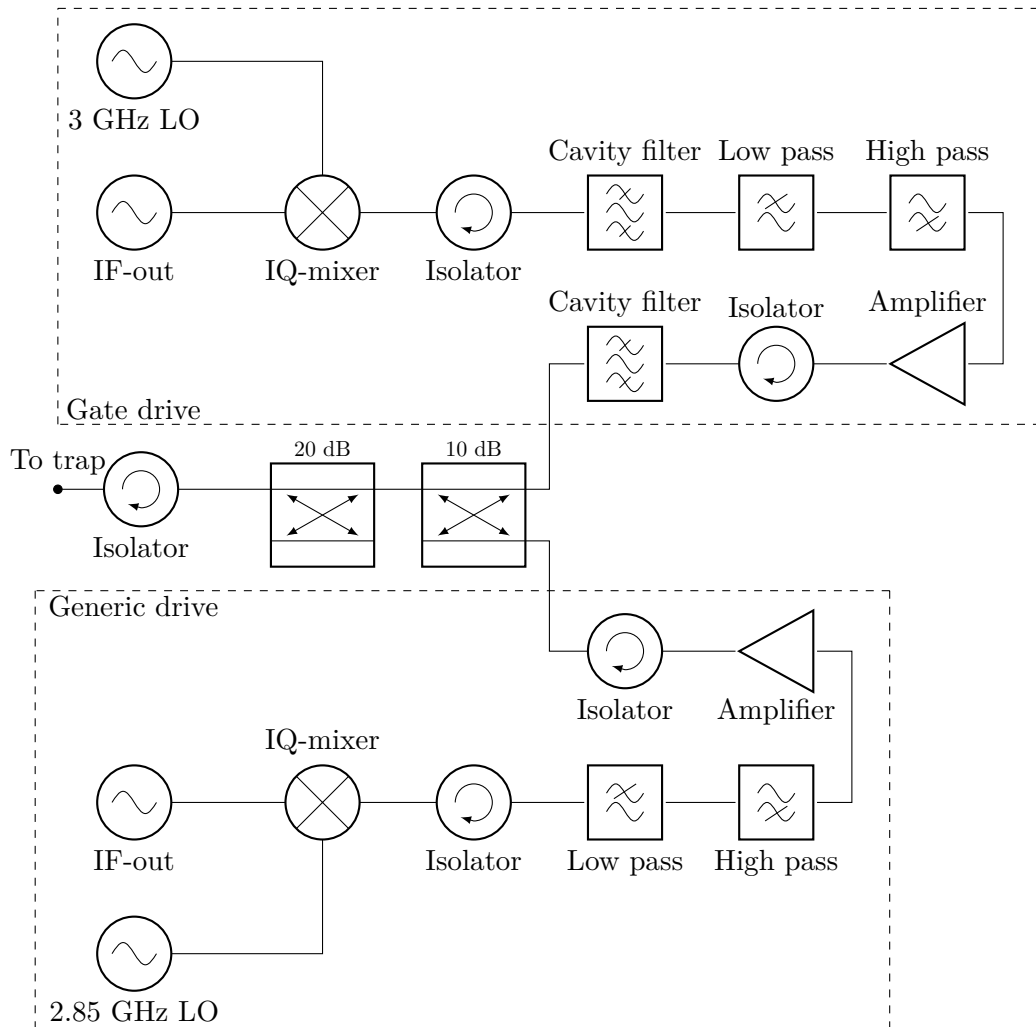


Figure 4.10: **Microwave chain** - The microwave chain consists of two arms which are driven from two independent Phaser quadrature outputs. These two arms are indicated by dashed boxes labelled ‘Gate drive’ and ‘Generic drive’. In each of these arms the intermediate frequency (IF) analogue quadrature signal generated by Phaser is up-converted to the desired microwave frequency using an IQ-mixer and a local oscillator (LO) generated by a bench top frequency synthesiser. After up-conversion, the signal in of each arms passes through an isolator into reflective frequency filters. The isolator directs back reflected signal into a matched termination load. Once spurious tones have been filtered, the microwave signal in both arms is amplified. To protect the amplifiers from potential high power back reflections, the amplifiers are followed by another set of isolators. For the gate drive, the isolator is followed by a second cavity filter, preventing amplifier noise from coupling to qubit spectator transitions. After this second cavity filter, the generic drive signal is merged into the gate drive signal via a 10 dB directional coupler. The 4th port of the coupler, which corresponds to the ‘pass through’ port of the generic drive, is terminated to avoid back reflections. The merged signal passes through a 20 dB directional coupler. The pass-through port of this coupler delivers the microwave signal to the vacuum chamber feed-through. The remaining two ports of the directional coupler are terminated and may be used for diagnostic equipment.

of amplifier for the ‘gate’ drive is more critical. We use a Mini-Circuits HL-100W-352+ 100 W amplifier. This amplifier has a high third-order-intercept (OIP3)⁶⁸, reducing inter-modulation distortion of the multi-tone WMDMS gate signal. After the IQ-mixer the ‘gate’ drive signal strength is ~ -10 dBm. This is primarily limited by the phaser output power. To attain the desired gate speeds from this limited signal strength, we require 50 dB gain. Due to this high amplifier gain, care must be taken to select a low noise amplifier as broadband amplifier noise causes qubit errors. The selected amplifier is specified with a noise figure (NF) < 10 dB. To suppress amplifier noise coupling to qubit spectator transitions, the ‘gate’ drive amplifier is followed by a further isolator⁶⁹ and second 3.123 ± 0.010 GHz cavity filter⁷⁰. Without this post-amplifier cavity filter, gate drive associated broadband microwave noise drives incoherent qubit spectator transitions with a ~ 450 ms decay time constant. Through the addition of this post amplifier filter, the broadband microwave noise no longer couples to spectator transitions and the induced qubit decay time increases to > 20 s.

Having prepared the ‘generic’ and ‘gate’ drive signals as described above, the ‘generic’ signal is coupled into the ‘gate’ signal using a 10 dB coupler. This is followed by a 20 dB coupler providing a pickoff signal which is used for diagnostic purposes. Finally, the combined signal is delivered to the vacuum system using a test and measurement cable⁷¹.

Characterisation measurements for the ‘gate’ drive phase noise and amplitude stability are presented in figures 4.11 and 4.12 respectively.

The characteristics of the in-vacuum components of the microwave chain

⁶⁸Typically +55 dBm

⁶⁹Bocen, custom circulator which was optimised for 2.5 GHz to 3.2 GHz signal frequencies and is rated for 10 W continuous input power. The third port of the circulator was 50Ω terminated.

⁷⁰ETL Systems, custom product

⁷¹Mini-Circuits, CBL-1.5M-SMSM+

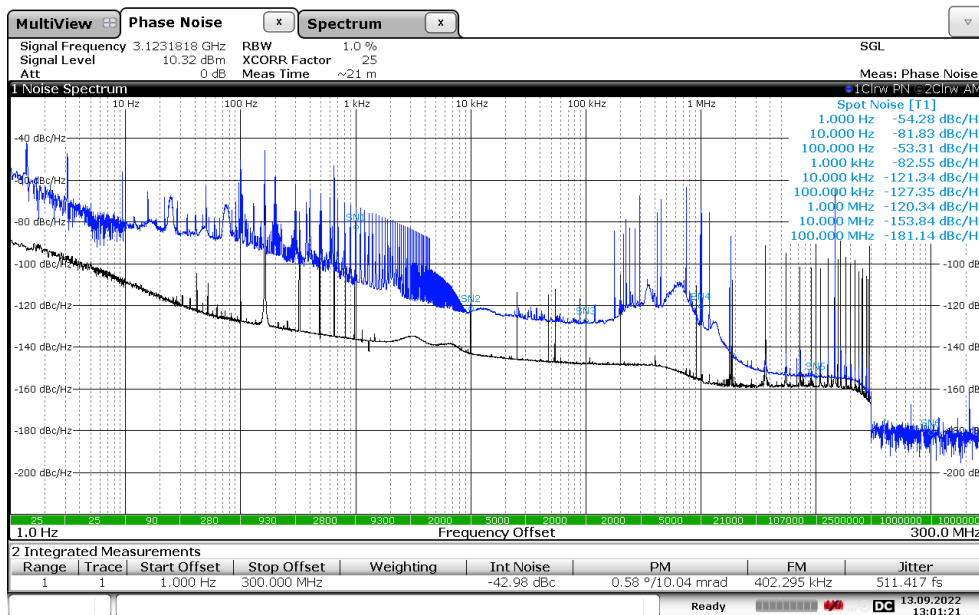


Figure 4.11: **Microwave noise** - Phase and amplitude noise of the ‘gate’ microwave chain measured using R&S FSWP phase noise analyser (referenced to the lab frequency standard, see section 4.7.2). The phase noise measurement is indicated in blue and amplitude modulation noise is indicated in black. The microwave line shape is dominated by phase noise. The majority of phase noise spurs below 10 kHz detuning from the carrier arise due to the lock of the 3.0 GHz Rohde & Schwarz SMA100A local oscillator to the lab frequency standard.

depend on the temperature at which the cryogenic pillbox is operated. Figure 4.13 shows the S_{11} parameter characterising the back reflection from the in-vacuum microwave components for different pillbox temperatures. Notably, the trap resonance shifts with pillbox temperature.

Some data presented in this work was taken after a major upset of the apparatus. This changed the microwave resonator properties, details of this upset are presented in appendix B.




Figure 4.12: **Microwave amplitude drift** - Measured amplitude drift of the microwave ‘gate’ drive. A $100\ \mu\text{s}$ gate drive pulse was repeatedly performed and measured using a temperature stabilised custom designed ADL5511 based RF envelope detector connected to the 20 dB pickoff. The envelope detector provides an output voltage proportional to the microwave amplitude which was measured using a Sampler channel set to unity gain. Each point of the plotted data indicates the measured sense voltage in the centre of the pulse envelop, averaged over 1000 repeats. The error bards indicate the standard deviation of the averaged measurements, which is consistent with the Sampler intrinsic noise. The measurements bounds the fractional microwave amplitude drift on the 20 min timescale to $< 2 \times 10^{-3}$.

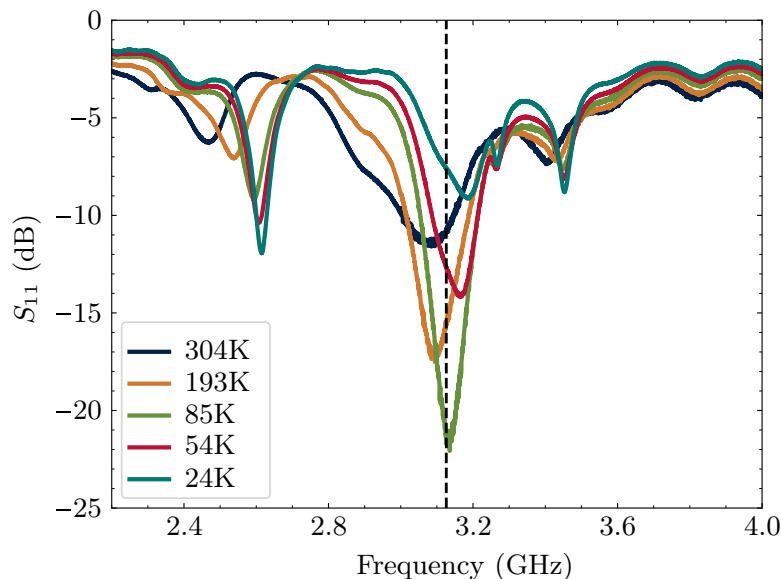


Figure 4.13: **Microwave resonance** - Measured S_{11} parameter (back reflection) of the microwave electrode as a function of microwave frequency. Measurements were performed at a range of pillbox temperatures, which are indicated in the legend. The vertical dashed line indicates the 288 G clock qubit resonance frequency. Notably, the microwave resonance shifts with pillbox temperature. Further, the depth of the reflection minimum varies with temperature. This variation is dominated by the resonator only being critically coupled at ~ 85 K. At higher/lower temperatures the resonator becomes under/over coupled. Figure reproduced from Weber et al. [62].

4.9.1 Microwave duty cycle management

When operating the experiment with the microwave pulse repetition rates and energies used for the demonstrated two-qubit logic gate, significant microwave duty-cycle effects were observed. Besides changing the overall gain of the microwave field experienced by the ion, the duty cycle also significantly impacted the differential gain between the two sideband tones used to drive the gate. Figure 4.14 showcases such differential gain changes when applying microwave pulses with powers and durations comparable to those employed for the demonstrated logic operations. Further, it was observed that adding a microwave pre-pulse of $\sim 400 \mu\text{s}$ during optical state preparation, reduced




Figure 4.14: **Microwave gain duty cycle dependence** - The figure shows the qubit phase accumulated in a spin echo experiment where a variable duration two-tone microwave pulse is applied in one of the arms of the echo sequence. The applied microwave tones are symmetrically detuned by ~ 5 MHz from the qubit frequency with a power of ~ 0.5 W per tone. The power in the blue detuned tone was empirically tuned to cancel the resultant Zeeman shift during short pulse durations. Crucially the accumulated phase scales quadratically with microwave pulse duration. This scaling is inconsistent with a fixed Zeeman shift and was traced to frequency dependent microwave field gain dependence on the applied microwave duty cycle.

microwave amplitude chirping effects on Rabi oscillations.

Consequently, we manage the microwave duty cycle delivered via the ‘gate’ microwave chain. During different phases of the experiment operation we employ one of two duty-cycle management approaches. The first of these approaches performs two microwave pulses matching the expected two-qubit gate tone frequencies and powers. One of these pulses is performed during optical operations before the qubit state is prepared and acts as a pre-pulse.

The second pulse is performed after state-preparation and before readout procedures are started. This pulse may perform logic operations or be sufficiently detuned from resonance frequencies to not drive appreciable dynamics. When changing the duration or power of this second tone the duration of the prepulse is adjusted to maintain a constant time averaged power delivery.

The second duty-cycle management approach simply applies a constant gate drive power. This is split between two microwave tones which are symmetrically detuned from the qubit frequency. The detuning of these tones is similar to those used during the microwave gate. Details of the tune-up procedure of these duty-cycle management procedures are presented in section 6.2.1.

4.10 Ground loop control

The apparatus used within this work makes up a room scale electrical system. Ground loop currents flowing in this system have the potential to generate both electric and magnetic field noise. The magnetic field noise associated with ground loops may degrade the coherence times of qubits. However, the use of a clock qubit in this work mitigates this sufficiently to not be of significant concern. Further ground loops can also result in electric field noise which couples to the collective ion motion via the ion charge. As discussed in chapter 3, the ion motion is leveraged to perform two-qubit gates. Environmental electric fields interacting with this collective ion motion result in logic operation errors. For example, electric field noise near the motional resonance frequency may excite the ion motion and lead to heating from the quantum ground state of the motion. To enable high fidelity logic operations the ground configuration of the apparatus was redesigned resulting in a significant reduction in the heating rate of the ion motion.

4.10.1 Noisy Kasli switch mode power supply

During characterisation work of the Fastino DAC, it was found that the switch mode power supply of the Kasli FPGA board had design flaws resulting in radiated, and common-mode conducted emissions. These emissions occurred at harmonic, and sub-harmonic frequencies of the ~ 450 kHz supply switching frequency. Though radiated emissions could be mitigated by screening and physically distancing sensitive components from Kasli, conducted emissions were found to be critically dependent on the ground configuration of the apparatus. As these conducted emissions are in the critical 100 kHz to 6 MHz frequency range containing resonances of the collective ion motion, a dummy setup was used to establish a ground configuration with satisfactory robustness against common-mode conducted emissions from Kasli. This dummy setup consisted of a Kasli board connected to a single Fastino (DAC) EEM. The analogue output of the Fastino DAC was measured using a Agilent E4405B spectrum analyser which is grounded through its mains connection to the same extension cord powering Kasli. The $50\ \Omega$ spectrum analyser is a proxy for the filter placed before the ‘DC’ trap electrodes⁷². A detailed investigation into noise coupling mechanisms using this setup is presented in appendix C.

4.10.2 Grounding configuration rules

Using the Fastino dummy setup detailed above, the following rules for the experiment ground configuration were empirically determined:

- Isolate Kasli-crate from the mains supply in sensitive frequency ranges.
 - Power Kasli from a class II DC supply⁷³. Typically, at frequencies

⁷²Safety regulations require the ion trap to be grounded, justifying the choice of a grounded load.

⁷³This work uses the XP Power, ALM65US12C2-8

above a few 100 kHz there is significant capacitive coupling across the transformer in a class II supply. We therefore use a common mode current choke⁷⁴ in series with the class II supply.

- Make sure there is a low impedance path to the load ground directly from Kasli board ground.
 - The Kasli front panel is connected to its board ground. We achieve a low impedance ground path by rack mounting Kasli in an electrically conductive crate and rack. This rack is bolted onto a common ground plane with the ion trap.
- Use common mode current chokes on the signal cables and reduce loop area with the ground plane where possible.
 - We employ a variety of clip-on ferrites, ferrite rings and common mode current chokes. The key is to make the common mode impedance large compared to the ground plane.

In addition to the rules above, electrical connection points to the trap table were reduced as far as practical. Remaining connections were bundled and powered through Tacima CS947/BP mains conditioners which contain common-mode noise filtering.

4.10.3 Heating rate improvement

Redesigning the electrical configuration of the experiment to respect the rules outlined above, resulted in a substantial reduction in heating of the ion motion. Before the electrical redesign, heating rates of the single ion ~ 5.6 MHz in-plane radial mode as high as 1.4×10^4 were measured using Raman sideband

⁷⁴Pulse Electronics Power, PH9455.705NL

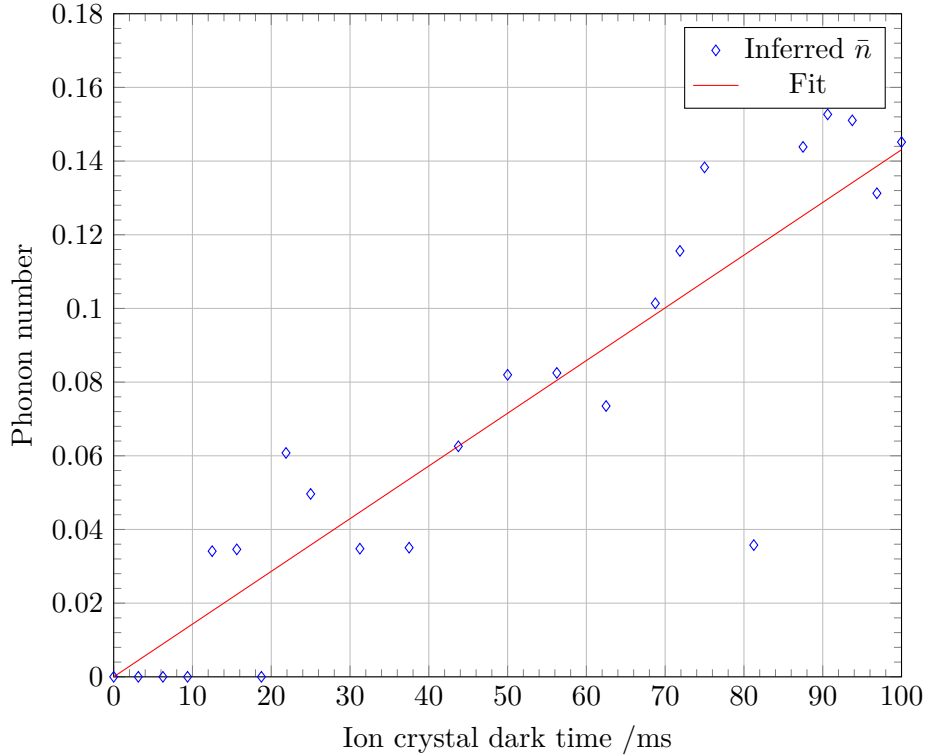


Figure 4.15: **Rocking mode heating rate** - Measured rocking mode phonon number for a range of ion “dark time”, where the ion is left uncooled after ground state cooling. The red line is a linear least-squares fit after eliminating outlier points at ~ 19 ms and ~ 82 ms dark times. This fit gives a $1.431(7)$ quanta per second heating rate.

thermometry. After the redesign this was reduced to a typical single ion in-plane radial mode heating rate of 300 quanta per second.

As the rocking mode of two identical trapped ions correspond to equal and opposite motion of the ions, the rocking mode is intrinsically immune to common electric field noise experienced by the two ions. As only the electric field noise difference between the ions couples to the in-plane rocking mode, we expect its heating rate to be significantly lower. Measuring the heating rate of the in-plane rocking mode we confirm a heating rate of 1 to 3 quanta per second, see figure 4.15. This rocking mode heating rate is sufficiently low not to limit the error of the demonstrated two-qubit gates.

Chapter 5

Time Dependent Trap

Control

From our discussion in 1.3.1, it became apparent that scaling of ion trap based quantum processors is likely to require ion transport capabilities. Besides these scalability considerations, control of the ion trap potential is desirable for a host of applications. Examples include using ions to sense spatial field distributions, tuning the normal mode structure of an ion crystal, and tuning individual ion locations for optical addressing. Recalling our previous discussion in 2.1, the trap potential is produced by combining RF and “static” electric fields. The shape of the potential produced by the RF field is given by the geometry of the RF electrode and is not easily manipulated in our apparatus. However, ion traps currently in use can have over 100 DC electrodes [98], allowing us to shape the trapping potential by applying different voltages to the DC electrodes.

In this thesis, trap-agnostic software tools for calculating trap potential dynamics were developed and tested with both the microwave integrated trap described in 4.1, and a Sandia HOA2 trap [99]. Within the context of the

microwave integrated trap from 4.1, the dynamic ion control was used to map the generated microwave fields, and improve multi ion loading rates through transport loading [100]. We are also planning to use the dynamical potential control to implement single-ion addressing of readout and microwave driven logic operations [81].

5.1 Computational approach & goals

Since the proposal of the QCCD architecture for a scalable ion trap quantum processor [12, 34], generating time dependent trap potentials to transport ions is of increased research interest. In the 20 years since, many ion transport schemes have been experimentally demonstrated [101, 102, 103, 104]. This includes landmark work in Mainz by A. Walther et al., demonstrating diabatic transport of an ion over $280\ \mu\text{m}$ in $3.6\ \mu\text{s}$ without exciting significant ion motion [101]. In the Mainz work Tikhonov regularisation was used to find suitable electrode voltages from a model of the trap fields. The desire to control the potential over the spatial extent of the ion crystal while having a finite number of available trap electrodes for shaping the potential, makes finding suitable electrode voltages an over-specified problem, which is suitable for Tikhonov regularisation. However, despite its successful demonstration and a publicly available implementation [105], some more recent efforts avoid Tikhonov regularization in favour of other approaches [103, 104].

In practice, using Tikhonov regularisation within experimental constraints requires appreciable manual tuning. Tikhonov regularisation does not allow for constrained optimisation. For electrode voltages and voltage slew rates to stay within experimental limitations, a voltage penalising term with per-electrode weights needs to be tuned. Suitable choices of the requisite tuning parameters depend on the desired trajectories and initial voltages [105, 106].

More recent work at ETH Zürich adopts a quadratic programming approach to find electrode voltages that produce the desired potential [104, 106, 107]. The quadratic programming approach finds the desired electrode voltages by minimising a cost function with up to second-order dependence on the electrode voltages. This type of problem is a second-order cone representable optimisation problem, benefiting from existing efficient solvers. The main advantage of the quadratic programming is support of constrained optimisation, allowing us to directly impose electrode voltage and slew rate constraints. As we have introduced a cost function indicating the desired trade-offs in our potential, there are still parameters to tune within the quadratic programming approach. However, these parameters are given by the physical trade-offs we are willing to accept in our potential well and are insensitive to the chosen trajectory. Once set up, the quadratic programming approach requires little tuning to operate with a range of trap trajectories and initial voltages.

Either of these computational approaches are limited by the accuracy of the electrode field model. As discussed by R. Oswald [107], both Tikhonov regularisation and quadratic programming approaches can be enhanced through feedback from experimental data. In the following we will not discuss feedback schemes, focusing our attention on calculating suitable feed-forward electrode voltages.

Despite recent work using quadratic programming approaches, no tools implementing these calculations are widely available. With increasing complexity of ion trap structures and experiments using more ions, there is demand for a tool capable of producing on-demand electrode voltage waveforms without the tuning requirements of Tikhonov regularisation. In the following we develop such a package using the Julia language [108] and the ‘convex.jl’ [109] conic optimisation library. The package is integrated into the open source ARTIQ

control system [94] and experimentally characterised for selected transport operations.

5.1.1 Aims of the computational package

Though it is not feasible to anticipate all upcoming use-cases, the aim of the new package is to make the most common potential manipulations easy to perform. These are: producing a potential well with specific shape, transporting ions, rotating ion crystals and splitting/merging multiple potential wells. Best performance is likely to require tuning and implementation of feedback [107]. However, many experimental use cases don't need best performance. Catering to the latter use case, we aim to produce valid electrode voltages on-demand (after initial tuning of the cost function to suit the experimental tolerances). With on-demand use in mind, the electrode waveform calculation latency should be kept short to avoid inconveniently long wait times when using new waveforms in the lab.

To achieve fast (diabatic) transport, carefully crafted potential well trajectories may be used to control the excitation of the ion motion [105, 110]. However, many experimental use cases are also not sensitive to the detailed potential well trajectory, only requiring controlled evolution between two potential wells. Keeping with our approach of providing on-demand solutions for simple cases, we develop tooling to calculate trajectories that are suitable within the adiabatic transport regime.

5.1.2 Controlled potential parameters

To fully characterise an arbitrary potential well $\phi(\mathbf{r})$, one would need infinitely many parameters. However, only a few properties of the potential well have experimental significance. Beyond the well position, experimentally import-

ant well parameters are the trap compensation, centre-of-mass normal-mode frequencies and normal-mode orientation. The trap compensation consists of forces displacing the ion away from its equilibrium position in the RF-field. Expanding around the RF null, the trap compensation corresponds to the potential gradients $\frac{\partial\phi}{\partial r_i}$. On the other hand, normal-modes are determined by the curvature around the potential minimum $\frac{\partial^2\phi}{\partial r_i\partial r_j}$. Finally, we wish to ensure the potential retains the desired properties over some spatial extent, ensuring a minimum potential depth. Using the commutation of partial derivatives and including the less critical axial compensation, these considerations give us 11 trap parameters we wish to control in static and evolving wells.

As we will see in section 5.5, ion separation requires a quartic trap potential along the direction of the pseudo-potential minimum. In addition to the 11 parameters identified so far, we therefore add tooling to control the 3rd and 4th derivative of the potential along the pseudo-potential axis.

5.2 The trap model

To relate electrode voltages to forces experienced by the ions, we need a model of the force fields within the trap. There are many ways of producing such models including the analytical gapless plane approximation [111], and more generally finite element analysis. As our aim is to develop a trap agnostic tool, we will assume the trap fields can be calculated, and instead focus on how to represent the fields for use with the optimiser.

In many experiments not all forces experienced by the ion can be controlled in real time or with sufficient time resolution. We therefore distinguish between tunable and static force fields. While the latter is outside our control, the former is taken to be proportional to a parameter under our control. Commonly, the pseudo-potential generated by the trap RF-field is taken to

be static while the forces generated by the DC-electrode potentials are tunable.¹ Both static and tunable fields will need to be included in calculations to produce the desired potential landscape. We therefore include both static and tunable force fields in our model.

To not overly restrict the applicability of the tool, we must also adopt a numerical representation of the trap fields. To choose a suitable numerical representation, we need to consider the typical positions explored by the ions and compare them to the characteristic scales of the trap force fields. Common ion trap designs have a trap axis around which a confining pseudo-potential is provided using an RF field (see section 2.1). The ion position along the trap axis can typically be tuned over the entire length of the ion trap electrode array. Usually, this distance is much larger than the trap electrode size. As substantial changes in the force fields occur on the scale of the trap electrodes, we choose to sample the force fields along the trap axis on a grid with small spacing compared to the trap electrode size [104]. Conversely, typical radial ion displacements (in the plane perpendicular to the trap axis) are small compared to the size of the electrode structures and the ion electrode distance. We therefore choose to represent the radial field components by Taylor expanding from the trap axis. Thus, we settle on sampling the trap fields on a one dimensional grid along the trap axis.

As we represent the radial fields with a Taylor expansion, our model needs to include sufficiently many terms to calculate the trap properties we wish to control (see section 5.1.2). In our case, controlling the radial trap frequencies and normal-mode tilt, requires the first order expansion of the 3D force fields. As we sample the field along the trap axis we might consider not including the gradients along the axial direction. However, as we will see in section

¹Though the RF field is often taken to be static, there is no fundamental reason preventing us from treating the RF amplitude as a tunable parameter.

5.5, there are situations in which we wish to Taylor expand along the axial direction. For these cases, we choose to include first order axial derivatives of the radial force components, and up to third order axial derivatives of the axial force component.

The trap model therefore consists of the force per volt generated by the electrodes and the force from the ‘static’ pseudo-potential. These are sampled on multiple points along the trap axis with spacing smaller than the electrode scale. At each sample point the first order expansion of the 3D force field and up to third-order terms along the trap axis are evaluated.

5.3 Calculating electrode voltages

In the preceding section 5.2 we decided on a numerical model of the trap fields. This model can be used to relate electrode voltages to the resulting trap potential. As discussed in 5.1, we adopt a quadratic programming approach to find practical electrode voltage sets that produce the desired trap potential. Quadratic programming requires us to specify a cost function which is minimised by the solver. In the following sections we will explore the cases of finding voltages to produce: a static trap potential, an evolving trap potential, and finally separation of the ions within a single potential well into distinct potential wells. As each of these scenarios has somewhat different requirements we will define different cost functions for each of these.

On physical grounds, it is not possible to produce an arbitrary potential landscape with a given trap design. The resulting trap potentials are therefore a best-effort trap that minimises the chosen cost function.

5.3.1 Static cost-function

The ions in a potential well typically sit in a chain along the trap axis. For experiments with a very large number of ions, this chain may become long compared to the electrode size. Consequently, a Taylor expansion may not be suitable for describing the variation of the trap fields along the axial direction. Analogously to our trap model considerations in 5.2, we consider the trap fields at multiple locations along the trap axis [104], while using Taylor expansion in the radial directions.

Developing a cost function, we start by considering a single well n . For a given axial position sample i we can choose to define a cost function as shown in equation 5.1. This cost function controls the forces and the force gradients the ion experiences, see section 5.1. Note that we are defining the axial direction to be along the z -axis.

$$\begin{aligned}
C_{\text{static}, n, i} = & c_0 (M_{x, i} - T_{x, n, i})^2 + c_1 (M_{y, i} - T_{y, n, i})^2 & (5.1) \\
& + c_2 (M_{z, i} - T_{z, n, i})^2 + c_3 \left(\frac{\partial M_{x, i}}{\partial x} - \frac{\partial T_{x, n, i}}{\partial x} \right)^2 \\
& + c_4 \left(\frac{\partial M_{x, i}}{\partial y} - \frac{\partial T_{x, n, i}}{\partial y} \right)^2 + c_5 \left(\frac{\partial M_{y, i}}{\partial y} - \frac{\partial T_{y, n, i}}{\partial y} \right)^2
\end{aligned}$$

In this equation $T_{x/y/z, n, i}$ represents the target force that the potential-well n should produce on an ion located at the axial position z_i of the i^{th} field sample. Meanwhile, $M_{x/y/z, i}$ represents the force on the ion that is produced by the electrode voltages within the trap model (see section 5.2). The values c_{0-5} are constants that may be used to tune the relative importance of different terms in the cost-function (see appendix E). The indices $x/y/z$ indicate the spatial components of the forces. By including only one of the derivative cross-terms, we assume the forces to be produced by a conservative potential. Further, the cost function does not include z -derivatives of the forces. Such

terms are unnecessary as we are sampling the field at multiple points along the trap axis. Another key distinction of this cost-function is that it does not operate on the potential, but on the forces experienced by the ions. This distinction eliminates the arbitrary potential offsets which are also optimised in preceding works.

Comparing different axial positions, we note that not all axial field sample positions are equally important. The purpose of the potential is to confine ions and define the properties of their experienced potential well. We therefore do not care about the fields far from the location of ions. We include this in our cost function by multiplying a weight factor $w_{n,i}$ with the cost function from each axial position sample $C_{\text{static},n,i}$. Depending on the use case, the best choice of $w_{n,i}$ may vary. Simple choices are a top-hat function or a Gaussian around the trap centre. The advantage of choosing Gaussian weights is that forces at the ion location are weighted most heavily, while the tails of the distribution ensure the potential well has a minimum width and depth [104]. In practice, the details of this weight function only matter when the well size approaches the electrode size.

Finally, we may wish to avoid the solver using electrode voltages near the experiment hardware voltage limits. We therefore include a term penalising electrode voltage deviation from a specified default voltage (V_{def}). From these considerations we arrive at the final cost function for producing static potential wells, which is shown in 5.2.

$$C_{\text{static}} = \sum_{n,i} (w_{n,i} C_{\text{static},n,i}) + r_0 \sum_j (V_j - V_{\text{def}})^2 \quad (5.2)$$

In this expression the sums over n , i and j respectively are over the specified potential wells, axial field samples and the electrodes. V_j is the voltage of the j^{th} electrode and r_0 is a constant to set the scale of the voltage term.

r_0 is somewhat redundant as all other terms already include a scale factor. It is included nonetheless as, in practice, the solver benefits from all terms being around unity in the vicinity of the solution.

5.3.2 Computation speed optimisation

As we aim to build a tool useful for on-demand electrode voltage calculations, we should be mindful to not impose impractically long calculation latencies. With equation 5.2 in mind, we note that the contribution to the cost function by an axial position sample depends on the weight associated with that sample and target well. In many common situations the field at most locations within the ion trap only contributes negligibly to the overall cost-function. We can therefore improve return latencies by discarding all cost terms with a sufficiently small weight factor.

5.4 Evolving potential wells

Besides crafting potential wells with specific properties, another important capability in upcoming ion trap experiments is the ability to evolve the fields experienced by the ion in a controlled way, see 1.3.1. Examples include: rotating ion crystals, smoothly evolving trap frequencies, and transporting ions between different locations.

For sufficiently slow changes in the trap potential it is possible to find a suitable sequence of electrode voltages using the static cost function described in 5.3.1. However, as faster operations are performed, experimental limitations on the electrode voltage rate of change become significant. To avoid creating large voltage steps between voltage sets, we are forced to consider the voltage trajectory as a whole [104].

5.4.1 Dynamic cost-function

The static cost function we introduced in 5.3.1 only considers the target potential at a single instant and has no notion of the voltage evolution. Commonly, trap electrode voltages are provided by a DAC. As the output of a DAC is updated in discrete time steps, it is natural for us to adopt a discretised notion of time, where the voltage trajectory is divided into N_{step} discrete voltage steps. Within this picture we can extend the definition of $C_{\text{static},n,i}$, the static cost function of well n at location i , from equation 5.1 to become $C_{\text{static},n,i,t}$, the static cost function of well n at location i at discrete time step t . The extended definition is shown in equation 5.3, which is a relabelling of equation 5.1.

$$\begin{aligned}
C_{\text{static},n,i,t} = & c_0 (M_{x,i,t} - T_{x,n,i,t})^2 + c_1 (M_{y,i,t} - T_{y,n,i,t})^2 \\
& + c_2 (M_{z,i,t} - T_{z,n,i,t})^2 + c_3 \left(\frac{\partial M_{x,i,t}}{\partial x} - \frac{\partial T_{x,n,i,t}}{\partial x} \right)^2 \\
& + c_4 \left(\frac{\partial M_{x,i,t}}{\partial y} - \frac{\partial T_{x,n,i,t}}{\partial y} \right)^2 + c_5 \left(\frac{\partial M_{y,i,t}}{\partial y} - \frac{\partial T_{y,n,i,t}}{\partial y} \right)^2
\end{aligned} \tag{5.3}$$

Continuing analogously to section 5.3.1 and the work by V. Negnevitsky [104] we define the cost function for the dynamical solver as shown in equation 5.4. As the position and size of the potential wells are not constant during well evolution, this definition allows for time dependent weighting $w_{n,i,t}$ of the cost term $C_{\text{static},n,i,t}$.

$$\begin{aligned}
C_{\text{dynamic}} = & \sum_{n,i,t} (w_{n,i,t} C_{\text{static},n,i,t}) + r_0 \sum_j \sum_t^{N_{\text{step}}} (V_{j,t} - V_{\text{def}})^2 \\
& + r_1 \sum_j \sum_t^{N_{\text{step}}-1} (V_{j,t+1} - V_{j,t})^2 + r_2 \sum_j \sum_t^{N_{\text{step}}-2} (V_{j,t+2} - 2V_{j,t+1} + V_{j,t})^2
\end{aligned} \tag{5.4}$$

As before the indices n , i and j represent the n^{th} potential well, i^{th} axial field sample and the j^{th} electrode respectively. The newly introduced index t indicates the time step out of N_{step} steps in the evolution. Crucially, the voltage $V_{j,t}$ represents the voltage present on electrode j at time step t . As the electrodes are typically low-pass filtered, $V_{j,t}$ only corresponds to the DAC output voltage when voltage changes are slow compared to the filter cut-off frequency. Options for compensating for the trap filters or directly optimising the DAC output voltages are discussed in appendix D. We are only looking to perform dynamics at timescales approaching the filter cut-off (see 5.1). It is therefore sufficient for us to select low electrode voltage slew rate waveforms by including two terms penalising electrode voltage slew rate and slew rate changes². In our cost function equation, these are scaled by the constants r_1 and r_2 respectively.

This cost function can be used as-is, finding electrode waveforms that produce the desired wells throughout the evolution. In practice, we wish to compose multiple transport sequences, smoothly continuing from the final electrode voltages of the previous evolution. We may also wish to make different trade-offs between well parameters during transport than are desired in the initial/final well of the trajectory. We therefore choose to constrain the initial and final electrode voltages in the waveform to be fixed to user specified values. Within the context of our tool box, these voltages may be calculated using the static solver.

5.4.2 Well trajectories

To use the dynamical cost function detailed in section 5.4.1 we need to provide the solver with $T_{x/y/z,n,i,t}$, the target force on the ion for each potential

²For inductive trap filters, sudden slew rate changes are associated with electrode voltage ringing

well, time step and field sample location. For advanced transport schemes this may be complicated. However, as discussed in 5.1.1 many use cases are insensitive to the transport details between a start and end potential. As a final component to performing potential well evolution within an ion trap, we should therefore provide tooling to find adiabatic potential well parameter trajectories between specified initial and final potential wells. There are 3 objectives for our chosen potential well trajectories, these are listed in order of priority.

1. Reliability: the ion should remain within the desired potential well. Ions leaving their potential well usually necessitates lengthy ion loading procedures and may lead to experimental errors.
2. Speed: the evolution should be fast compared to other experimental operations. Within the bounds set by the adiabatic limit and the electrode filters, we wish to avoid well evolution dominating the experimental duty-cycle. We therefore desire potential well trajectories that perform reliable transport in as few time steps as possible.
3. Avoid motional excitation: In many contexts of ion trap operation motional excitation of the ion is undesirable. We therefore wish to limit the motional excitation produced by the potential well dynamics. Producing excessive ion motion requires us to include additional re-cooling steps, increasing the experimental duty-cycle associated with well evolution.

Fortunately, as we are operating within the adiabatic limit and below the trap filter cut-off, these goals go hand in hand. Loss of the ion from the target potential is often associated with background gas collisions of the ion. By reducing the time spent evolving the well, the probability of such a collision occurring during the transport operation is reduced. The time spent evolving

the well also places a limit on the motional excitation as electric field noise continuously heats the ion motion. Finally, the ion leaving the desired potential is an extreme form of motional disturbance. It stands to reason that trajectories giving low motional excitation should leave the ion less vulnerable to ion-loss related reliability issues. However, there is a trade-off between the evolution speed of trap parameters and motional excitation. Accelerating the trap potential results in a change of the equilibrium position of the ion within the potential. Fast changes in the well acceleration (jerk) therefore excite oscillations of the ions within the well (see section 5.3.2 in [104] for the resulting displacement operator). Transport can be tuned to cancel the excited motion [101, 110]. However, as we discussed in section 5.1.1, we adopt the simpler adiabatic approach where we avoid exciting motion in the first place. This requires us to evolve the ion equilibrium position acceleration slowly compared to the trap frequency. Suitable examples of the ion equilibrium position trajectory include \sin^2 and polynomial ramps (see [101, 112]).

Though polynomial ramps minimise the jerk in a trajectory, we adopt the trajectory presented for axial ion transport by V. Negnevitsky [104]. This potential well trajectory balances smooth ion acceleration with a constant velocity around the trajectory centre. The constant velocity section is desirable as it reduces voltage slew rate requirements for a given transport time. The potential well parameter $p(t)$ trajectory between times t_0 and t_f is shown in equation 5.5 and figure 5.1.

$$p_n(t) - p_n(t_0) = [p_n(t_f) - p_n(t_0)] P \left(\frac{t - t_0}{t_f - t_0} \right) \quad (5.5)$$

$$P(x) = \frac{\rho(x) - \rho(0)}{\rho(1) - \rho(0)} \quad (5.6)$$

$$\rho(x) = \ln \left| \frac{\exp(ab(2x-1)) - ie^{-a}}{\exp(ab(2x-1)) - ie^a} \right| \quad (5.7)$$

Using this trajectory, the trade-off between smooth acceleration and constant velocity may be tuned using the a and b parameters. Here we use the values $a = 3.0$ and $b = 1.5$ which were found to be suitable by V. Negnevitsky [104]. We extend the use of this trajectory for all controlled parameters of the potential well described in 5.1.2. The only exception to the smooth evolution is the target well width. As the well width does not accelerate the ion position, we choose a linear trajectory for this parameter.

5.5 Ion separation

With the approach laid out in section 5.4 we enable nearly all operations we desire to perform (see section 5.1.1). However, our choice of cost-function and trajectory in section 5.4 is not well suited for ion separation and merging³ of potential wells. In the following we develop a separate cost-function and trajectory for separating potential wells. Noting that merging is reverse separation, this approach also allows us to combine multiple wells.

5.5.1 Ion-separation trajectory

During the separation procedure of ions within a potential well, a potential barrier needs to be raised between the ions inside the potential well. This

³Merging is the combination of multiple potential wells and their contained ions into a single potential well.

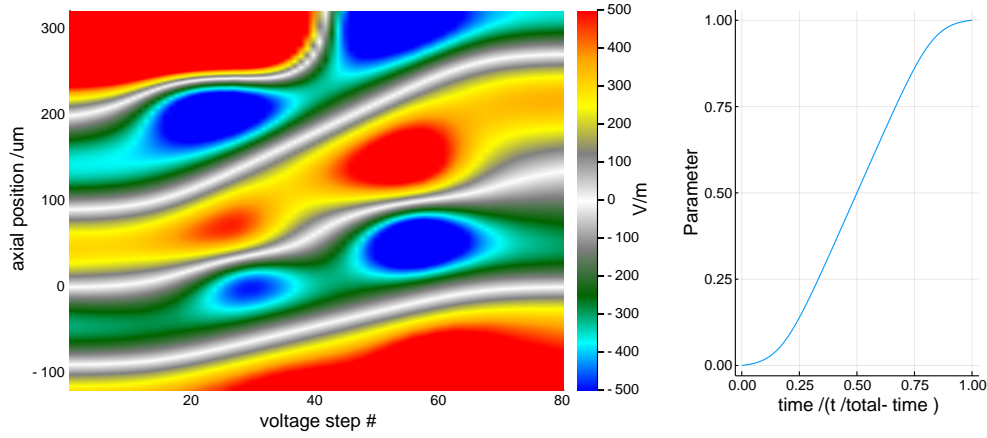


Figure 5.1: **Transport trajectories** - Left: Heat-map showing the axial force-field during simultaneous transport of two potential wells in the microwave integrated trap. The x-axis shows the time index of the voltage step. The y-axis shows the position within the trap. The two potential minima are indicated by the white lines flanked above by dark green and below by yellow. The opposite ordering corresponds to potential maxima. In the figure, two potential wells are initially located at $\pm 90 \mu\text{m}$ from the trap centre. The wells are then transported such that one is located in the trap centre and the other is located in the far trap zone at $271.55 \mu\text{m}$. Right: Shape of the smooth parameter trajectory described in equation 5.5. The function accelerates and decelerates smoothly, while maintaining a high velocity during evolution.

procedure is described by a quartic potential transitioning from a single anharmonic potential well to a double-well [104, 112, 113].

Considering a quartic potential energy per unit charge as given in equation 5.8, two positive, singly charged ions of mass m within the potential will be separated by a distance s and have a centre-of-mass normal-mode angular frequency ω . The separation s and angular frequency ω are given by equations 5.9 and 5.10 respectively [113].

$$V(z) = \alpha z^2 + \beta z^4 + \gamma z + \delta z^3 + V_0 \quad (5.8)$$

$$\beta s^5 + 2\alpha s^3 = \frac{e}{2\pi\epsilon_0} \quad (5.9)$$

$$\omega^2 = 2\alpha + 3\beta s^2 \frac{e}{m} \quad (5.10)$$

For separating two ions in a well γ is usually chosen to be zero. Commonly, the “axial compensation” δ is tuned experimentally to compensate stray axial forces not included within the trap model. For a single, infinitely deep, quartic confining well, we have $\alpha, \beta \geq 0$. This transitions to an infinitely deep double-well as α becomes negative. As the well transitions from a single to a double well, the axial normal-mode frequency is set by the quartic term β . As low confinement frequencies leave the ions vulnerable to motional excitation by jerk and electric field noise, we wish to maximise β while the wells are being separated [112, 113]. None the less, the attainable axial normal-mode frequencies during ion separation are significantly lower than are otherwise achievable. It is therefore especially important to choose fast and low jerk ion trajectories. We use a polynomial ion-separation trajectory $s(t)$ between times t_0 and t_f as is shown in equation 5.11 [112].

$$s(t) = s(t_0) + (s(t_0) - s(t_f)) \left(-10 \left(\frac{t - t_0}{t_f - t_0} \right)^3 + 15 \left(\frac{t - t_0}{t_f - t_0} \right)^4 - 6 \left(\frac{t - t_0}{t_f - t_0} \right)^5 \right) \quad (5.11)$$

Notably, the ion is stationary at the start and end of this trajectory. During ion separation, all controlled trap parameters from section 5.1.2 that do not contribute to the ion separation procedure are held constant. Figure 5.3 shows the calculated potential evolution, ion separation and jerk in such a trajectory. The resulting fields and trajectories are shown in figures 5.2 & 5.4.

At the end of this splitting trajectory the ion separation approaches the electrode scale and strong axial confinement is restored.




Figure 5.2: **Single to double well transition** - Sequence of 21 potentials resulting from the voltage calculated for ion separation in the microwave integrated trap. To aid the eye, the potentials at different times are shown in different colours. For ion separation the well starts as a near quadratic potential which smoothly evolves into a double-well. As the potential offset changes substantially during splitting, the potential at the trap centre has been subtracted at each time step.

5.5.2 Separation algorithm

During our splitting trajectory we wish to maximise the axial quartic confinement of the ions. The cost function laid out in section 5.4.1 is not suited to the task of maximising the axial curvature. We therefore need to use a modified cost function for potential wells undergoing ion separation. We note that, during the critical phases of splitting, the ion separation is much smaller than the electrode scale. We therefore adopt a Taylor expansion approach for all fields of a separating well. Generalising from V. Negnevitsky [104], equation 5.12 shows the cost function $C_{\text{sep},t}$ of the well being separated at time step t .




Figure 5.3: **Ion-separation evolution** - (top) evolution of the ion–ion distance during the single to double well transition. (centre) The jerk (rate of acceleration change) applied to the equilibrium ion position. (bottom) Frequency of the centre-of-mass axial normal-mode for $^{43}\text{Ca}^+$ ions during the single to double-well transition. The lower this normal mode frequency is, the more sensitive the ion is to jerk. For suitable choices of the well curvature α at the start and end of the separation procedure, the jerk zero crossing is near the minimum normal-mode frequency.




Figure 5.4: **Ion-separation force-field** - Heat-map showing the axial force-field during separation of a single potential well into two potential wells. The x-axis show the time index of the voltage step. The y-axis shows the position within the trap. Potential minima are indicated by the white lines flanked above by dark green and below by yellow. The opposite ordering corresponds to potential maxima. In the figure, a potential produced by the static solver is placed at the trap centre ($z = 0$). In the first 50 steps the static solver electrode voltages are interpolated to the first separation solver voltages. Steps 50 to 150 show the evolution calculated by the separation solver. Note the creation of a potential maximum in the well centre between steps 50 and 150. The voltage steps 150 to 200 show the interpolation from the final splitting solver electrode voltages to the electrode voltages calculated by the static solver for two wells located at $\pm 30 \mu\text{m}$ from the trap centre.

$$\begin{aligned}
C_{\text{sep},t} = & \sum_{\epsilon \in \{x,y,z\}} a_{\epsilon} (M_{\epsilon,t} - T_{\epsilon,t})^2 \\
& + \sum_{\epsilon \in \{x,y,z\}} \sum_{\rho \in \{x,y,z\}} b_{\epsilon\rho} \left(\frac{\partial M_{\epsilon,t}}{\partial \rho} - \frac{\partial T_{\epsilon,t}}{\partial \rho} \right)^2 \\
& + c \left(\frac{\partial^2 M_{z,t}}{\partial z^2} - \frac{\partial^2 T_{z,t}}{\partial z^2} \right)^2 + d \frac{\partial^3 M_{z,t}}{\partial z^3}
\end{aligned} \tag{5.12}$$

In this equation $T_{x/y/z,t}$ represents the target force in the $x/y/z$ -direction at time step t . Meanwhile, $M_{x/y/z,i}$ represents the $x/y/z$ force component produced by the electrode voltages within the trap model (see section 5.2). Both T and M are defined as the force acting on an ion located at the trap centre. The values a_{ϵ} , $b_{\epsilon\rho}$, c , and d are constants that may be used to tune the relative importance of different terms in the cost-function.⁴

In general, ion separation may be performed in the presence of other “spectator” potential wells. These are held fixed during ion separation and can therefore be represented by the static solver cost-function C_{static} presented in equation 5.2. At a given time step t , we define a cost-function suitable for ion separation in the presence of spectator wells by summing $C_{\text{sep},t}$ and the static cost-functions associated with the spectator wells $C_{\text{static},n,i}$ (see equation 5.1).

$$C_{\text{sep-multi-well},t} = C_{\text{sep},t} + \sum_{n,i} (w_{n,i} C_{\text{static},n,i}) \tag{5.13}$$

In this expression the sums over n and i respectively are over the spectator potential wells and axial field samples. $w_{n,i}$ is a weight for each spectator well over the spectator well field-sample positions.

As the achievable axial potential quartic confinement (β in equation 5.8) is only known after optimisation via the cost function, our discussion of the ion separation trajectory (section 5.5.1) does not provide us with the required

⁴As $\frac{\partial T_{\epsilon,t}}{\partial \rho}$ and $\frac{\partial M_{\epsilon,t}}{\partial \rho}$ are symmetric in ϵ and ρ there are only six physically significant parameters in $b_{\epsilon\rho}$

axial force gradient $\frac{\partial T_{z,t}}{\partial z}$ (α in equation 5.8). The desired axial force gradient depends on the achievable axial potential quartic confinement (β in equation 5.8). The latter is unknown at the time of calculating the cost function. However, re-arranging equation 5.9 we know that α may be calculated from β and the target separation s as shown in equation 5.14.

$$\alpha = \frac{e}{4\pi\epsilon_0 s^3} - \beta \frac{s^2}{2} \quad (5.14)$$

Identifying $-2\alpha e = \frac{\partial T_{z,t}}{\partial z}$ and $-24\beta e = \frac{\partial^3 M_{z,t}}{\partial z^3}$, we could perform a substitution in 5.12. As equation 5.14 is linear in β , the resulting cost function would still be compatible with quadratic programming. This approach would allow us to proceed similarly to the treatment in section 5.4.1.

However, for the current implementation we follow the interpolation approach taken by V. Negnevitsky [104]. This approach uses the cost function shown in equation 5.15, where we have added a voltage penalising term to equation 5.13.

$$C_{\text{sep-sample},t} = C_{\text{sep},t} + \sum_{n,i} (w_{n,i} C_{\text{static},n,i}) + r_0 \sum_j (V_j - V_{\text{def}})^2 \quad (5.15)$$

V_j is the voltage of the j^{th} electrode and V_{def} a user specified default voltage. r_0 is a constant to set the scale of the voltage term.

This cost function can be used to find a sequence of static wells with a range of values for $\frac{\partial T_{z,t}}{\partial z}$. For monotonic progression of the ion separation, a 4th order spline is used to determine the electrode voltage as a function of ion separation. With $V_j(s)$ in-hand, we know the electrode voltage waveform for the ion separation trajectory. An ion separation waveform calculated using the spline interpolation method is shown in figure 5.4.

Adopting this spline interpolation approach has 2 major drawbacks:

1. It does not allow us to simultaneously evolve independent parameters.

To do so, would require us to find $V_j(\rho_0, \dots, \rho_i)$ with ρ_i being the i^{th}

independent parameter. Establishing this multidimensional function is impractical.

2. We are unable to discourage high voltage slew rates as each trajectory sample is calculated independently.

Due to these drawbacks, it is desirable to adopt the substitution approach in further development of the package.

Finally, the potential wells confining the ions at the start and end of the ion separation procedure are optimised to maximise the curvature around the separation centre. This potential is dissimilar to potential wells produced by the usual cost function for well evolution, see 5.4.1. Interpolating between suitably placed static solver potential wells and the splitting potential allows a smooth transition between these different potential landscapes [104]. In our case we adopt a polynomial interpolation between the static and separation solver electrode voltages at the start and end of the trajectory. Using interpolation of the same form as the separation trajectory in equation 5.11 gives suitably smooth transitions. Figure 5.4 shows the evolution of the axial forces during this interpolation.

Putting everything together, figure 5.5 shows the force field evolution during ion transport of multiple wells and well separation in the presence of a spectator well.

5.6 Characterisation

Having implemented a toolbox that calculates trap electrode voltages based on the considerations presented in sections 5.1 to 5.5, we demonstrate the capabilities of the tool in several ion trap experiments. To this end, a driver and mediator for the open-source ARTIQ control system were developed [94].




Figure 5.5: **Demonstration of toolbox capabilities** - Heat-map showing the axial force field during evolution of multiple potential wells. The x-axis show the time index of the voltage step. The y-axis shows the position within the trap. Potential minima are indicated by the white lines flanked above by green and below by yellow. The opposite ordering corresponds to potential maxima. In the figure, two potential wells are initially located at $\pm 90 \mu\text{m}$ from the trap centre. In the first 100 time steps, the wells are transported such that one is located in the trap centre and the other is located in the far trap zone at $271.55 \mu\text{m}$. During this transport the axial mode frequency of the two modes is increased smoothly. The following 200 voltage steps are associated with separation of the central potential well. Steps 100 to 150 show the interpolation from the last transport electrode voltages to the separation solver voltages. Voltage steps 150 to 250 show the separation of the central potential well in the presence of the fixed spectator well at $271.55 \mu\text{m}$. In the final 50 voltage steps of the separation procedure, the separation solver electrode voltages are interpolated to the electrode voltages calculated by the static solver for two wells located $\pm 30 \mu\text{m}$ from the trap centre. Finally, voltage steps 300 to 400 show transport of the two created wells to $\pm 90 \mu\text{m}$. During this transport, the confining frequency of the well at $271.55 \mu\text{m}$ is ramped up further.

Given the trap agnostic nature of the toolbox and the wide spread use of ARTIQ, the toolbox, driver, and mediator are suitable for on-demand calculations within many ion trap experiment setups and is now in use within three ion trap experiments using different ion trap designs.

We now demonstrate ion separation and transport in the microwave integrated trap described in section 4.1, and a Sandia HOA2 trap [99]. As discussed in section 5.4.2, we desire ion dynamics to be reliable, fast, and avoid motional excitation. However, sensitivity to these factors varies by use case. As our use cases are insensitive to the latter two points the characterisation focuses on the operation reliability. Further, these two setups were not designed with ion dynamics in mind. This limited the characterisations that could be performed within these setups.

- When the characterisation was performed within the HOA2 trap, speed was limited by a low update rate DAC and ion thermometry was unavailable. Consequently only the reliability could be characterised.
- The microwave integrated trap is designed for cryogenic operation. However, the toolbox was characterised with the trap operating at room temperature. This limited measurements of radial motion excitation and precluded measurement of axial motion due to excessive environmental ion heating.

5.6.1 Separation and transport reliability

The reliability of ion separation and transport was demonstrated by combining ion separation and transport in a single experiment. Two ions were laser cooled within a single well. This was followed by separating the two ions into distinct wells. The resulting two potential wells are transported through the length of

the trap. The wells may then be brought back to the trap centre individually, allowing the number of ions in each well to be read out. Finally, the wells are placed symmetrically around the trap centre and merged. At this point the total number of ions is detected and the ions are re-cooled. During this procedure two types of error are possible. Firstly, an ion may enter the wrong potential well during separation. Secondly, an ion may escape the potential well.

The following section presents the observations within the two ion traps. It is worth stressing that these tests were performed with limited tuning of the transport parameters on the experiment. The number of transport steps were chosen to be sensible, but were not optimised. Cost function weights were determined using the rules of thumb presented in appendix E. The only experimental tuning was to find a suitable axial compensation that resulted in the two ions entering separate wells, see figure 5.6.

5.6.1.1 HOA2

After separation of the cooled ions, one well was moved to the trap centre for readout, while the other was moved 210 μm to the far end of the available electrodes. After counting the ions in the central well, the wells were merged and the total ion number detected. No errors were observed when performing 10^4 repeats of this sequence. This gives a 3.7×10^{-4} upper bound of the splitting error probability with 95% confidence.

5.6.1.2 Microwave integrated trap

Due to issues with two-ion readout when the characterisation was performed, the scheme used in the HOA2 trap was modified for use in the microwave integrated trap. After separation of the cooled ions, one well was moved to




Figure 5.6: **Axial compensation tuning** - Measured axial compensation calibration data in the microwave integrated trap. The well separation protocol was applied to a single trapped ion. The y-axis indicates the probability of finding the ion in the well which evolves towards the negative-z direction. The x-axis indicates the axial compensation field used in the separation procedure. There is a sharp transition from the ion being in the negative-z well to being found in the positive-z well at $E_z \approx 31.75 \text{ V m}^{-1}$. The axial compensation identified using this single-ion method was then used for two-ion separation characterisation.

the trap origin for readout, while the other was moved to $z = 271.55 \mu\text{m}$ at the far end of the available electrodes. After detecting the presence of ions in the central well, the two wells are transported back to the position they had pre-splitting. From there the well previously located at $z = 271.55 \mu\text{m}$ is transported to the trap centre. Simultaneously, the remaining well is transported along the trap axis to $z = -271.55 \mu\text{m}$. The time taken for the entire splitting and transport trajectory was around 1.6 ms. This is faster than the multi-ms cooling time, achieving the aim of not dominating the experiment duty cycle. As previously noted, this transport time was not optimised. Repeating this sequence 2.3×10^4 times, no ion separation errors were observed. This gives a 1.6×10^{-4} upper bound of the splitting error probability with 95% confidence. Meanwhile 2 ion-loss events occurred in 2.5×10^4 repeats. This rate of ion loss is consistent with the two-ion lifetime within the microwave integrated trap at room temperature, thus there is no evidence that the loss rate is increased by the transport operations.

5.6.2 Radial mode heating during transport

Within the microwave integrated trap there are plans to perform single qubit addressing by adjusting the trap potential. If this is to be interleaved with two qubit gates, it is important for trap evolution not to cause excessive excitation of the radial ion motion (see 3.2). We establish a bound on the radial normal-mode excitation by shuttling a single ion from the trap centre ($z = 0$) along the trap axis to a ‘far’ well position. Once at the far position, there is a $350 \mu\text{s}$ delay after which, the ion is then transported back to the central position.

To distinguish motional excitation caused by transport from other heating sources, the motional excitation for fixed duration transport over a range of transport distances is measured. Using 100 voltage steps at the minimum




Figure 5.7: **Radial excitation from transport** - Measured radial mode excitation when performing $39.2\ \mu\text{s}$ duration transport over a range of distances within the microwave integrated trap. The x-axis indicates the position of the far well to which the ion is transported in $39.2\ \mu\text{s}$ from the trap centre at $z = 0$. Once at the far position, the ion is left in place for $350\ \mu\text{s}$ after which it is transported back to the trap centre with a second $39.2\ \mu\text{s}$ transport a) Average motional temperature as estimated from Raman sideband thermometry. This method assumes the motion to be in a thermal state and may therefore be misleading when measuring non-thermal excitations that may be caused by transport. Irrespective of the transport distance, Raman sideband thermometry gives average occupation numbers of $\bar{n} \approx 0.8$ quanta. b) Measured transition probability driven by the red sideband pulse of the Raman sideband thermometry measurement. The Rabi frequency of red sideband transitions between the motional Fock states $|n\rangle$ and $|n+1\rangle$ is given by $\Omega_{n,n+1} = \sqrt{n+1}\Omega_{0,1}$ [12]. The red sideband transition probability is therefore given by a superposition of different frequency oscillations determined by the Fock state superposition. The transition probability driven by the red Raman sideband varies between ~ 0.18 and ~ 0.26 over the range of transport distances. This $\sim \sqrt{2}$ variation in the transition probability suggests that any coherent excitation of the ion motion is on the single phonon level.

Fastino sample update interval of 392 ns results in a 39.2 μ s transport duration. Within the microwave integrated trap, this transport duration presents a reasonable trade off between transport speed and resulting motional excitation. Figure 5.7 shows the measured motional excitation when performing such a 39.2 μ s transport to and from a range of far well positions. The measurement indicates that the radial motion experiences some excitation on the single phonon level. Radial motional excitation from transport (without tuning of compensations or solver parameters) is therefore bound on the single phonon level.

Chapter 6

Two-Qubit Gate

Implementation

We now discuss implementing two-qubit logic gates within the apparatus discussed in chapter 4. This includes both the sequence of operations performed in a gate characterisation experiment (section 6.1) and the tune-up procedure used to configure relevant experimental parameters (section 6.2).

6.1 Gate sequence

In this section, we discuss the experimental sequence of operations used to perform and characterise two-qubit logic gates. Besides performing the desired two-qubit logic operations as presented in chapter 3, a characterisation experiment must also perform state preparation, readout, and laser cooling as discussed in chapter 2. Furthermore, our discussion of the experimental apparatus in section 4.6 established a requirement for repeated laser intensity stabilisation throughout long running experiments. Finally, our discussion of the microwave delivery apparatus in sections 4.7.1.3 & 4.9.1 indicated that the

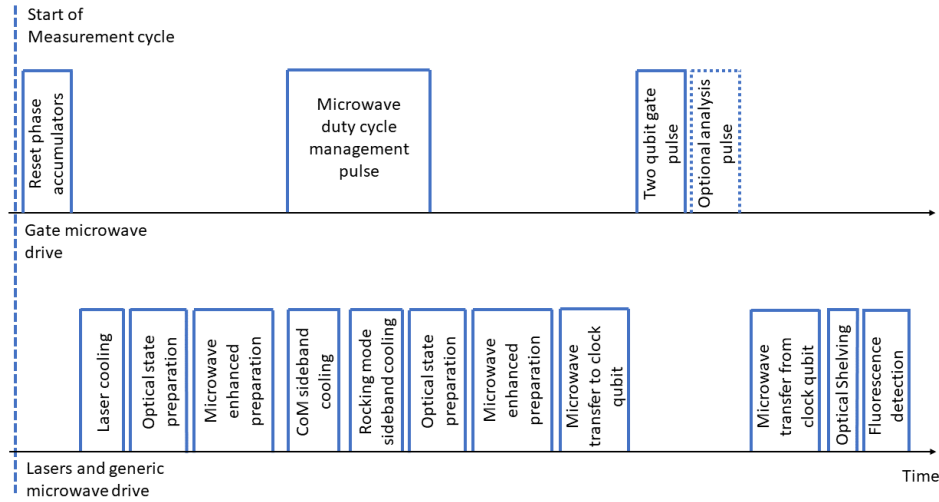


Figure 6.1: **Gate measurement sequence** - The above pulse sequence indicates the order of operations within the 10 ms gate-measurement experiment cycle. The dashed line indicates the start of a new cycle. As the microwave ‘gate’ drive duty cycle management pre-pulse is concurrent with laser and ‘generic’ drive operations, ‘gate’ drive operations are shown on a separated timeline to aid visual clarity. Characterisation of the entangled ion state produced by the microwave gate, requires sequences with and without an additional $\pi/2$ analysis pulse. This pulse is not present in all measurement sequences and indicated by a dotted box. Note that the width of boxes is not to scale with the duration of the indicated operations.

timing and power delivery of microwave pulses must be managed throughout the experimental sequence.

Figure 6.1 shows an overview of the measurement sequence employed when characterising the two-qubit gates; the timing of this sequence is discussed in more detail below.

6.1.1 Cycle alignment

To start any given measurement cycle within the experimental sequence we must ensure that the pulse sequence has deterministic alignment to the 320 ns digital communication cycle between the “Phaser” AWG and “Kasli” control-

system FPGA. Further, we wish to control the microwave power delivery to the surface trap. To simplify the microwave power management, we determine a fixed cycle repetition time of 10 ms, which is enforced across all measurement sequences. As discussed in 4.9.1, the fixed cycle time allows us to manage the average microwave power delivery by controlling the energy delivered in each cycle, see 6.1.7.

The 10 ms duty cycle is an integer multiple of the 320 ns Phaser-Kasli communication cycle¹, ensuring that alignment with this cycle is maintained throughout an experiment. To ensure deterministic alignment to the 320 ns communication cycle between experiments, the start of the 10 ms cycle is set to have a fixed delay after the response arrival time when reading a Phaser register. We choose to start the first 10 ms measurement cycle 20 μ s after receipt of the register content.

Finally, to obtain a deterministic phase relationship between the numerically controlled oscillators (NCOs) of Phaser, the phase accumulator of each NCO is cleared at the start of each 10 ms cycle.

6.1.2 Laser cooling

The measurement cycle is started with a 3 ms dark resonance laser cooling pulse as described in subsection 2.4.2. After a 10 μ s delay, the ion is optically pumped to the $4S_{1/2}, |F = 4, m = +4\rangle$ state as described in section 2.5. This microwave enhanced optical pumping procedure lasts ~ 1.2 ms. With the ion thus prepared, pulsed Raman resolved sideband cooling is employed to cool the ion to the motional ground state. Both the radial in-plane centre-of-mass and rocking modes are cooled sequentially. First, the centre of mass mode is cooled using 20 cycles of spectrally-addressed, 8 μ s Raman-sideband and 20 μ s

¹10 ms/320 ns = 31250

397 nm σ^+ optical pumping pulses. Next, the in-plane rocking mode is cooled with 30 cycles of similarly addressed pulses of equal pulse duration. Including control system related dead times, the resolved sideband cooling procedure of both modes requires ~ 1.5 ms.

6.1.3 State preparation

The optical pumping employed during resolved sideband cooling is optimised for cooling speed rather than state preparation fidelity. We therefore follow the resolved sideband cooling with a second round of microwave enhanced optical state preparation (see section 2.5.1) occupying a further ~ 1.2 ms of the measurement cycle. Within the microwave enhanced state preparation sequence, 500 μ s are spent on initial optical pumping with the EOM enabled. As five cycles of microwave enhancement were employed, microwave pulses occupy a further ~ 10 μ s and 397 nm σ^+ pulses with the EOM disabled require a further 20 μ s. The remainder of the microwave enhanced state preparation time is accounted for by dead times.

Having prepared the $4S_{1/2}, |F = 4, m = +4\rangle$ state, we transfer the ion population to the $4S_{1/2}, |F = 3, m = +1\rangle$ clock qubit state using the sequence of π -pulses described in section 2.5.2. This procedure takes ~ 0.08 ms and is dominated by inter-pulse dead times (see section 4.7.1.3). With the transfer pulses completed, the ion preparation is complete and quantum logic operations may be performed on the clock qubit. However, control system limitations necessitate a further 20 μ s delay before clock qubit operations are commenced.

6.1.4 Readout

Depending on the desired clock qubit logic operations, the time required before the qubit state may be read out varies. As stressed in section 6.1.1, we wish to

keep deterministic alignment to the 320 ns Kasli-Phaser communication cycle. Further, our discussion of microwave duty cycle effects in section 4.9.1 raised concerns of Rabi frequency variations throughout the measurement sequence. As will become apparent from our duty cycle procedure discussed in section 6.1.7, we satisfy both deterministic Kasli-Phaser cycle alignment and alleviate Rabi frequency variation concerns by allowing a fixed microwave operation time of ~ 0.7 ms . The ~ 0.7 ms operation time was chosen to allow sufficient time for any desired microwave operations while not imposing an overly long delay before reading back results.

The readout sequence starts with up to two optional clock qubit $\pi/2$ -pulses. These are phase coherent with the microwave logic operations and may be used for calibration (section 6.2.4) and parity analysis (section 3.2.2). After this step the readout procedure outlined in section 2.6 is applied. The readout procedure starts with a sequence of microwave pulses which map the qubit states into the $4S_{1/2}$, $|F = 3, m = +1\rangle$, and $|F = 4, m = +4\rangle$ states. The microwave transfer requires ~ 0.09 ms² and is dominated by inter-pulse dead times (see section 4.7.1.3). Once the qubit state re-mapping is complete, optical shelving is performed. To achieve high fidelity optical shelving we use ten 850 nm re-pumping cycles. The 60 μ s shelving pulse duration is evenly split across eleven 393 nm pulses. The ten 850 nm σ^+ , and π pulses respectively occupy 10 μ s, and 6 μ s respectively. Including dead times, the shelving procedure requires ~ 0.11 ms .

Once the ion is shelved, we allow a 5 μ s dead time before fluorescence detection is commenced. The exact duration of fluorescence detection operation is subject to calibration. For two-ion readout detection times are approximately 0.5 ms.

²The additional transfer duration when compared to state-preparation is due to the use of a Ramsey sequence on transition ‘c’.

6.1.5 Clock qubit operations

During the clock qubit operation phase of the measurement sequence we perform a MS style gate described in section 3. As several gates with different gate durations and modulation schemes are presented in this work, the timing of the gate pulses necessarily varies. We reserve presentation of the timing characteristics of each gate for the discussion of the gate performance in chapter 7. Here we present the general structure and relative timing of different aspects of the implemented WMDMS gates. Gates performed without dynamical decoupling follow the same scheme, but have the decoupling drive amplitude set to zero.

The implemented WMDMS gates may be broken down into three phases. At the start of the WMDMS gate, no microwave field is applied to the ion. As discussed in section 3.3.1, we wish to pulse shape the microwave sideband tones driving the motional interaction which underpins the MS gate. In the first stage of the microwave gate the sideband amplitude is smoothly increased over the rise time t_{edge} . During this stage the amplitude envelope of the sideband tones is increased along a $\sin^2\left(\frac{\pi}{2} \frac{t}{t_{\text{edge}}}\right)$ trajectory. During this “ramp up” period, no dynamical decoupling is applied. The next stage of the implemented WMDMS gate maintains a constant sideband tone amplitude throughout its duration t_{high} . During this stage of the gate Walsh modulated dynamical decoupling is implemented. When implementing the dynamical decoupling, care must be taken that all dynamical decoupling pulses maintain the same phase alignment to the $t_{\text{cycle}} = 320$ ns Kasli-Phaser communication cycle. As time resolution of the Phaser AWG gateway is limited to $t_{\text{step}} = 40$ ns (see section 4.7.1.3), the following procedure is used to determine suitable decoupling tone timing and dead-times when performing n dynamical decoupling pulses:

1. The inter-decoupling pulse $t_{\text{inter dd pad}}$ delay is set to the shortest time not resulting in significant ringing of the Phaser FIR filters $t_{\text{inter dd pad}} = 4t_{\text{step}}$.
2. Determine the padding time t_{pad} needed before/after the decoupling pulse sequence via $\frac{t_{\text{pad}}}{t_{\text{step}}} = \frac{t_{\text{high}} + t_{\text{inter dd pad}}}{t_{\text{step}}} \bmod \left(\frac{nt_{\text{cycle}}}{t_{\text{step}}} \right)$. This padding time is split equally between the start and end of this stage of the WMDMS gate. In case of t_{pad} not being an even multiple of t_{step} , the padding at the start/end of the decoupling pulses are rounded down/up respectively.
3. The final duration of each dynamical decoupling pulse $t_{\text{dd pulse}}$ is given by $t_{\text{dd pulse}} = \frac{t_{\text{high}} + t_{\text{inter dd pad}} - t_{\text{pad}}}{n} - t_{\text{inter dd pad}}$.
4. Finally, we wish to avoid ringing of the Phaser FIR filters. This is achieved by bandwidth limiting the decoupling pulse amplitude changes. Decoupling tone amplitude changes are therefore ramped into/out of over the time $t_{\text{dd edge}} = 3t_{\text{step}}$. This ramp time is included in the above definition of $t_{\text{dd pulse}}$, but sets a technical limitation of $t_{\text{dd pulse}} \geq 6t_{\text{step}}$.

The third and final stage of the WMDMS pulse ramps down the sideband amplitude similarly to the ramp up in the first stage detailed above. No dynamical decoupling is performed during this stage of the WMDMS gate.

6.1.6 Housekeeping operations

Beyond the 10 ms measurement sequence detailed above, we interrupt the measurement cycle every 800 sequence repetitions to perform ion loss detection and laser intensity stabilisation. Laser intensity stabilisation is discussed in section 4.6. Ion loss detection is performed by preparing the ions in the $4S_{1/2}$ ground state and observing their fluorescence similarly to state readout. If

ion loss has occurred, data since the last ion number verification is discarded and the experiment sequence is terminated.

6.1.7 Duty-cycle management

As noted in our preceding discussion of the apparatus in section 4.9.1, it is critical to manage the microwave duty-cycle delivered to the ion trap. Mismanagement of the microwave duty-cycle leads to thermal transients affecting trap properties. Such thermal transients are large enough to be experimentally relevant and too slow to conveniently wait out. We therefore employ two distinct duty cycle management approaches during the measurement sequence and “idle” microwave periods such as is the case during housekeeping operations detailed in section 6.1.6 and at times between experimental runs.

As we choose to fix the repetition rate of the experiment measurement sequence at 100 Hz, we may control the overall microwave power delivered to the trap by matching microwave delivery in all measurement cycles. As we may wish to vary the two qubit gate duration or powers, we choose to add a microwave pre-pulse before the clock qubit state is initialised. Due to the pre-pulse occurring before qubit initialisation, it does not perform qubit operations. This pre-pulse consists of the three microwave tones used in the WMDMS gate. The amplitudes of each tone in the pre-pulse are set to those expected for the WMDMS gate. Changes in the logic operation amplitude or gate duration may now be compensated by suitably scaling the pre-pulse duration to give the same overall power delivery. To protect from non-linearities in the trap power dissipation, amplitude deviations from the pre-pulse amplitude are kept small. The timing of the microwave pre-pulse is chosen to coincide with initial optical state preparation and Raman sideband cooling (see section 6.1.2). This allows the pre-pulse to occur shortly before the logic pulse while

not interfering with the final state-preparation procedure, see section 4.9.1.

Complementary to the approach we have outlined for the measurement cycle, the microwave duty cycle was maintained during microwave “idle” times by applying the red and blue detuned microwave sideband tones continuously. A suitable “idle” power delivery was ensured by using an empirically tuned equivalent sideband amplitude, see section 6.2.1.

6.2 Calibration procedure

For successful gate operation a multitude of experimental parameters must be tuned. To perform gate operations we followed the tune-up procedure presented in this section. As many steps in the tune-up procedure depend on completion of preceding steps, we present the tune-up operations in chronological order.

6.2.1 Duty-cycle management

The gate tune-up procedure starts by deciding the microwave input power at which gates are to be performed. Combined with prior measurements of the microwave field of the trap [62] this allows a rough determination of the MS gate duration with a given number of phase-space loops. Based on this gate duration estimate and the desire to maintain at least a 400 μs microwave pre-pulse (see section 4.9.1), we determine the microwave duty time within the measurement loop outlined in 6.1. Though the exact value of the chosen duty time is non-critical, the duty time is chosen to minimally fit a gate with twice as many phase-space loops as the target gate while maintaining a 400 μs pre-pulse. A suitable amplitude of the resonant decoupling tone is determined using the heuristics presented in section 3.3.4. Finally, a coarse calibration of




Figure 6.2: **Pillbox temperature stability** - Measured pillbox temperature as measured via the pillbox resistive temperature detector (RTD). The measurement shows the pillbox temperature over two hours when performing the WMDMS gate presented in chapter 7.

state preparation and readout microwave pulses is performed based on the expected microwave field.

Having chosen the powers and durations of the microwave pulses within the measurement sequence, the “idle” microwave duty-cycle management is empirically tuned to match the chosen measurement sequence duty-cycle. We perform an initial estimate of the microwave power dissipated within the trap by assuming quadratic scaling of the Phaser amplitude scale with dissipated microwave power. This initial estimate is corrected using the temperature of the cryogenic pillbox as measured by the pillbox RTD, see section 4.1.2. The iterative “idle” microwave power correction is carried out as follows:

1. Use the RTD mounted to the pillbox to measure the steady state temperature T_0 when no microwaves are applied to the trap.
2. Run the gate measurement sequence until the cryogenic pillbox reaches the measurement sequence steady state temperature T_m .
3. Apply the “idle” microwave power to the pillbox and note down the associated steady state temperature T_i
4. Correct the microwave amplitude y applied during idle periods according to the first order expansion $y_{\text{new}} = y_{\text{old}} \left(1 - 0.5 \frac{T_i - T_m}{T_i - T_0}\right)$. The factor of 0.5 arises from the quadratic scaling of power with amplitude microwave amplitude.
5. Apply the new “idle” microwave power to the pillbox and note down the associated steady state temperature T_i
6. If the pillbox temperature deviation is not within tolerance, continue from point 4.

This sequence is continued until the pillbox temperature deviation is satisfactorily small. When operating the cryostat at 24.6 K, employing sideband tones of approximately 2 W per sideband and 8% duty cycle results in $T_m - T_0 \approx 200$ mK. Using the tuned “idle” power, the resulting steady state temperature difference between T_m and T_i can be suppressed below the 1 mK quantisation limit of our temperature measurement. Long term temperature stability with this duty cycle management in place is within 10 mK, see figure 6.2.

6.2.2 Pulse shaping

The desired pulse shaped edges of the microwave sideband tones for the two qubit gate are determined such that sufficient suppression of off-resonant excitations is achieved. Further, the pulse edges of the resonant decoupling drive are set such that ringing of the Phaser FIR filters is suppressed.

6.2.3 Transfer pulse calibration

After fixing the microwave duty cycle properties, we calibrate the microwave pulses used within state preparation (see section 2.5) and readout (see section 2.6). The calibration is performed on a single trapped ion by measuring Rabi oscillations on the different hyperfine transitions. The transfer pulses are calibrated with an experimental sequence following the sequence outlined in section 6.1. The sequence is performed as it would be when performing a two qubit gate with the duration estimated in section 6.2.1. To avoid systematic transfer errors due to different thermal environments in different stages of the microwave cycle, the state-preparation and readout pulses are calibrated independently. Further, to keep sequence timing consistent, the pulse durations of each microwave pulse are fixed. Rabi oscillations are observed by scanning the amplitude of the fixed duration pulses. To ensure good readout contrast, the amplitude of the complimentary state preparation/readout pulse on the transition being calibrated is set to zero.

6.2.4 Analysis pulse calibration

After calibrating the microwave transfer pulses, the gate analysis pulse is calibrated by preparing a single ion in the clock qubit. Similarly to the transfer pulse calibration outlined above, the gate characterisation sequence, including the expected two qubit gate pulse, are applied. The gate pulse is applied with

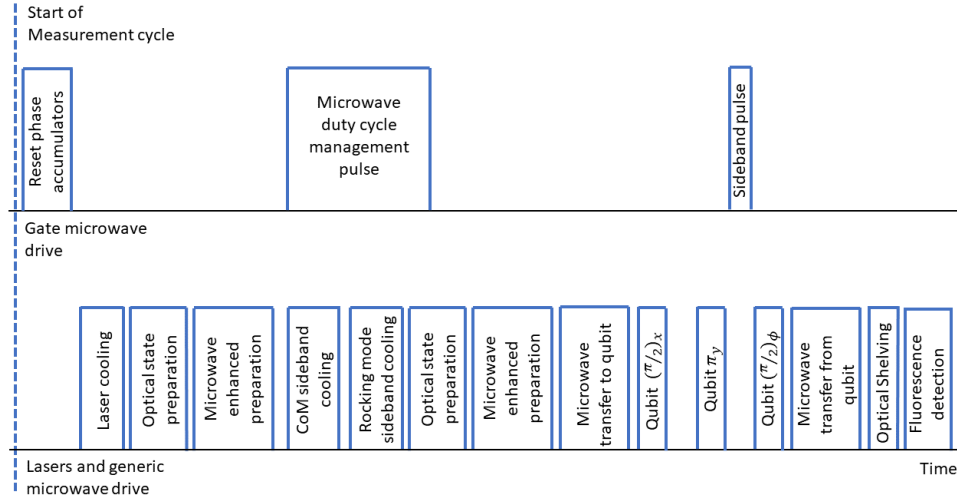


Figure 6.3: **AC-Zeeman shift measurement sequence** - The above pulse sequence indicates the order of operations within the 10 ms spin-echo experiment cycle. The dashed line indicates the start of a new cycle. As the microwave ‘gate’ drive duty cycle management pre-pulse is concurrent with laser and ‘generic’ drive operations, ‘gate drive operations are shown on a separated timeline to aid visual clarity. The spin-echo microwave pulses are generated using the ‘generic’ microwave chain. The phases of these pulses are indicated by a subscript. Note that the width of boxes is not to scale with the duration of the indicated operations.

no decoupling tone and the microwave sideband tones are detuned from the single ion motional mode, such that negligible dynamics are driven. The usual gate characterisation sequence is modified by applying two identical, in-phase analysis pulses instead of the single analysis pulse used during gate characterisation, see section 6.1.4. The analysis pulse is calibrated by scanning the amplitude of the two analysis pulses to produce an overall π qubit rotation.

6.2.5 Microwave field characterisation and AC-Zeeman shift

As the microwave field experienced by the ion is sensitive to the applied duty cycle, we must measure the microwave fields experienced at this duty cycle to determine the expected AC-Zeeman shift caused by the two qubit gate tones.

The AC-Zeeman shift is caused by the combined effect of the σ and π polarised microwave fields. Notably, the shift caused by the π polarised microwave field of each sideband tone is much larger than the shift resulting from the σ polarised microwaves. However, due to the equal and opposite microwave detuning, there is near perfect cancellation of the π field AC-Zeeman shifts caused by the two microwave sideband tones. Consequently, any measurement of the AC-Zeeman shift is critically dependent on the amplitude balancing of the microwave sideband tones experienced by the ion. To avoid sensitivity to power balancing errors, we apply a single microwave sideband tone, while managing the microwave duty cycle as described in section 6.1.7. This sideband tone is applied with the same amplitude and comparable frequency settings as in the two-qubit gate which is being tuned up. As this measurement is performed on a single ion, which does not have a radial rocking mode, the sideband tone does not drive appreciable sideband dynamics. We measure the resulting AC-Zeeman shift on transitions ‘a’, ‘b’, ‘c’, ‘q’, ‘e’, and ‘f’ (see section 2.3.2) by varying the experimental sequence described in section 6.1. Specifically, we omit the state-preparation/readout pulses which are not required to prepare/read out the ion state of the desired transition. We then perform a spin echo sequence on the transition being interrogated where the single sideband tone is enabled for varying duration in one half of the spin-echo sequence, see figure 6.3. Finally, the state is read out using a sequence of microwave pulses which map the two transition states into the $|F = 4, m = +4\rangle$ and a convenient $F = 3$ state.

By fitting to the measured AC-Zeeman shifts of transitions ‘a’, ‘b’, ‘c’, ‘q’, ‘e’, and ‘f’, it is possible to ascertain the σ and π components of the microwave amplitude experienced by the ion, see figure 6.4. By assuming that the σ to π amplitude ratio is not significantly different for the two sideband tones, we




Figure 6.4: **Zeeman shift fit** - Residuals of two-parameter fit to the measured, single-sideband Zeeman shifts, as measured before performing the highest fidelity demonstrated gate. The different measured transitions are indicated by the labels on the x-axis. The fractional deviation of the measured shift from the fitted shift is indicated on the y-axis. The error bars indicate the statistical uncertainty of the measured Zeeman shift. The fitted Zeeman shift is consistent within the $\sim 1\%$ measurement uncertainty of the AC-Zeeman shift of the different transitions.

now calculate the AC-Zeeman shift expected when the second sideband tone is added.

6.2.6 Sideband amplitude balancing

During the presentation of the microwave field characterisation, we noted that the observed AC-Zeeman shift is sensitive to the amplitude balancing of the sideband tones used for the gate. Following the field characterisation presented in section 6.2.5, we calculated the AC-Zeeman shift expected when sideband




Figure 6.5: **Sideband amplitude Rabi spectroscopy** - Calibration measurement determining the AC-Zeeman shift and balancing the microwave sideband amplitudes. The measurement uses a fixed amplitude red sideband tone and scans the amplitude of the blue sideband tone. The red and blue sidebands are symmetrically detuned around the qubit frequency accounting for the expected AC-Zeeman shift, see section 6.2.5. A third microwave is applied at this expected qubit frequency. The pulse duration of the three tones is set to match the expected gate duration. The amplitude of the central third tone is set to drive approximately a π pulse over the pulse duration when the qubit is resonant. The y-axis indicates the probability of detecting the ion in the prepared qubit state. The x-axis indicates the amplitude scale factor of the blue sideband which is proportional to the generated blue sideband amplitude. The red line indicates a parabolic least squares fit to the measured data. The fit gives the blue sideband amplitude scale factor resulting in the expected qubit frequency as 0.404770(3).

amplitudes are balanced. We now leverage the known AC-Zeeman shift to balance the microwave sideband powers using Rabi spectroscopy on a single ion.

The Rabi spectroscopy experiment follows the gate characterisation sequence outlined in section 6.1. However, during the clock qubit logic phase of the sequence, we choose not to apply a WMDMS gate. Instead, we apply the two gate sideband tones and a weak non-modulated ‘carrier’ tone half way between the frequencies of the two sideband tones. The duration of this pulse is matched to the expected two-qubit gate duration. The red sideband tone amplitude is fixed to the one calibrated in section 6.2.5. The blue sideband amplitude is coarsely set to be comparable to the red sideband amplitude. Finally, the ‘carrier’ tone amplitude is set to produce approximately a π pulse over the duration of the three-tone pulse. The set of these three tones is detuned from the qubit transition frequency by the AC-Zeeman shift calculated in section 6.2.5. The sideband power is now balanced by tuning blue sideband power such that the AC-Zeeman shifted qubit frequency becomes resonant with the applied carrier tone, see figure 6.5. As the Lamb-Dicke parameter is of order $\tilde{\eta} \approx 10^{-3}$, the AC-Zeeman shift generated by the sideband tones is $\sim 10^6$ times larger than the shift generated by a resonant microwave tone driving Rabi oscillations with frequency comparable to the gate duration. Similarly, for our chosen dynamical decoupling powers during the WMDMS gate we have $t_{\text{gate}}\Omega_{Q,n}^{\text{DD}} \approx 2\pi \times 60$ resulting in a factor $\sim 10^3$ difference in AC-Zeeman shift caused by the decoupling tone compared to the shift from the microwave sidebands. Such small frequency errors are inconsequential to the gate operation. Consequently, this method of balancing the sidebands has the benefit of also controlling the average AC-Zeeman shift across the gate pulse.




Figure 6.6: **WMDMS gate sideband detuning scan** - Qubit readout probability when scanning the sideband detuning. Both ions are prepared in the bright state before application of a two loop, $331 \mu\text{s}$ WMDMS gate and no analysis pulse is performed. Measured readout probabilities and standard error intervals for detecting two, one, or no bright ions after application of the WMDMS pulse. The gate detuning for the two-loop $331 \mu\text{s}$ WMDMS pulse is at the $P_{db} + P_{bd}$ minimum near $+5.55 \text{ kHz}$ and is indicated by the dotted line.

6.2.7 Dynamical decoupling drive phase

After balancing the microwave sideband powers, the remainder of the tune-up procedure is performed using two ions. The first of the calibrations performed with two ions is the phase calibration of the dynamical decoupling tone. The decoupling tone only commutes with the gate dynamics when it matches the average phase of the sidebands. As the decoupling term $\Omega_{Q,n}^{\text{DD}}$ is significantly stronger than the driven motional dynamics $\Omega'_{Q,n}$ ($\Omega_{Q,n}^{\text{DD}} \gg \Omega'_{Q,n}$), mismatch

of the phase suppresses the gate dynamics.

Calibration of the decoupling drive phase uses the gate characterisation sequence outlined in section 6.1. The sequence is performed with the desired WMDMS gate and without any analysis pulses. Using the coarsely tuned gate duration, the sideband detuning giving the largest qubit transition probability is found, see figure 6.6. Using this gate detuning the phase of the dynamical decoupling drive is scanned. The resulting scan is symmetric around the desired decoupling drive phase, allowing the matched decoupling phase to be found by symmetry, see figure 6.7.

6.2.8 Fine gate duration adjustment & sideband detuning

We now fine-tune the two-qubit gate duration based on the driven qubit dynamics. Similarly to the calibration of the decoupling drive phase in section 6.2.7, we scan the microwave sideband tone detuning from the motional mode without analysis pulses. However, instead of focusing on the point of maximum qubit transfer probability, we focus on the expected two-qubit gate point. As both qubits are both initially prepared into the $|F = 3, m_f = +1\rangle$ state, an ideal gate should only produce two bright or dark ions. These two outcomes are expected to be equally probable, giving $P_{11} = P_{00} = 0.5$. The crossing point where $P_{11} = P_{00}$ should therefore coincide with the minimum probability of measuring a single bright ion $P_{10} + P_{01}$. From simulation, gate duration errors shift the relative position of the $P_{11} = P_{00}$ crossing point relative to the $P_{10} + P_{01}$ minimum in sideband detuning scans, see figure 6.8. Fine adjustment of the gate duration is performed by iteratively adjusting the gate pulse duration to overlap $P_{11} = P_{00}$ and the $P_{10} + P_{01}$ minimum.

The sideband frequency at which the $P_{11} = P_{00} \approx 0.5$ crossing and $P_{10} + P_{01}$ minimum occur is determined to be the optimal sideband frequency for gates.

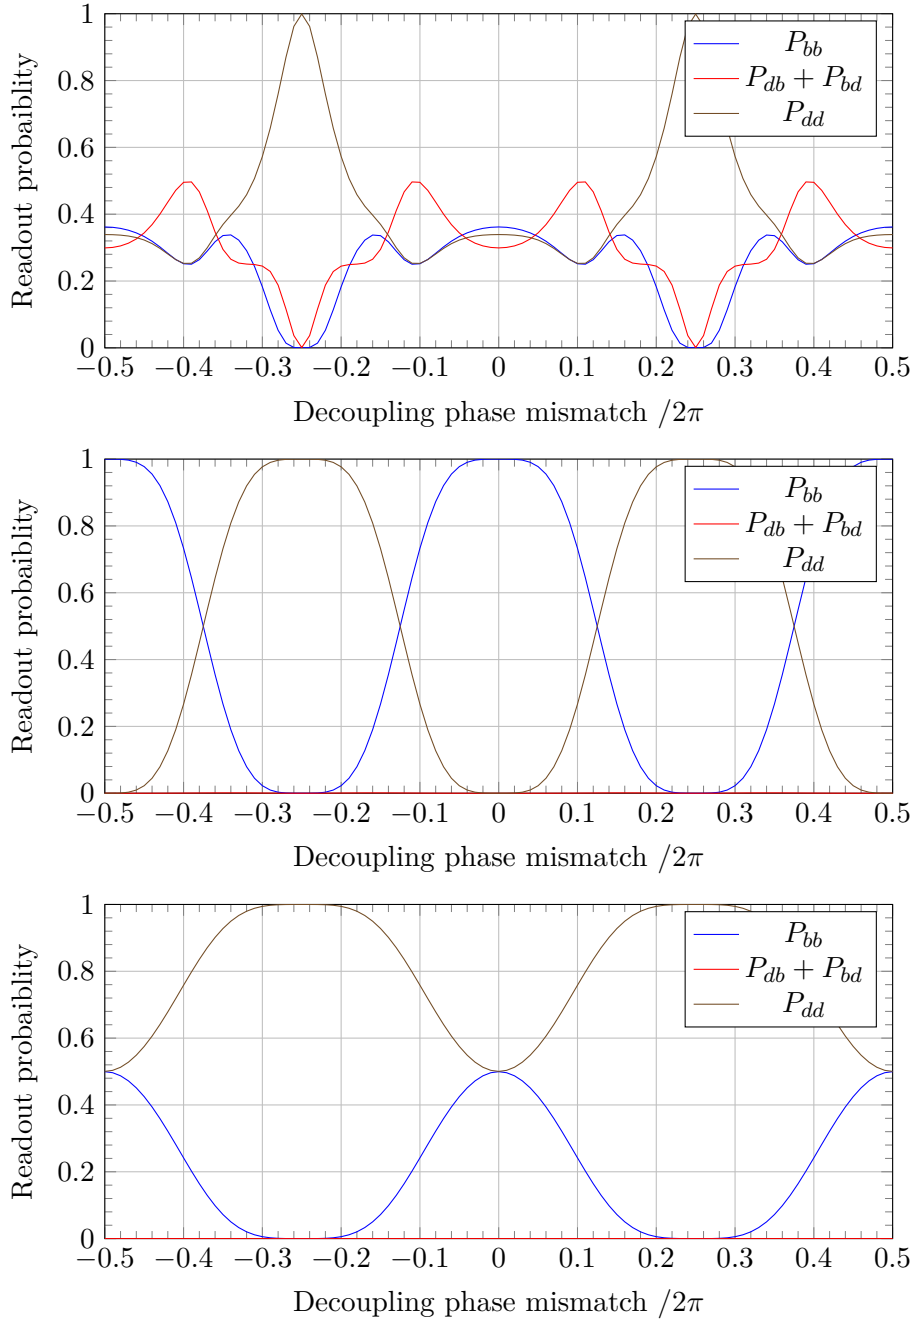


Figure 6.7: **Decoupling phase scans** - Simulated dynamical decoupling phase scans for a range of sideband detunings for a $330 \mu\text{s}$, two-loop W7MDMS gate with $\Omega_{Q,n}^{\text{DD}} = 2\pi \times 180 \text{ kHz}$. From top to bottom, the carrier phase scans are performed with ‘zero sideband detuning’, ‘half the gate detuning’, ‘at the gate detuning’. Notably, all scans are symmetric around zero decoupling phase mismatch. Further, gate dynamics are suppressed at 0.5π decoupling phase mismatch.




Figure 6.8: **Fine gate duration adjustment** - Simulated sideband detuning scan for a two-loop W7MDMS gate with $\Omega_{Q,n}^{\text{DD}} = 2\pi \times 180 \text{ kHz}$. The y-axis indicates the qubit measurement outcome probabilities where blue, red and purple lines indicate P_{bb} , $P_{bd} + P_{db}$, and P_{dd} respectively. The solid lines indicate the measurement probabilities for a gate with correctly calibrated $330 \mu\text{s}$ gate duration. The dashed and dotted lines indicate measurement probabilities when the gate duration is 10% too short (dotted) or too long (dashed). Notably, the position of the $P_{dd} = P_{bb}$ crossing point relative to the $P_{bd} + P_{db}$ minimum is indicative of the gate duration error.

Chapter 7

Two-Qubit Gate

Characterisation

We now present the two-qubit logic gates demonstrated using the procedures outlined in chapter 6. Characterisation measurements of the implemented two-qubit gates are presented in section 7.1. This is followed by estimates of the contributing error sources presented in section 7.2.

7.1 Gate measurements

The fidelity of the implemented two qubit gate was estimated using parity measurements of the resulting Bell state, see sections 3.2.2 and 7.1.1. Beyond the two-qubit gate error, the parity measurement is also subject to state preparation and measurement errors. We therefore measure and correct for state preparation and measurement errors as presented in section 7.1.2. For the implemented WMDMS gate, the Rabi frequency associated with the dynamical decoupling tone is also of interest. Associated characterisation measurements are discussed in section 7.1.3. Finally, microwave input power driving the

two-qubit gate is characterised in section 7.1.4. An overview of these measurements for the implemented two-qubit gates is presented in table 7.1.

7.1.1 Parity analysis

The fidelity of the implemented two-qubit gates was estimated using parity analysis as outlined in section 3.2.2. As both qubits are prepared in identical states before the two-qubit gate is applied, performing a measurement after a perfect two-qubit gate should give no chance of the two ions being in different states. Further, applying a $\pi/2$ analysis pulse should result in full contrast parity oscillations when scanning the phase of the analysis pulse. Using equation 3.29 presented in section 3.2.2, we require a minimum of three measurement sequences to perform parity analysis. The experimental sequence used for these measurements is presented in section 6.1. The three measurements are the ‘population’ measurement where no analysis pulse is performed, and two ‘parity’ measurements where analysis pulses with a $\pi/2$ phase difference are performed at the extreme points of the parity oscillation. To reduce systematics due to drifting experimental parameters, we interleave n measurement sequences of population and parity measurements as indicated in equation 7.1.

$$\left(\text{population, parity}_\phi, \text{parity}_{\phi+\frac{\pi}{2}}\right)^n \quad (7.1)$$

A fidelity estimate using this sequence depends on choosing an analysis phase ϕ matching the extreme points of the parity oscillation. Two different procedures were employed to determine the maximum parity contrast.

In the first method we perform a scan of the analysis pulse phase, allowing the phase of maximum contrast to be determined. This scan was performed with points of different phases being acquired in randomised order and may be performed with reduced statistics compared to the gate characterisation

Gate duration	Fidelity	loop #	$\Omega_{Q,n}^{DD}$	Sideband edge time	Microwave power (per sideband)	Duty Cycle	$\bar{\epsilon}_{\text{SPAM}}$	Raw fidelity
331 μs	99.51(11)%	2	180 kHz	1 μs	1.8 W	8%	$2.8(2) \times 10^{-3}$	98.67(9)%
216.8 μs	99.15(20)%	2	0	2.8 μs	3.3 W	5%	$1.2(1) \times 10^{-3}$	98.78(20)%
153.8 μs	99.02(21)%	1	0	2.8 μs	3.3 W	5%	$1.2(1) \times 10^{-3}$	98.66(21)%

Table 7.1: **Table of demonstrated two-qubit gates** - The table show the duration of the demonstrated gate operation, as well as the SPAM-corrected fidelity and number of driven phase-space loops for the three demonstrated two qubit gates. For the demonstrated W7MDMS gate, the Rabi frequency of the dynamical decoupling drive is given as $\Omega_{Q,n}^{DD}$. The measured average SPAM error probability per ion is indicated by $\bar{\epsilon}_{\text{SPAM}}$. The pre SPAM correction “raw fidelity” is also given. The duration of the \sin^2 shaped sideband pulse edge for each gate is specified as the “Sideband edge time”. Finally, the power per microwave sideband tone delivered to the vacuum assembly and the microwave duty cycle are also shown. Notably, the W7MDMS and the two MS gates were performed in apparatus A and B respectively.

measurement, as we only seek to sufficiently determine the extreme phase, rather than resolve the parity contrast. Having estimated the analysis phase of maximum contrast, we collect sufficient measurement cycle repetitions at the maximum contrast phase to resolve the two-qubit gate error.

In the second method, a similar coarse scan analysis phase scan is performed. However, significantly reduced statistics compared to the first method are acquired. The rough analysis phase scan is followed up by a higher resolution scan of the analysis phase around the maximum and minimum of the parity signal. This higher resolution phase scan is performed in randomised order of analysis phase value ϕ and acquires measurements in the same interleaved fashion indicated in equation 7.1. Figure 7.1 shows the measured $P_{10} + P_{01}$ probability of such a analysis phase scan of the demonstrated 216.8 μs two-loop MS gate.

This difference in methods is due to a difference in experiment stability before and after the upset of the apparatus outlined in appendix B. The highest reported fidelity two-qubit gate performed in 331 μs was characterised before this upset and uses the first method. The other reported gates were characterised after the upset and employed the second method due to increased drifts of the apparatus. In the following the pre-and post-upset apparatus will be referred to as apparatus A and apparatus B respectively.

7.1.2 SPAM correction

Errors in the parity measures may not only be caused by errors in the implemented two-qubit gate, but also errors in the state preparation and measurement (SPAM) procedure. To accurately estimate the error originating from the two qubit gate, we characterise the state preparation and measurement errors and correct for the impact of imperfect SPAM on the measured fidelity.

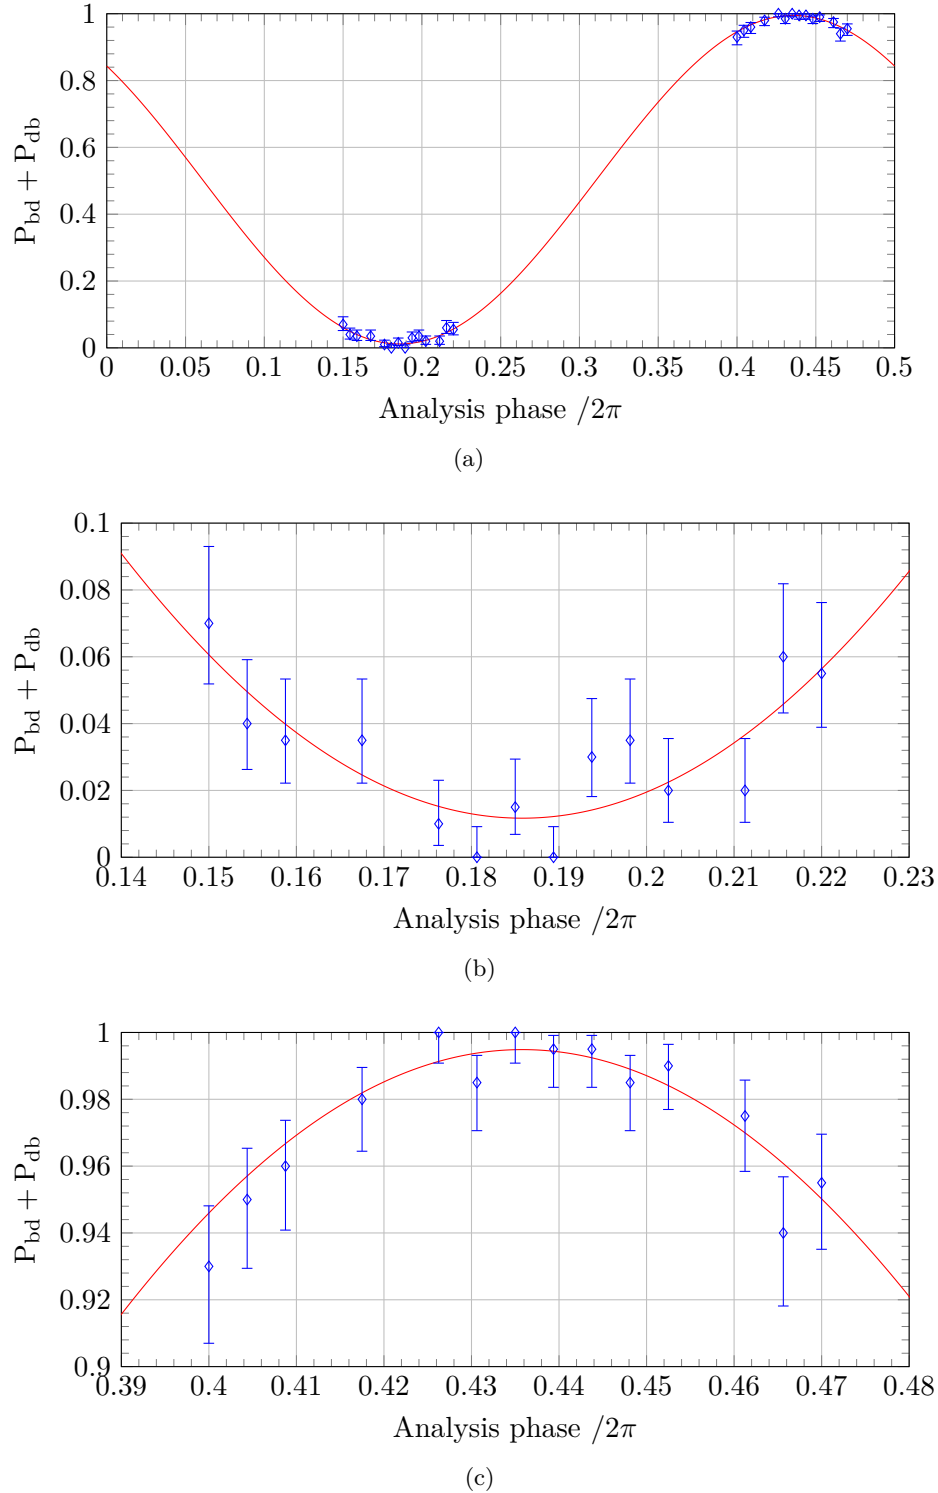


Figure 7.1: **Analysis phase scan** - Parity measurement of the demonstrated $216.8\ \mu\text{s}$ two-loop Mølmer-Sørensen gate. The x-axis indicates the phase of the analysis pulse. The measured $P_{bd} + P_{db}$ population after performing a the analysis pulse is indicated on the y-axis. Measured populations with statistical 68% confidence intervals are indicated in blue. The red line indicates the sinusoidal fit obtained via maximum likelihood estimation. a) This figure shows an overview of the measured analysis phase scan. b) This figure is a zoomed in section of a), showing the measured population minimum. c) This figure is a zoomed in section of a), showing the measured population maximum.

\	Per ion SPAM error	Expected average thresholding error
Apparatus A	$2.8(2) \times 10^{-3}$	3×10^{-4}
Apparatus B	$1.2(1) \times 10^{-3}$	3×10^{-4}

Table 7.2: **SPAM error vs thresholding error** - Measured clock qubit state preparation and measurement (SPAM) error compared to the expected mean error probability arising from the threshold state detection method, see section 2.6.3. The apparatus A errors were measured when performing the 331 μs W7MDMS gate. The apparatus B SPAM error was measured when performing the 154 μs & 217 μs MS gates. Notably, the measured per ion SPAM error is significantly larger than the average two-ion thresholding error probability. The improved SPAM fidelity of apparatus B compared to apparatus A is due to improved fidelity of the microwave transfer pulses through the introduction of a 20 μs delay after switching the Phaser DAC-NCO frequency, see 4.7.1.3.

In our case, the SPAM error is dominated by qubit independent error sources such as the reliability of the microwave transfer pulses, optical pumping, and optical shelving. Notably, we are not limited by thresholding errors, see table 7.2. As our SPAM errors are independent, the reduction of the measured fidelity caused by the single qubit SPAM error probability $\bar{\epsilon}_{\text{SPAM}}$ is estimated by equation 7.2 [114].

$$1 - \mathcal{F} \approx 3\bar{\epsilon}_{\text{SPAM}} \quad (7.2)$$

As both the SPAM error and gate error are small, we may estimate the first order SPAM corrected two qubit gate fidelity $\mathcal{F}_{\text{corrected}}$ from the raw parity fidelity \mathcal{F}_{raw} and the average single qubit SPAM error $\bar{\epsilon}_{\text{SPAM}}$ as shown in equation 7.3.

$$\mathcal{F}_{\text{corrected}} = \mathcal{F}_{\text{raw}} + 3\bar{\epsilon}_{\text{SPAM}} \quad (7.3)$$

7.1.3 Dynamical decoupling strength

The 331 μs gate in apparatus A was performed with a Walsh modulated decoupling tone. The strength of the employed dynamical decoupling was meas-

ured by observing Rabi oscillations of a single ion. This was done by replacing the Walsh modulated resonant pulse sequence during the gate with a single resonant pulse throughout the constant sideband power section of the two-qubit gate pulse. The amplitude and frequency settings of the resonant and sideband tones were kept identical to those used during the two-qubit gate. The gate pulse duration was then scanned around the pulse duration of the two-qubit gate, allowing Rabi oscillations driven by the resonant carrier tone to be observed. As the microwave sideband tones are frequency tuned to the two-ion rocking mode sideband detuning, they drive negligible interaction on a single ion. The Rabi oscillation characterisation measurement of the 331 μs WMDMS gate is shown in figure 7.2.

7.1.4 Gate input power

Once the measurements above were completed, the gate input power was characterised by measuring the output power of the microwave chain¹ when emitting single microwave tones of equal amplitude setting to those used during the gate pulse. Power of each microwave tone was measured using an Agilent V3500A power meter set 3.12 GHz and a 30 dB attenuator². The resultant input powers are shown in table 7.3. These powers are systematically different from those produced during the microwave gate as the Agilent V3500A does not facilitate the short integration times and triggering needed to measure the input powers when maintaining the same duty cycle as used during the two-qubit gate. As these powers are only meant for rough comparison to laser free gate mechanisms no attempt is made to compensate for this systematic uncertainty of the microwave input power.

¹This measurement was taken at the last accessible point before the microwave signal enters the vacuum system

²Mini-Circuits, BW-S30W20+




Figure 7.2: **WMDMS gate decoupling** - Measured WMDMS gate decoupling strength. The measurement was performed by applying a W0MDMS pulse, with amplitudes matching the ones used for the demonstrated 331 μs gate, to a single ion. The duration of the pulse was then scanned to show the average decoupling tone Rabi frequency across the gate pulse. The y-axis shows the probability of observing a dark ion. Measured probabilities and symmetric standard error estimates are shown in blue. A sinusoidal fit to the data gives a Rabi frequency of 180.70 ± 0.08 kHz and is shown in red.

\	Red sideband	Blue sideband	Carrier tone
Apparatus A	1.8 W	1.7 W	6 mW
Apparatus B	3.2 W	3.4 W	0

Table 7.3: **Microwave input powers** - Measured microwave power incident on the vacuum chamber from each of the microwave tones used in the demonstrated gates. The powers were measured using an Agilent V3500A power meter set to 3.12 GHz. As the sideband tone powers exceed the rated power of this power meter, the sideband powers were measured using a 30 dB attenuator.

7.2 Error characterisation

We now turn our attention to estimating the error contributions to our demonstrated gates. As the experiment conditions between the higher fidelity 331 μs gate performed in apparatus A and the faster gates performed in apparatus B are not the same, estimates for both sets of gates are provided where applicable.

7.2.1 Gate duration calibration error

There are two effective sources of gate duration miscalibration errors. Firstly, we empirically tune the gate duration to within $\pm 1 \mu\text{s}$ accuracy during the gate tune-up procedure, as presented in section 6.2.8. Secondly, drifts of the microwave interaction, such as are caused by microwave power fluctuations, correspond to drifts of the ideal gate duration. The latter is described by equation 7.4.

$$\frac{\Delta t_g}{t_g} = \frac{dt_g}{d\Omega'_Q} \frac{\Delta\Omega'_Q}{t_g} = -\frac{\Delta\Omega'_Q}{\Omega'_Q} \quad (7.4)$$

With the duty cycle management in place, fractional microwave field drifts are bounded at the 10^{-3} level (see section 7.2.6). Consequently, for our gate durations, gate duration errors are dominated by the $\pm 1 \mu\text{s}$ calibration accuracy. Corresponding gate errors are shown in table 7.4.

7.2.2 Mode frequency stability

The demonstrated gates did not adopt a sideband Walsh modulation strategy to mitigate motional detuning errors. Consequently, the error due to gate detuning errors is described by equation 3.30. Motional detuning errors may be caused by miscalibration of the optimal sideband detuning or by drift of

Gate type	W7MDMS	MS	MS
Error source	2-loop, 331 μ s	2-loop, 217 μ s	1-loop, 154 μ s
Apparatus	A	B	B
gate duration calibration error	$< 1.8 \times 10^{-5}$	$< 4.2 \times 10^{-5}$	$< 7.0 \times 10^{-5}$
Mode frequency drift	2.0×10^{-3}	5.1×10^{-3}	7.0×10^{-3}
Motional heating rate	$< 1.7 \times 10^{-4}$	$< 1.1 \times 10^{-4}$	$< 1.5 \times 10^{-4}$
Crosstalk to radial CoM mode	$< 1.6 \times 10^{-3}$	$< 1.6 \times 10^{-4}$	$< 1.6 \times 10^{-4}$
Kerr coupling to axial motion	6×10^{-5}	4×10^{-4}	4×10^{-4}
Kerr coupling out-of plane motion	2×10^{-5}	1.7×10^{-3}	1.7×10^{-3}
Motional dephasing	?	?	?
Sideband imbalance.	4.6×10^{-8}	2.8×10^{-8}	1.4×10^{-9}
Qubit frequency error	$< 6 \times 10^{-5}$	8×10^{-4}	3.5×10^{-4}
Analysis pulse error	$< 1 \times 10^{-5}$	$< 1 \times 10^{-5}$	$< 1 \times 10^{-5}$
Decoupling drive qubit rotation	?	n/a	n/a
Off-resonant carrier excitation	$< 1.6 \times 10^{-6}$	$< 1.2 \times 10^{-6}$	$< 9.0 \times 10^{-7}$
Broadband microwave noise	$< 1.6 \times 10^{-5}$	$< 1.0 \times 10^{-5}$	$< 7 \times 10^{-6}$
Qubit spectator transition excitation	$< 2 \times 10^{-11}$	$< 3 \times 10^{-13}$	$< 3 \times 10^{-13}$
Magnetic field noise	4.8×10^{-5}	5.6×10^{-5}	4.0×10^{-5}
Raman photon scatter	3.3×10^{-4}	2.2×10^{-4}	1.5×10^{-4}
Estimated total error	$\sim 3.5 \times 10^{-3}$	$\sim 8.5 \times 10^{-3}$	$\sim 1.0 \times 10^{-2}$
Measured error	$4.9(11) \times 10^{-3}$	$8.5(20) \times 10^{-3}$	$9.8(21) \times 10^{-3}$

Table 7.4: **Gate error budget** - Breakdown of estimated gate error contribution by a variety of gate error sources. Gate error sources which have been bounded to be below a threshold, rather than estimating the error in the gate, are indicated by $<$ signs. Error sources which were not estimated are marked by ‘?’ and sources not applicable to a given gate are marked ‘n/a’. Notably, mode frequency drifts are believed to contribute significant errors to all demonstrated gates. For each gate a total estimated error is also provided. This estimate sums the estimated errors linearly and uses typical values for bounded error sources. It is believed that the error discrepancy between the estimated and measured 331 μ s gate error is accounted for by global qubit rotations driven by the dynamical decoupling tone.

the mode frequency in the time since the mode detuning was last calibrated. In our case, we are dominated by slow drifts of the normal mode frequency. The motional frequency stability was significantly degraded in apparatus B, used for the faster gates, when compared to apparatus A, which was used for the high fidelity 331 μs gate. For the high fidelity gate, the mode frequency was measured to have drifted ~ 130 Hz over the course of the gate characterisation measurement. For gates performed in apparatus B, the stability was significantly degraded. The drift over the gate duration was estimated to be ~ 200 Hz min^{-1} amounting to ~ 390 Hz over the gate characterisation measurement. Measurement of this drift was hampered by the reduced ion lifetime in apparatus B and systematic mode frequency changes associated with our reloading process.

Empirically, the mode frequency drift was observed to be monotonic on the experiment timescales. The gate detuning of gates during gate tune-up was therefore manually chosen to anticipate some of the mode frequency drift. Estimates of the resulting gate errors, including the anticipated drift correction, are presented in table 7.4. Mode frequency fluctuations contribute significantly to the observed error in all demonstrated gates. Future efforts may mitigate this error through use of sideband Walsh modulation, see 3.3.2. Alternatively, the mode frequency drift could be tracked during gate characterisation.

7.2.3 Heating rate

As the MS gate depends on entangling the qubit state with the ion motion, processes decohering the ion motion result in gate errors. One such process is thermalisation of the ion motion with environmental electric field noise. As the ion is laser cooled to ~ 10 μK temperatures, this process is represented by heating of the ion motion characterised by a heating rate. Within our

experiment this heating rate was measured using the standard technique of Raman sideband thermometry. Before performing the 331 μs gate a rocking mode heating rate of $1.3 \text{ quanta s}^{-1}$ was measured. However, the heating rate drifts somewhat, while typically remaining $< 2 \text{ quanta s}^{-1}$. We therefore estimate the error based on the typical heating rate range. The two-qubit gate error associated with this heating rate is presented in table 7.4.

7.2.4 Coupling to spectator modes

The ideal MS gate only couples to a single mode of the collective ion motion. However, in practice, the spectating motional modes impact the gate dynamics both through spectral crosstalk and direct mode-mode interactions. In the following we discuss both of these effects and their likely impact on the gate fidelity.

7.2.4.1 Frequency addressing of motional mode

Our treatment of the spin dependent force in section 3.1.1 neglected interactions with other motional modes under the justification that the gate tones were far detuned from the motional sidebands associated with other motional modes. Or more specifically that the motional interaction strength Ω'_Q is much weaker than the detuning from the spectator mode sideband transition Δ_s . As the speed of driven two qubit gates increases, the validity of the $\frac{\Omega'_Q}{\Delta_s} \ll 1$ approximation decreases. As we target the in-plane rocking mode with our gate interaction, the spectator mode with the smallest detuning from the sideband tones is the in-plane centre-of-mass (CoM) mode. Considering this detuning we note that the gates we demonstrate are still in the far detuned regime with $\frac{\Omega'_Q}{\Delta_s} \lesssim 0.02$. However, with this interaction strength-detuning ratio some mode crosstalk effects become apparent.

Considering the effect of the spin dependent force on the CoM mode, phase space trajectories similar to the rocking mode are driven. However, due to the significantly larger detuning, the spin dependent phase space loops driven by the interaction are smaller and close on the timescale of the spectator detuning Δ_s . In section 3.2.1 we noted that rocking mode phase is only accumulated for the $|+-\rangle$ and $| -+\rangle$ states. Due the in-phase motion of the CoM mode, the driven CoM dynamics accumulate geometric phase on the $|++\rangle$ and $|--\rangle$ states. The rate of accumulating geometric phase ϕ via the CoM mode in the limit of the spectating CoM mode detuning Δ_{CoM} being much larger than the gate detuning Δ_g is expressed by equation 7.5.

$$\frac{\partial}{\partial t} \left(\frac{\phi_{\text{CoM}}}{\phi_g} \right) \approx \frac{\Delta_g}{\Delta_{\text{CoM}}} \quad (7.5)$$

Due to the geometric phase from the CoM and gate modes accumulating on different qubit states, the necessary gate duration is increased. From equation 7.5 we attain a 5% gate speed reduction for the WMDMS gate performed in apparatus A. Though this gate speed reduction is undesirable, it is calibrated for during our gate tune-up procedure, presented in section 6.2. Beyond the gate speed effect, the accumulated geometric phase now also depends on the CoM motion. Therefore, loop closure errors of the CoM mode result in entanglement of the qubit states with the CoM motion. Such residual entanglement results in two-qubit gate errors. We make no effort to ensure loop closure on the CoM mode and present simulated worst case errors in, table 7.4. As the worst case residual entanglement error would have increased to 3.8×10^{-3} when performing the $154 \mu\text{s}$ and $217 \mu\text{s}$ gates, the axial confinement strength was increased from $\omega_z = 2\pi \times 1.1 \text{ MHz}$ to $\omega_z = 2\pi \times 2.2 \text{ MHz}$ for these measurements. The increased axial confinement strength resulted in an increase of the CoM spectator mode detuning from $\sim 2\pi \times 120 \text{ kHz}$ to $\sim 2\pi \times 560 \text{ kHz}$, reducing the residual entanglement effects to the 10^{-4} level.

Mode	Apparatus A	Apparatus B
Axial centre of mass	1.1 MHz	2.2 MHz
Axial breathing	1.9 MHz	3.8 MHz
Out-of-plane centre of mass	6.1 MHz	6.5 MHz
Out-of-plane rocking	6.0 MHz	6.1 MHz
In-plane centre of mass	5.7 MHz	4.9 MHz
Gate mode (In-plane rocking mode)	5.6 MHz	4.3 MHz

Table 7.5: **Motional mode frequencies** - Complete list of motional modes of the trapped ions and their associated frequencies in both apparatus A and apparatus B.

An overview of all modes of the ion motion and their frequencies is presented in table 7.5.

7.2.4.2 Kerr coupling

As discussed in section 3.2.3.3 direct mode-mode interactions may also cause gate errors through the Kerr non-linearity where the motional state of spectator modes imparts a shift on the frequency of the gate mode frequency. Due to symmetry, the Kerr interaction only occurs with the axial stretch and out-of-plane rocking modes. The respective sensitivities of the gate mode frequency to these spectator modes are given by equations 3.34 & 3.35 in section 3.2.3.3. Both the axial and out-of-plane Kerr sensitivities scale unfavourably with increased axial confinement as it reduces the ion-ion distance. Respectively, these errors scale approximately as $\omega_z^{2+2/3}$ and $\omega_z^{6+2/3}$. As the axial confinement frequency was increased twofold from the high fidelity gate in apparatus A to the faster gates in apparatus B, these sensitivities of the faster gates are increased approximately 6-fold and 100-fold respectively.

As the axial-stretch and out-of-plane rocking mode occupation numbers were not measured, we make an estimate of their impact on the gate fidelity. The gate demonstrated in apparatus A, used an axial confinement $\omega_z = 2\pi \times$

1.1 MHz with associated radial mode frequencies $\omega_{\text{oop-rocking}} = 2\pi \times 6.01$ MHz, and $\omega_g = 2\pi \times 5.6$ MHz. Of these modes, only the gate mode is ground state cooled. The occupation of the other modes is assumed to be in a thermal state with average occupation numbers determined by the dark resonance cooling limit. This limit is mode dependent and influenced by the mode frequency, and the dark resonance cooling projection onto the mode. From these we estimate the respective average occupation numbers of the axial stretch and out-of-plane rocking modes to be 4 and 19 quanta, respectively. In apparatus B, an axial confinement with $\omega_z = 2\pi \times 2.2$ MHz was employed, see table 7.5 for frequencies of other modes. In this configuration we estimate the spectating stretch and rocking mode average occupation numbers as 2 and 19 quanta, respectively. The corresponding gate errors are presented in table 7.4.

7.2.5 Motional dephasing

Beyond the motional dephasing mechanisms discussed so far, it is possible for other mechanisms, such as mode frequency fluctuations on the gate timescale, to cause decoherence of the ion-motion. No direct measurement of the motional coherence time has been performed. In the absence of technical limitations, the motional coherence time would be limited by the motional heating rate. We therefore make the assumption that motional dephasing does not limit the gate fidelity.

7.2.6 Microwave field

Beyond the motional error mechanisms discussed so far, fluctuations in the microwave field used to drive the gate dynamics also lead to errors in the two-qubit gate. There are multiple types of microwave field fluctuations that might cause errors in the gate dynamics. Here we will consider the impact

of slow drifts in the microwave field amplitude and microwave field gradient as well as variations of these on the timescale of the gate operation. Besides variations in the microwave power delivered to the trap, such fluctuations can also be caused by variations of the properties of the trap integrated microwave resonator, see section 4.9 & appendix B.

7.2.6.1 Microwave field characterisation error

As part of the gate tune-up procedure the σ and π microwave field components are characterised. This measurement is used to set the power balancing of the microwave sideband tones via the resultant AC-Zeeman shift. A mischaracterisation of the σ and π microwave field components therefore results in an amplitude imbalance between the microwave sideband tones experienced by the ion. Assessing the fit quality, we can bound the fractional error in the σ and π microwave field to be $< 1\%$. To relate this to the resulting sideband imbalance, we must consider the origin of the Zeeman shift experienced in the presence of symmetrically detuned sideband tones from the qubit frequency.

Due to the equal and opposite detuning of the microwave sideband tones from the qubit transition, the shifts associated with the microwave π field cancel near perfectly when the sidebands are balanced. Consequently, the two-tone Zeeman shift is dominated by the microwave σ field. However, the ratio between the σ and π polarised shifts resulting from a single sideband tone is $\frac{\Delta f_{\text{Zeeman } \sigma}}{\Delta f_{\text{Zeeman } \pi}} \approx 10^{-3}$. Consequently, the change in Zeeman shift arising from a sideband imbalance is dominated by the microwave π field which is not cancelled by the opposite sideband π field. Therefore, a $\sim 1\%$ field mischaracterisation of the balanced AC-Zeeman shift only results in a $\sim 10^{-5}$ fractional sideband imbalance. The resulting sideband imbalance is insignificant as it is dominated by slow drifts of the sideband balancing, see section 7.2.6.2.

7.2.6.2 Slow microwave field drifts

As part of the gate tune-up procedure, the AC-Zeeman shift generated by the sideband tones is characterised, see sections 6.2.5 & 6.2.6. Comparing these measurements over time, we estimate slow drifts in the AC-Zeeman shift. For the 331 μs gate in apparatus A, we estimate the Zeeman shift miscalibration to be < 600 Hz. Figure 7.3 shows the dependence of the expected gate error for the demonstrated two-loop, 7th order WMDMS gate. Compared to a standard MS gate, the WMDMS scheme results in a ~ 4 order of magnitude error reduction. Consequently, a 600 Hz AC-Zeeman shift error results in 6×10^{-5} gate error, which is much smaller than motional error sources. For the gates in apparatus B the Zeeman shift drift was reduced to ~ 50 Hz. As the AC-Zeeman shift stability was improved, standard MS gates were performed. Simulated gate errors resulting from these qubit frequency detuning errors are summarised in table 7.4.

Beyond the qubit detuning errors, such AC-Zeeman shift fluctuations are associated with fluctuations of the microwave field amplitude experienced by the ion. From the bound on the AC-Zeeman shift variation, we can derive bounds on the slow common mode and differential drifts of the microwave sideband amplitudes. In apparatus A these common mode and differential amplitude drift bounds are $< 1 \times 10^{-2}$ and $< 3 \times 10^{-4}$ respectively. However, we can establish a tighter bound on common mode amplitude drifts as the calibrated amplitude of the analysis pulse stayed within a $< 3 \times 10^{-3}$ fractional range for over an hour. This variation is consistent with the analysis pulse calibration precision.

The common sideband amplitude drift corresponds to a change of the ideal gate duration and detuning. From our discussion of gate duration calibration error in section 7.2.1 fractional gate duration fluctuations of $\sim 3 \times 10^{-3}$ are

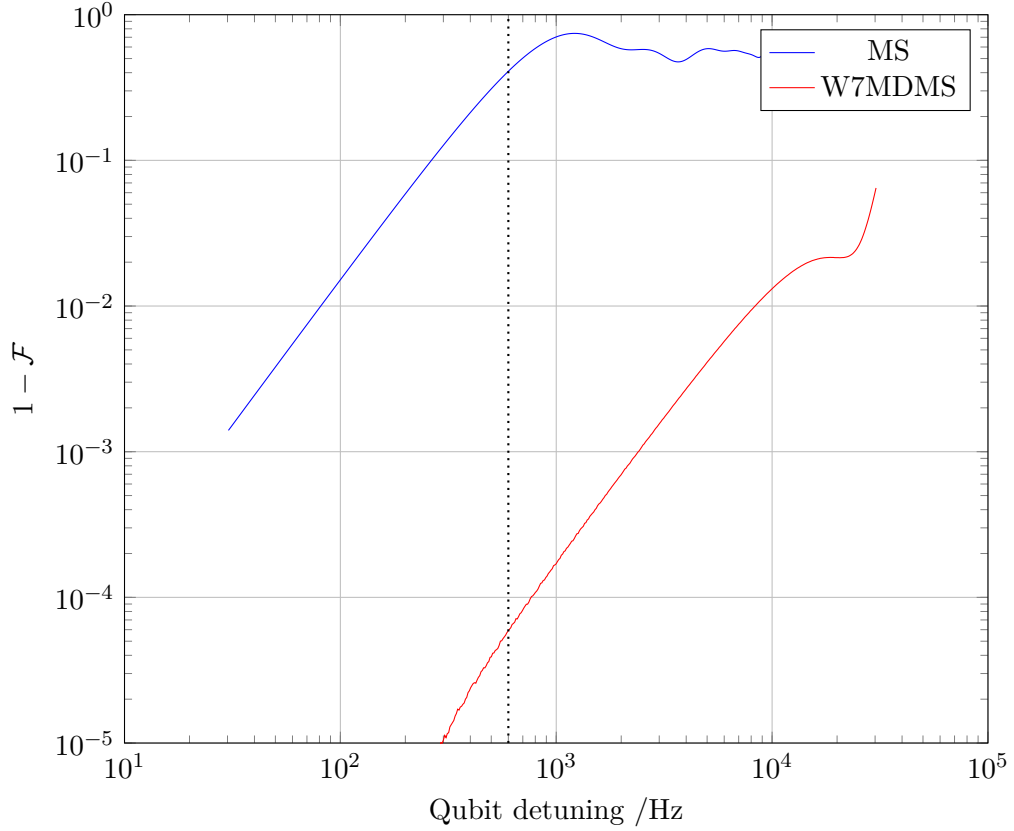


Figure 7.3: **2 loop W7MDMS detuning error** - Simulated gate error dependence on qubit frequency detuning for a $331 \mu\text{s}$ two-loop W7MDMS gate with $\Omega_{Q,n}^{\text{DD}} = 2\pi \times 180 \text{ kHz}$ and standard $331 \mu\text{s}$ two-loop MS gate. The vertical dotted line indicates the estimated qubit detuning upper bound, which corresponds to $\sim 6 \times 10^{-5}$ two qubit gate error.

on the same scale as the inaccuracy of our gate duration calibration. Further, the resulting gate detuning fluctuations are dominated by the drift of the gate mode frequency, see section 7.2.2. Finally, the drift of the ‘analysis’ pulse will degrade the measured parity of the produced Bell state. Table 7.4 shows the apparent gate error such an analysis pulse degradation would result in.

Considering the differential sideband amplitude drift, we note that it results in an imbalance of the sideband amplitudes. The associated sideband imbalance error is shown in table 7.4.

7.2.6.3 Fast microwave field drifts

Besides the long timescale microwave field errors discussed so far, it is possible for the microwave field to fluctuate within the application of a gate pulse. Such a fluctuation would result in a time-varying qubit frequency error, two-qubit gate speed, and decoupling tone strength. The shape and reproducibility of such fast fluctuations has not been characterised. Generally, the self cancellation achieved by Walsh modulating the decoupling tone phase relies on Rabi frequency variations throughout the gate being described by a sufficiently low order polynomial. However, it is believed that the fidelity of the demonstrated WMDMS gate may have been limited by imperfect self cancellation due to fast microwave field fluctuations.

Measurements performed on a single ion indicate the presence of intra gate Rabi fluctuations by showing imperfect self cancellation. These measurements used the microwave tone powers and pulse durations for a single loop $162\ \mu\text{s}$ WMDMS gate with $\Omega_{Q,n}^{\text{DD}} \approx 1\ \text{MHz}$. The self cancellation error of the Walsh modulated carrier was observed to scale quadratically with the decoupling strength $\Omega_{Q,n}^{\text{DD}}$. Further, the error was measured to be dependent on the Walsh modulation order and was at a maximum for the Walsh-7 modulation, coinciding with the modulation order used for with the demonstrated $331\ \mu\text{s}$ WMDMS gate. Decreasing the modulation to Walsh-3 or increasing the Walsh modulation to Walsh-31 both resulted in significant reductions in qubit rotation error, see figure 7.4. This behaviour is consistent with a significant cubic and/or quartic component to the microwave amplitude variation throughout the applied gate pulse. As these characterisation measurements were performed at a different sideband power and pulse duration than were used in the WMDMS gate they do not allow straightforward estimation of the corresponding error in the demonstrated $331\ \mu\text{s}$ WMDMS gate. Unfortu-

nately, further characterisation was precluded by the apparatus upset detailed in appendix B. However, it seems reasonable that the unexplained $\sim 2 \times 10^{-3}$ WMDMS gate error may arise through this mechanism.

In future gates, this error could be mitigated by choosing a more suitable Walsh modulation order. As the MS gates which were demonstrated in apparatus B do not use a decoupling drive, intra-gate carrier Rabi frequency fluctuations do not contribute a significant error in these gates.

7.2.6.4 Broadband microwave noise

In our preceding discussion of the microwave signal chain we noted that the ion experienced broadband microwave noise, see section 4.9. We have bounded the transition timescale associated with this broadband noise to be > 20 s. For few 100 μ s MS gates this bounds the broadband noise induced MS gate error on the $\sim 10^{-5}$ level. Bounds for each gate are presented in table 7.4.

7.2.7 Off-resonant qubit excitation by sidebands

The demonstrated gates adopted pulse shaping to mitigate the effect of qubit rotations driven by the microwave sideband tones, see section 3.3.1. From the chosen softening of the pulse edges, we are able to bound the worst case error from off-resonant carrier excitation. The corresponding bounds for the different implemented gates are presented in table 7.4.

7.2.8 Qubit spectator transitions

Besides off-resonantly driving qubit Rabi oscillations by the π component of the microwave sideband field, the σ component of the microwave sideband field couples to ~ 100 MHz detuned spectator transitions connected to the $|F = 3, m_f = +1\rangle$, and $|F = 4, m_f = +1\rangle$ qubit states. This interaction causes




Figure 7.4: **Decoupling drive qubit rotation error** - Measured SPAM corrected single qubit error probability when applying a single loop, $162\ \mu\text{s}$ WMDMS gate pulse with decoupling strength $\Omega_{Q,n}^{\text{DD}} \approx 1\ \text{MHz}$. The x-axis indicates the decoupling tone Walsh modulation order n_{Walsh} . The y-axis shows the measured SPAM corrected single qubit transition probability after application of the WMDMS pulse. The error bars represent the standard confidence interval. Notably, the self cancellation error has a local maximum at the W7MDMS gate, which was used in the demonstrated $331\ \mu\text{s}$ two qubit gate. The measured error probability when using Walsh-3 or Walsh-31 modulation is consistent with no error and significantly outperforms the error observed when using Walsh-7 or Walsh-15 modulation. The increased error when using Walsh-7 and Walsh-15 modulation is consistent with a cubic and quartic Rabi frequency fluctuation during the application of the gate pulse.

the initial clock qubit states to become dressed states that have been mixed with other hyperfine states of the $4S_{1/2}$ ground level. The dressing with the spectator states results in two error mechanisms.

Firstly, diabatic enabling and disabling of the of the microwave sideband tones will cause the initial qubit state to be projected onto the dressed states. Due to the precession of the dressed states, the state produced by projecting back onto the initial hyperfine states at the end of the sideband interaction, generally yields a superposition of qubit and spectator hyperfine states. This causes an error where the ion state has leaked from the qubit subspace. Fortunately, careful accounting for the precession of the dressed states can be avoided by adiabatically ramping the sideband pulse amplitudes. As we already ramp into and out of our sideband pulses on the $1\ \mu\text{s}$ timescale to suppress the off-resonant excitation of the qubit transition, the transition into and out of the dressed states is well into the adiabatic regime yielding a worst case error of order 10^{-11} , see table 7.4.

Secondly, though our initial qubit states were on a clock transition (see section 2.3.3), the dressed states produced in the presence of the microwave field are not magnetic field insensitive. Consequently, during the operation of the microwave gate, broadband noise in the environmental magnetic field couples into the dressed state energy, leading to qubit dephasing errors. Using the magnetic field noise characterisation presented in section 4.5.1, we estimate the associated gate error caused by this induced magnetic field sensitivity to be $\sim 5 \times 10^{-5}$. Estimates for each implemented gate are presented in table 7.4.

7.2.9 Raman laser photon scatter

Though we demonstrate “laser free” two-qubit gates, we employ two 397 nm Raman laser beams with ~ 800 saturation intensities for Raman sideband cooling. As outlined in section 4.6, we employ a single pass AOM for optical switching of these beams. This configuration has limited extinction, resulting in Raman laser leakage during gate operations. Out of the two qubit states, the Raman leakage light predominantly couples to the $4S_{1/2}, |F = 3, m_f = +1\rangle$ state, which has a characteristic photon scattering time of ~ 1 s per ion. The corresponding gate errors for each gate are tabulated in 7.4.

7.2.10 Error summary

Combining the gate error contributions from the sources considered above, we find that the estimated total gate errors broadly align with the measured gate error probabilities. In the case of the 331 μ s W7MDMS gate, we measure a gate error probability of $4.9(11) \times 10^{-3}$ where we expect $\sim 3.5 \times 10^{-3}$ error from the combined effect of the error estimates above. In the case of the 154 μ s and 217 μ s MS gates performed in apparatus B, the respective measured gate errors of $9.8(21) \times 10^{-3}$ and $8.5(20) \times 10^{-3}$ match well with the 1.0×10^{-2} and 8.5×10^{-3} estimated gate errors.

For all demonstrated gates, we estimate the gate error to be dominated by mode frequency drifts. The next most significant error in the W7MDMS gate is believed to be crosstalk to the in-plane radial centre-of-mass mode or a potential qubit rotation error caused by cancellation failure of the decoupling tone Walsh modulation. In the case of the two MS gates, the next most significant error is estimated to be Kerr interaction with the out-of-plane radial mode.

Chapter 8

Conclusions

In this work we have presented two advances towards scaling up trapped-ion quantum computers using chip-based ion traps. Firstly, we have developed a toolbox for trap agnostic ion transport calculations and successfully demonstrated ion transport as well as splitting and merging operations within two different ion traps using the waveforms calculated by this ion transport toolbox. Secondly, we have developed a novel error mitigation strategy for near field microwave gradient driven two-ion quantum logic gates, which suppresses errors from both fluctuations in the qubit frequency and imperfection in the decoupling drive itself. Using this novel error mitigation method we have demonstrated 331 μs two-ion quantum logic gates, with $4.9(11) \times 10^{-3}$ logic error probability which is below the 1% error threshold required for quantum error correction. This represents a $9\times$ gate duration reduction when compared to previously demonstrated near field gradient driven microwave gates exceeding the 1% error probability threshold. Additionally, two faster Mølmer-Sørensen style gates were demonstrated. Respectively, the two Mølmer-Sørensen gates had gate operation times of 216.8 μs & 153.8 μs and measured gate error probabilities of $8.5(20) \times 10^{-3}$ & $9.8(21) \times 10^{-3}$. When compared to previous




Figure 8.1: **Laser free two-ion gates** - Selection of laser free two ion gate demonstrations using the near field microwave gradient, near-motional-frequency magnetic field gradient and MAGIC gate mechanisms. The gates performed in this work are shown in black. Previously demonstrated gates using near field microwave gradients are shown in blue. Results using the MAGIC (far field) and near-motional-frequency magnetic field gradient methods are shown in green and red respectively. Prior to this work, the fastest near field microwave gradient driven gate demonstrations which are below the 1% error probability required for error correction have gate durations $t_g = 2938 \mu\text{s}$. The fastest previously-demonstrated error correctable laser free gate was performed using the near-motional-frequency magnetic field gradient method with a gate duration $t_g = 740 \mu\text{s}$. The two ion gates demonstrated in this work achieve over $2\times$ shorter gate durations while retaining error correctable fidelity. Alternatively, when compared to previous demonstrations of laser-free gates with comparable operation times, the gate error probabilities measured in this work are approximately $5\times$ lower.

demonstrations of laser-free gates with comparable operation times, the gate errors probabilities measured in this work are over $5\times$ lower. These gate operation times are over $2\times$ faster than all previous demonstrations of error correctable laser free two ion gates and make laser free gate operation times comparable to the $\sim 100\mu\text{s}$ gate durations of common laser driven two-ion gates [115]. An overview of the improved laser-free two-ion gate speed and fidelity is shown in figure 8.1.

The dominant error sources degrading the fidelity of the demonstrated laser free gates are believed to be drifts in the radial motional mode frequency, Kerr coupling to spectator modes and potential global qubit rotation errors arising from the dynamical decoupling tone. The motional mode frequency stability of the apparatus was significantly degraded when performing the $216.8\mu\text{s}$ & $153.8\mu\text{s}$ two-ion gates, future work may improve the demonstrated gate fidelities by restoring the apparatus to previously demonstrated mode frequency stabilities. Additionally, errors arising from both motional mode frequency drifts and Kerr coupling could be suppressed by employing sideband Walsh modulation. Finally, future work could avoid errors arising from qubit rotations driven by the dynamical decoupling tone by choosing the optimum Walsh modulation order for the decoupling tone.

This work used a surface trap with an integrated microwave resonator. This allowed resonant enhancement to significantly increase the input power efficiency of the apparatus, enabling faster gate operation. However, power dissipation through the application of microwaves may cause changes in the properties of this resonator. Future designs may reduce this effect through the use of superconducting traps as these can avoid the dissipation of microwave energy on the trap chip. Investigation of traps fabricated from high temperature superconductors, operated well below their critical temperature may also

be of interest as such materials can exhibit much higher critical currents and magnetic fields than conventional superconductor materials [116].

Bibliography

- [1] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, jun 1982.
- [2] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2(July):79, 2018.
- [3] Lov 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*, pages 212–219, may 1996.
- [4] Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, 2003.
- [5] David P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.
- [6] Pieter Kok. Photonic quantum information processing. *Contemporary Physics*, 57(4):526–544, 2016.
- [7] G. Wendin. Quantum information processing with superconducting circuits: A review. *Reports on Progress in Physics*, 80(10), 2017.

-
- [8] Gang Qin Liu and Xin Yu Pan. Quantum information processing with nitrogen-vacancy centers in diamond. *Chinese Physics B*, 27(2), feb 2018.
- [9] Xiaoling Wu, Xinhui Liang, Yaoqi Tian, Fan Yang, Cheng Chen, Yong Chun Liu, Meng Khoon Tey, and Li You. A concise review of Rydberg atom based quantum computation and quantum simulation. *Chinese Physics B*, 30(2), feb 2021.
- [10] Colin D. Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M. Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), jun 2019.
- [11] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091–4094, 1995.
- [12] D.J. Wineland, C. Monroe, W.M. Itano, D. Leibfried, B.E. King, and D.M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259, may 1998.
- [13] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476(7359):181–184, 2011.
- [14] H. C. Nägerl, Ch. Roos, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Investigating a qubit candidate: Spectroscopy on the S1/2 to D5/2 transition of a trapped calcium ion in a linear Paul trap. *Physical Review A*, 61(2):023405, jan 2000.
- [15] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. Demarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leib-

-
- fried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-lived qubit memory using atomic ions. *Physical Review Letters*, 95(6), aug 2005.
- [16] M. A. Sepiol, A. C. Hughes, J. E. Tarlton, D. P. Nadlinger, T. G. Ballance, C. J. Ballance, T. P. Harty, A. M. Steane, J. F. Goodwin, and D. M. Lucas. Probing qubit memory errors at the part-per-million level. *Physical Review Letters*, 123(11), sep 2019.
- [17] Ye Wang, Mark Um, Junhua Zhang, Shuoming An, Ming Lyu, Jing Ning Zhang, L. M. Duan, Dahyun Yum, and Kihwan Kim. Single-qubit quantum memory exceeding ten-minute coherence time. *Nature Photonics*, 11(10):646–650, oct 2017.
- [18] T. P. Harty, D. T.C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical Review Letters*, 113(22), nov 2014.
- [19] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-Fidelity Universal Gate Set for Be⁹⁺ Ion Qubits. *Physical Review Letters*, 117(6), aug 2016.
- [20] C J Ballance, T P Harty, N M Linke, M A Sepiol, and D M Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. *Physical Review Letters*, 117(6):1–6, 2016.
- [21] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter. High-fidelity laser-free

- universal control of two trapped ion qubits. *Nature*, 597(7875):209–213, sep 2021.
- [22] Craig R. Clark, Holly N. Tinkey, Brian C. Sawyer, Adam M. Meier, Karl A. Burkhardt, Christopher M. Seck, Christopher M. Shappert, Nicholas D. Guise, Curtis E. Volin, Spencer D. Fallek, Harley T. Hayden, Wade G. Rellergert, and Kenton R. Brown. High-Fidelity Bell-State Preparation with Optical Qubits. *Physical Review Letters*, 127(13), sep 2021.
- [23] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits. *Nature*, 555(7694):75–78, 2018.
- [24] Daniel Gottesman. Theory of fault-tolerant quantum computation. *Physical Review A*, 57(1):127–137, 1998.
- [25] Robert Raussendorf and Jim Harrington. Fault-tolerant quantum computation with high threshold in two dimensions. *Physical Review Letters*, 98(19), may 2007.
- [26] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(3), sep 2012.
- [27] Warren Nagourney, Jon Sandberg, and Hans Dehmelt. Shelved optical electron amplifier: Observation of quantum jumps. *Physical Review Letters*, 56(26):2797–2799, 1986.
- [28] Fangzhao Alex An, Anthony Ransford, Andrew Schaffer, Lucas R. Sletten, John Gaebler, James Hostetter, and Grahame Vittorini. High

- Fidelity State Preparation and Measurement of Ion Hyperfine Qubits with $I > 1/2$. *Physical Review Letters*, 129(13):130501, sep 2022.
- [29] Justin E. Christensen, David Hucul, Wesley C. Campbell, and Eric R. Hudson. High fidelity manipulation of a qubit built from a synthetic nucleus. 2:3–7, 2019.
- [30] Nicolai Friis, Oliver Marty, Christine Maier, Cornelius Hempel, Milan Holzäpfel, Petar Jurcevic, Martin B. Plenio, Marcus Huber, Christian Roos, Rainer Blatt, and Ben Lanyon. Observation of Entangled States of a Fully Controlled 20-Qubit System. *Physical Review X*, 8(2), apr 2018.
- [31] Kortny Rolston-Duce. Quantinuum announces quantum volume 4096 achievement, Apr 2022.
- [32] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, jun 2018.
- [33] Frank Arute and Et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, oct 2019.
- [34] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002.
- [35] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.
- [36] J. Stuart, R. Panock, C. D. Bruzewicz, J. A. Sedlacek, R. McConnell, I. L. Chuang, J. M. Sage, and J. Chiaverini. Chip-Integrated Voltage

-
- Sources for Control of Trapped Ions. *Physical Review Applied*, 11(2), feb 2019.
- [37] Zak David Romaszko, Seokjun Hong, Martin Siegele, Reuben Kahan Puddy, Foni Raphaël Lebrun-Gallagher, Sebastian Weidt, and Winfried Karl Hensinger. Engineering of microfabricated ion traps and integration of advanced on-chip features. *Nat Rev Phys*, 2(6):285–299, jun 2020.
- [38] Karan K. Mehta, Colin D. Bruzewicz, Robert McConnell, Rajeev J. Ram, Jeremy M. Sage, and John Chiaverini. Integrated optical addressing of an ion qubit. *Nature Nanotechnology*, 11(12):1066–1070, dec 2016.
- [39] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A - Atomic, Molecular, and Optical Physics*, 75(4), apr 2007.
- [40] Samuel Vizvary, Matthew Boguslawski, Zachary J Wall, Isam D Moore, Eric R Hudson, and Wesley C Campbell. Fundamental Limits on Gate-Laser Scattering Errors in Barium-133. In *53rd Annual Meeting of the APS Division of Atomic, Molecular and Optical Physics Volume 67, Number 7*, Orlando, Florida, 2022. Bulletin of the American Physical Society.
- [41] T P Harty, M A Sepiol, D. T.C. Allcock, C J Ballance, J E Tarlton, and D M Lucas. High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves. *Physical Review Letters*, 117(14):1–6, 2016.

-
- [42] G. Zarantonello, H. Hahn, J. Morgner, M. Schulte, A. Bautista-Salvador, R. F. Werner, K. Hammerer, and C. Ospelkaus. Robust and Resource-Efficient Microwave Near-Field Entangling Be+ 9 Gate. *Physical Review Letters*, 123(26):260503, 2019.
- [43] Florian Mintert and Christof Wunderlich. Ion-trap quantum logic using long-wavelength radiation. *Physical Review Letters*, 87(25):257904–1–257904–4, 2001.
- [44] M. Johanning, A. Braun, N. Timoney, V. Elman, W. Neuhauser, and Chr Wunderlich. Individual addressing of trapped ions and coupling of motional and spin states using rf radiation. *Physical Review Letters*, 102(7), feb 2009.
- [45] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger. Trapped-Ion Quantum Logic with Global Radiation Fields. *Physical Review Letters*, 117(22), nov 2016.
- [46] Shannon X. Wang, Jaroslaw Labaziewicz, Yufei Ge, Ruth Shewmon, and Isaac L. Chuang. Individual addressing of ions using magnetic field gradients in a surface-electrode ion trap. *Applied Physics Letters*, 94(9), 2009.
- [47] A. Khromova, Ch Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and Ch Wunderlich. Designer spin pseudomolecule implemented with trapped ions in a magnetic gradient. *Physical Review Letters*, 108(22), jun 2012.

-
- [48] N. Timoney, I. Baumgart, M. Johanning, A. F. Varón, M. B. Plenio, A. Retzker, and Ch Wunderlich. Quantum gates and memory using microwave-dressed states. *Nature*, 476(7359):185–188, aug 2011.
- [49] Lorenza Viola and Seth Lloyd. Dynamical suppression of decoherence in two-state quantum systems. *Phys. Rev. A*, 58(4):2733–2744, oct 1998.
- [50] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland. Trapped-ion quantum logic gates based on oscillating magnetic fields. *Physical Review Letters*, 101(9):1–4, 2008.
- [51] R. T. Sutherland, R. Srinivas, S. C. Burd, D. Leibfried, A. C. Wilson, D. J. Wineland, D. T.C. Allcock, D. H. Slichter, and S. B. Libby. Versatile laser-free trapped-ion entangling gates. *New Journal of Physics*, 21(3), mar 2019.
- [52] R. Srinivas, S. C. Burd, R. T. Sutherland, A. C. Wilson, D. J. Wineland, D. Leibfried, D. T.C. Allcock, and D. H. Slichter. Trapped-Ion Spin-Motion Coupling with Microwaves and a Near-Motional Oscillating Magnetic Field Gradient. *Physical Review Letters*, 122(16), apr 2019.
- [53] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A - Atomic, Molecular, and Optical Physics*, 61(6):8, 2000.
- [54] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous heating in ion traps. *Physical Review Letters*, 97(10):1–4, 2006.

-
- [55] Ivan A. Boldin, Alexander Kraft, and Christof Wunderlich. Measuring Anomalous Heating in a Planar Ion Trap with Variable Ion-Surface Separation. *Physical Review Letters*, 120(2):23201, 2018.
- [56] Benjamin Foulon, Keith G. Ray, Chang-Eun Kim, Yuan Liu, Brenda M. Rubenstein, and Vincenzo Lordi. How Correlated Adsorbate Dynamics on Realistic Substrates Can Give Rise to $1/w$ Electric-Field Noise in Surface Ion Traps. *Phys. Rev. A*, 105(1):013107, jan 2022.
- [57] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, St Schulz, F. Schmidt-Kaler, and H. Häffner. Fabrication and heating rate study of microscopic surface electrode ion traps (New Journal of Physics (2011) 13 (013032)). *New Journal of Physics*, 14, 2012.
- [58] H. G. Dehmelt. Radiofrequency Spectroscopy of Stored Ions I: Storage. *Advances in Atomic, Molecular and Optical Physics*, 3:53–72, 1968.
- [59] S. Seidelin, J. Chiaverini, R. Reichle, J. J. Bollinger, D. Leibfried, J. Britton, J. H. Wesenberg, R. B. Blakestad, R. J. Epstein, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Physical Review Letters*, 96(25):1–4, 2006.
- [60] T. G. Ballance, J. F. Goodwin, B. Nichol, L. J. Stephenson, C. J. Ballance, and D. M. Lucas. A short response time atomic source for trapped ion experiments. *Review of Scientific Instruments*, 89(5), 2018.
- [61] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J. P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-

-
- selective photoionization for calcium ion trapping. *Physical Review A - Atomic, Molecular, and Optical Physics*, 69(1):13, 2004.
- [62] M. A. Weber, C. Löschnauer, J. Wolf, M. F. Gely, R. K. Hanley, J. F. Goodwin, C. J. Ballance, T. P. Harty, and D. M. Lucas. Cryogenic ion trap system for high-fidelity near-field microwave-driven quantum logic. *arXiv:quant-ph*, jul 2022.
- [63] Jochen Wolf. *Cryogenic, near-field quantum logic chips with passive field nulling on $^{43}\text{Ca}^+$* . PhD thesis, University of Oxford, 2019.
- [64] B Keitch. *A quantum memory qubit in Calcium-43*. PhD thesis, University of Oxford, 2007.
- [65] Clemens Löschnauer. *A cryogenic trap for quantum logic using $^{43}\text{Ca}^+$ ions*. PhD thesis, University of Oxford, 2021.
- [66] Hugo Alexander Janacek. *Optical Bloch equations for simulating trapped-ion qubits*. PhD thesis, University of Oxford, 2015.
- [67] D. T.C. Allcock, T. P. Harty, M. A. Sepiol, H. A. Janacek, C. J. Ballance, A. M. Steane, D. M. Lucas, and D. N. Stacey. Dark-resonance Doppler cooling and high fluorescence in trapped Ca-43 ions at intermediate magnetic field. *New Journal of Physics*, 18(2), feb 2016.
- [68] Thomas P. Harty. *High-Fidelity Quantum Logic in Intermediate-Field $^{43}\text{Ca}^+$* . PhD thesis, University of Oxford, 2013.
- [69] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T.C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100(20):2–5, 2008.

-
- [70] D Szwerc. *High Fidelity Readout and Protection of a $^{43}\text{Ca}^+$ Trapped Ion Qubit*. PhD thesis, University of Oxford, 2010.
- [71] Alice Heather Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, University of Oxford, 2010.
- [72] Christopher E. Langer. *High Fidelity Quantum Information Processing with Trapped Ions*. PhD thesis, University of Colorado, 2006.
- [73] Klaus Mølmer and Anders Sørensen. Multiparticle entanglement of hot trapped ions. *Physical Review Letters*, 82(9):1835–1838, 1999.
- [74] M. A. Sepiol. *A high-fidelity microwave driven two-qubit quantum logic gate in $^{43}\text{Ca}^+$* . PhD thesis, University of Oxford, 2016.
- [75] M. V. Berry. Quantal phase factors accompanying adiabatic changes. *A Half-century of Physical Asymptotics and Other Diversions: Selected Works By Michael Berry*, 392:45–57, 1984.
- [76] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(6930):412–415, 2003.
- [77] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt. Process tomography of ion trap quantum gates. *Physical Review Letters*, 97(22):1–4, 2006.
- [78] Joseph Emerson, Robert Alicki, and Karol Zyczkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10), 2005.

-
- [79] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Physical Review A - Atomic, Molecular, and Optical Physics*, 77(1):1–7, 2008.
- [80] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland. Randomized benchmarking of multiqubit gates. *Physical Review Letters*, 108(26):1–5, 2012.
- [81] Aaron Leu. To be published. Master’s thesis, ETH Zürich, 2022.
- [82] C. H. Baldwin, B. J. Bjork, J. P. Gaebler, D. Hayes, and D. Stack. Subspace benchmarking high-fidelity entangling operations with trapped ions. *Physical Review Research*, 2(1), mar 2020.
- [83] Jahan Claes, Eleanor Rieffel, and Zihui Wang. Character Randomized Benchmarking for Non-Multiplicity-Free Groups with Applications to Subspace, Leakage, and Matchgate Randomized Benchmarking. *PRX Quantum*, 2(1), jan 2021.
- [84] C. Marquet, F. Schmidt-Kaler, and D. F.V. James. Phonon-phonon interactions due to non-linear effects in a linear ion trap. *Applied Physics B: Lasers and Optics*, 76(3):199–208, mar 2003.
- [85] X. Rebecca Nie, Christian F. Roos, and Daniel F.V. James. Theory of cross phase modulation for the vibrational modes of trapped ions. *Physics Letters A*, 373(4):422–425, jan 2009.
- [86] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe. Coherent error suppression in multiqubit entangling gates. *Physical Review Letters*, 109(2), jul 2012.

- [87] A. Lemmer, A. Bermudez, and M. B. Plenio. Driven geometric phase gates with trapped ions. *New Journal of Physics*, 15, aug 2013.
- [88] I. Cohen, S. Weidt, W. K. Hensinger, and A. Retzker. Multi-qubit gate with trapped ions for microwave and laser-based implementation. *New Journal of Physics*, 17(4):43008, 2015.
- [89] David Hayes, Kaveh Khodjasteh, Lorenza Viola, and Michael J. Biercuk. Reducing sequencing complexity in dynamical quantum error suppression by Walsh modulation. *Physical Review A - Atomic, Molecular, and Optical Physics*, 84(6):1–13, 2011.
- [90] T. Harty, D. Main, D. Nadlinger, and M. Weber. Stabilizer mezzanine board. *github.com/OxfordIonTrapGroup/stabilizer_current_sense*, 2019.
- [91] Sinrara-Hardware-Collaboration. Sinara Hardware.
- [92] David R. Leibrandt and Jason Heidecker. An open source digital servo for AMO physics experiments. *arXiv:physics*, aug 2015.
- [93] R W P Drever, J L Hall, F V Kowalski, J Hough, G M Ford, A J Munley, and H Ward. Laser Phase and Frequency Stabilization Using an Optical Resonator. *Appl. Phys. B*, 31:97–105, 1983.
- [94] Sébastien Bourdeauducq, Robert Jördens, Peter Zotov, Joe Britton, Daniel Slichter, David Leibrandt, David Allcock, Aaron Hankin, Florent Kermarrec, Yann Sionneau, Raghavendra Srinivas, Ting Rei Tan, and Justin Bohnet. ARTIQ 1.0. *Zenodo*, may 2016.
- [95] David P. Nadlinger. ndscan - N-dimensional scans for ARTIQ. *github.com/OxfordIonTrapGroup/ndscan*, jun 2018.

-
- [96] Thomas P. Harty. Squareatron-5000 v1.0. github.com/OxfordIonTrapGroup/Squareatron-5000, 2019.
- [97] Jochen Wolf. Squaratron stability characterisation. *private communication*, 2019.
- [98] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, 2021.
- [99] Peter Maunz, Craig R Clark, Raymond Haltli, Andrew Hollowell, Jonathan Mizrahi, Paul Resnick, Jonathan D Sterk, Daniel L Stick, Boyan Tabakov, and Matthew G Blain. Characterization of a High-Optical-Access surface trap optimized for quantum information processing. www.osti.gov/biblio/1239095, 2015.
- [100] Colin D. Bruzewicz, Robert McConnell, John Chiaverini, and Jeremy M. Sage. Scalable loading of a two-dimensional trapped-ion array. *Nature Communications*, 7:1–6, 2016.
- [101] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. Controlling fast transport of cold trapped ions. *Physical Review Letters*, 109(8):1–5, 2012.
- [102] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. Coherent diabatic ion transport and separation in a multizone trap array. *Physical Review Letters*, 109(8), aug 2012.

- [103] Daniel Lobser, Matthew G. Blain, Craig William Hogle, Melissa Revelle, Daniel Lynn Stick, Christopher G Yale, and Peter Lukas Wilhelm Maunz. Quantum and Classical Control of Ions in Sandia’s HOA Trap. *NACTI*, 2017.
- [104] Vlad Negnevitsky. *Feedback-stabilised quantum states in a mixed-species ion system*. PhD thesis, ETH Zürich, 2018.
- [105] Kilian Singer, Ulrich Poschinger, Michael Murphy, Peter Ivanov, Frank Ziesel, Tommaso Calarco, and Ferdinand Schmidt-Kaler. Colloquium: Trapped ions as quantum bits: Essential numerical tools. *Reviews of Modern Physics*, 82(3):2609–2632, 2010.
- [106] Ludwig E. De Clercq. Transport Quantum Logic Gates for Trapped Ions. Master’s thesis, 2015.
- [107] Robin Oswald. Velocity Control of Trapped Ions for Transport Quantum Logic Gates. Master’s thesis, 2015.
- [108] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- [109] Madeleine Udell, Karanveer Mohan, David Zeng, Jenny Hong, Steven Diamond, and Stephen Boyd. Convex optimization in Julia. *Proceedings of HPTCDL 2014: 1st Workshop for High Performance Technical Computing in Dynamic Languages - Held in Conjunction with SC 2014: The International Conference for High Performance Computing, Networking, Storage and Analysis*, 0:18–28, 2014.

-
- [110] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga. Shortcuts to adiabaticity: Concepts, methods, and applications. *Reviews of Modern Physics*, 91(4), 2019.
- [111] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Physical Review A*, 78(6):1–12, 2008.
- [112] H. Kaufmann, T. Ruster, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. Dynamics and control of fast ion crystal splitting in segmented Paul traps. *New Journal of Physics*, 16(7):073012, jul 2014.
- [113] J. P. Home and A. M. Steane. Electrode configurations for fast separation of trapped ions. *Quantum Information and Computation*, 6(4-5):289–325, 2006.
- [114] Christopher J. Ballance. *High-Fidelity Quantum Logic in Ca+*. PhD thesis, University of Oxford, 2014.
- [115] C. H. Baldwin, B. J. Bjork, M. Foss-Feig, J. P. Gaebler, D. Hayes, M. G. Kokish, C. Langer, J. A. Sedlacek, D. Stack, and G. Vittorini. High-fidelity light-shift gate for clock-state qubits. *Physical Review A*, 103(1), jan 2021.
- [116] Melanie Windridge. Smaller and quicker with spherical tokamaks and high-temperature superconductors. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 377(2141), mar 2019.

Appendix

Appendix A

Quantum-Jump Enhanced SPAM Optimisation

As part of this work high fidelity state preparation and readout operations in $^{43}\text{Ca}^+$ at 288 G were developed. As noted in chapter 2, this method closely follows the methods used with $^{43}\text{Ca}^+$ at lower magnetic fields. When setting up the apparatus to perform state-preparation and readout operations two quantum-jump enhanced protocols were developed that allowed for optimisation of the optical polarisations of the 397 nm σ^+ and 850 nm beams, as well as setting of the 850 nm cavity reference.

A.1 $3\text{D}_{3/2}$ to $4\text{P}_{3/2}$ quantum jump enhanced spectroscopy

To successfully repump the $3\text{D}_{3/2}$ population to the $4\text{P}_{3/2}$, $|F = 5, m_f = +5\rangle$ state during readout operations the 850 nm laser must be tuned to spectrally address the targeted transition of the three 850 nm beams. To identify the correct laser frequency we used two related optical-shelving schemes. These




Figure A.1: $^{43}\text{Ca}^+$ level structure - Figure reproduced from Weber et al. [62].

schemes use a 397 nm and 850 nm beam to shelve ions in the $3D_{5/2}$ level. The application of the 397 nm excites the ions to the $4P_{1/2}$ level from which the ion may spontaneously decay to the $3D_{3/2}$ level or return back to the $4S_{1/2}$ ground level. When the ion decays to the $3D_{3/2}$ level, it can be excited to the $4P_{3/2}$ level via the 850 nm beam. Spontaneous decay from the $4P_{3/2}$ level in turn allows the ion to become shelved in the $3D_{5/2}$ level. The rate at which this shelving process occurs depends on only the 397 nm driven rate of excitation to the $4P_{1/2}$ level and the 850 nm driven excitation to the $4P_{3/2}$ level. All other transition rates are fixed by the properties of the $^{43}\text{Ca}^+$ ion.

Within our apparatus we have both the Doppler cooling and σ^+ polarised optical pumping 397 nm beams. Using the Doppler beam for excitation to the $4P_{1/2}$ level results in the greatest shelving rate and populates many of the hyperfine states within the $3D_{3/2}$ level. On the other hand, using the σ^+ polarised optical pumping beam results in self terminating excitation due to

optical pumping to the $4S_{1/2}, |F = 4, m = +4\rangle$ state. However, we can drive we can readily produce ~ 5 MHz Rabi flops on the $4S_{1/2}, |F = 4, m = +4\rangle \leftrightarrow |F = 3, m = +3\rangle$ transition using microwaves. Continuously driving such Rabi oscillations, allows the 397 nm optical pumping beam to continuously excite transitions to the $4P_{1/2}$ level at ~ 5 MHz. Though this approach significantly reduces the shelving rate when compared to using the Doppler cooling beam, using the optical pumping beam only significantly populates $3D_{3/2}$ states which can be reached by decay from the $4P_{1/2}, |F = 4, m = +4\rangle$ state. Consequently much fewer $3D_{3/2}$ levels are populated, increasing the relative signal of the transitions of interest for readout.

The 397 nm & 850 nm scattering process results in the ion becoming shelved, which becomes more probable as the duration of the 397 nm and 850 nm pulses is increased. After application of such a pulse, shelving can be detected via the same fluorescence detection method used for qubit readout, see section 2.6.3. Figure A.2 shows a spectrum which was attained by shelving with the 397 nm optical pumping and 850 nm, π beam before the 850 nm, π polarisation was set.

Using this spectroscopy method, the desired 850 nm laser frequency can be identified from the spectrum. Further, we can set the desired 850 nm beam polarisations by minimising the features attributed to σ^- 850 nm transitions. Finally, the 850 nm beam alignment can be optimised by maximising the shelving rate.

A.2 Optical state-preparation optimisation

The above shelving protocol which uses the σ^+ polarised optical pumping 397 nm beams relies on leakage processes from the $4S_{1/2}, |F = 4, m = +4\rangle$ state. In the preceding section, such a leakage process was engineered by




Figure A.2: 850 nm **spectrum** - Measured quantum jump enhanced 850 nm spectrum using the optical pumping 397 nm beam. The x-axis shows the detuning of the 850 nm frequency relative to the $3D_{3/2} \leftrightarrow 4P_{3/2}$ transition centre. The y-axis indicates the measured probability of the ion being shelved after the application of the spectroscopy pulse. This spectrum was measured using the 850 nm, π beam with a misconfigured beam polarisation. The vertical bars indicate expected hyperfine transition frequencies. The bars are colour coded by transition polarisation. Further, their height is scaled by the product of the $4P_{1/2}$, $|F = 4, m = +4\rangle$ decay probability to the corresponding $3D_{3/2}$ state and the 850 nm transition matrix element. A selection of these expected transitions is labelled in the figure by the $|F, m\rangle$ quantum numbers of the associated hyperfine states.

driving the $4S_{1/2}$, $|F = 4, m = +4\rangle \leftrightarrow |F = 3, m = +3\rangle$ microwave transition. In the absence of such a microwave drive, shelving is contingent on other leakage processes from the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state. Such process include polarisation impurities of the σ^+ polarised 397 nm optical pumping beam and laser leakage of the higher power, non- σ^+ polarised 397 nm Doppler and Raman laser beams. Based of this diagnostic we were able to empirically optimise state preparation errors arising form polarisation impurities and 397 nm laser leakage.

After optimising laser leakage, as well as the 850 nm, and 397 nm, σ^+ beams via these spectroscopy methods we measured the optical state prepar-




Figure A.3: **Optical state preparation and shelving fidelity** - Measured error probability when optically pumping a single trapped ion to the $4S_{1/2}$, $|F = 4, m = +4\rangle$ state followed by optical shelving. a) Measured time series data showing the probability of not being shelved as inferred from 1000 measurements of the optical pumping and shelving procedure. b) Power spectral density of the auto-correlation of the time series data shown in a). Note the peak around 2.4 mHz. This corresponds to a 7 minute period and matches the duration of the laboratory air-conditioning cycle.




Figure A.4: **Microwave-enhanced state-preparation and shelving error** - Power spectral density of the of the auto-correlation of the time series data measured with microwave enhanced state preparation. Note the absence of the 2.4 mHz peak. The y-axis of this figure is not comparable to figure A.3.

ation and shelving error probability of $5.1(5) \times 10^{-4}$. This fidelity showed fluctuations which were correlated with the laboratory air conditioning cycle, see figure A.3. Enabling microwave enhanced state-preparation eliminated the correlation with the air-conditioning and improved the measured state preparation and shelving error probability to $4.8(2) \times 10^{-4}$, see figure A.4. As the error correlation with the air conditioning is resolved by microwave enhanced state preparation, we attribute the correlation to drift of the optical pumping beam polarisation with the laboratory temperature. We believe the remaining error probability to be limited by the optical shelving procedure.

Appendix B

Microwave Field and RF Resonance Post “Experiment Upset”

Between demonstration of the 331 μs , 99.51(11)% fidelity W7MDMS gate, and later demonstrations of 216.8 μs & 153.8 μs MS gates there was a major apparatus upset. This upset permanently degraded the stability of the RF field experienced by the ion. Further, this event also resulted in a modified microwave field and resonator properties. The precise nature of the upset is not fully understood at the time of writing. In the following we present some of the consequences of the event.

B.1 Spontaneous RF resonance jumps

After the apparatus upset, the resonance frequency of the trap and external RF resonator was measured to spontaneously jump between the pre-upset ~ 61.6 MHz resonance frequency and a new ~ 63.3 MHz. These two RF




Figure B.1: **RF resonance jumps after experimental upset** - Measured RF resonance spectrum before and after a spontaneous jump in the RF resonator frequency. The y-axis indicates the measured S_{11} parameter of the RF input to the RF resonator. This characterises the power reflection coefficient from the resonator input. The x-axis shows the applied RF frequency. After the upset, the RF resonance was measured to spontaneously jump between ~ 61.6 MHz and ~ 63.3 MHz. Such jumps resulted in ion loss and trapping was only possible after adjusting the RF frequency to match the new resonator frequency.

resonance statuses are referred to as apparatus A and B respectively. Such jumps were most frequently observed when performing temperature changes of the trap pillbox. However, a jump was also observed after approximately one day of operating the trap at ~ 27 K. Whenever such a jump occurred all trapped ions would be lost. Re-trapping ions was only possible after matching the RF drive frequency to the new RF resonance frequency. The measured RF resonance before and after such a jump is shown in figure B.1.

B.2 Mode frequency variation

After the apparatus upset the stability of the ion motional mode frequency was significantly degraded. This degradation took two forms. Firstly, spontaneous kHz level jumps of the mode frequency were observed after the upset, see figure B.2. Secondly, significantly larger mode frequency drifts were observed when applying high power microwave pulses.

B.3 Shift of microwave null

The RF resonance frequency jumps presented in section B.1 coincided with jumps of the trap microwave field. Notably the position of the partial microwave null, see section 4.1.1, shifted between being ~ 5.5 μm displaced from the RF null and being only ~ 2.2 μm displaced. The ~ 5.5 μm to ~ 2.2 μm null positions coincided with the RF resonance being at ~ 61.6 MHz and ~ 63.3 MHz respectively. The shift of the microwave field is shown in figures B.3 and B.4.

The shift of the microwave null position to 2.2 μm displacement from the RF null brings the measured microwave null position more in line with the ~ 2 μm displaced null position targeted in the design of the microwave in-




Figure B.2: **Post-upset mode frequency variation** - Measured mode frequency drift after the apparatus upset. The measurement series was performed on a single ion when not applying high power microwave pulses and without re-loading the ion. The measurement was performed when the RF resonance was in configuration A. The y-axis shows the measured mode frequency. The x-axis indicates the elapsed time since the start of measurement sequence. Note the discontinuous mode frequency jumps at ~ 7500 s and ~ 13000 s. When not applying high power microwave pulses, the mode frequency drift is around $\sim 30 \text{ Hz min}^{-1}$. However, after the apparatus upset, the mode frequency was significantly degraded when applying high power microwaves.




Figure B.3: **Pre-upset microwave field** - Microwave magnetic field, as measured by observing Rabi oscillation of a single ion within a potential well which is horizontally displaced from the RF-null. The green, orange, and blue circles represent the measured microwave field amplitude perpendicular ($|B_{\perp}|$, upper plot) and parallel ($|B_{\parallel}|$, lower plot) to the traps surface at trap temperatures of $T = 300$ K, 77 K, and 21K respectively. The x-axis indicates the in-plane radial displacement of the potential well from the RF-null. The measured microwave field amplitude is normalised for an input power of 1W. Figure reproduced from Weber et al. [62].




Figure B.4: **Post-jump microwave field** - Microwave magnetic field, as measured by observing Rabi oscillation of a single ion within a potential well which is horizontally displaced from the RF-null. The x-axis indicates the in-plane radial displacement of the potential well from the RF-null. The y-axis represents the measured microwave field amplitude parallel (B_{\parallel}) to the trap's surface at a trap temperature of 27 K. Notably, the $|B_{\parallel}|$ minimum has shifted from $\sim 5.5 \mu\text{m}$ to $\sim 2.2 \mu\text{m}$ displacement from the RF null when compared to the pre-upset field.

tegrated ion trap [62]. Weber et al. [62] proposed that the larger pre-upset null displacement may have been caused by resonant coupling of microwaves into the trap RF electrode. RF resonance jumps coinciding with shifts of the microwave null is consistent with this hypothesis and suggests the jump may be caused by changes of RF components inside the trap pillbox which are implicated in the formation of the modified microwave resonance.

The 216.8 μs & 153.8 μs MS gate demonstrations were performed when the microwave null was at $\sim 2.2 \mu\text{m}$. Due to the reduced null displacement from the RF null and ion position, the ion experiences a reduced π polarised microwave field during these gates. Consequently, the AC-Zeeman shift resulting from drifts in the microwave sideband power balance is reduced. This likely explains the significant AC-Zeeman shift drift reduction observed when performing these gates. The reduced shift sensitivity reduced the qubit frequency uncertainty sufficiently to make dynamical decoupling unnecessary. As dynamical decoupling was no longer required the gates performed post-upset adopted a standard MS scheme instead of the previous WMDMS scheme.

B.4 Shift of microwave resonance

After the upset, intermittent jumps in the Rabi frequency generated by a given microwave input power were observed. In particular, the Rabi frequency of the $4\text{S}_{1/2}$, $|F = 4, m = +2\rangle$ to $|F = 3, m = +1\rangle$ transition was measured to vary by as much as 10%. This is likely related to this transition frequency being on the edge of the microwave resonance. Consequently, the resonant enhancement of the signal addressed to this transition is particularly sensitive to changes in the resonator properties.

Appendix C

Fastino Output Noise

Characterisation & Kasli

Ground Loops

The following is adapted from the characterisation results I presented on the Sinara-HW issue discussion for Fastino noise, see <https://github.com/sinara-hw/Fastino/issues/56>.


C.1 Measurement setup

- I used the Agilent E4405B spectrum analyser with a 100 kHz DC-block and, where appropriate, a 1.9 MHz low-pass filter. Notably, the spectrum analyser input is grounded. Likely use-cases of Fastino also involve a grounded load. It is therefore important to know Fastino performance in such a configuration. Plots use a 1 kHz bandwidth and 10 sample averaging.
- Fastino v1.0 is controlled by a Kasli v1.1

- The spur performance is very grounding dependent. Care must be taken when comparing to measurements with different grounding configurations. For these measurements, the spectrum analyser and Kasli are powered from the same extension lead.
- Unless specified otherwise, measurements were made using an IDC-BNC adapter board. The board outputs are filtered with common mode current chokes.
- Measurements were performed with the board grounds not connected to a case and significant space between boards

C.2 Initial performance


Using a grounded power supply and no modifications to Fastino, there are many spurs in the channels 24-31.



We can see many spurs throughout the scan. It is worth noting that many spurs are only visible in some channels and their respective intensity varies between channels.

C.3 Fixes to noise sources which are part of the Fastino design

Issues #53 (Fastino switch-mode supply housing not grounded) and #54 (Oscillating Fastino voltage rail) are addressed as discussed in the respective issue threads. This removes some of the spurs. Looking at channel 31 the effects are demonstrated.



These issues have been resolved in subsequent Fastino revisions.


C.4 Kasli switch-mode noise (with common mode current chokes in Fastino outputs)

The switch mode supply on Kasli (Kasli IC6) operates at ~ 450 kHz. The supply generates many harmonics and sub harmonics at half the harmonic intervals. This noise is filtered on Kasli (for example there is a big drop-off across L6). However, the noise still couples to all rails across Kasli. In particular the switch-mode noise couples strongly to the Kasli P12V0 rail. Kasli is not a low noise design and the behaviour may therefore be expected.

The coupling of the Kasli switch-mode noise to Fastino outputs is heavily grounding dependent and accounts for all remaining spurs within the frequency scan. The coupling of these spurs is Fastino output channel dependent. Below, channel 31 is shown as it was particularly sensitive.

- Unlike the other switch-mode spurs on Kasli, the ~ 900 kHz 1st switch mode harmonic is larger on the Kasli side of FL1 than the PSU side. (Measurement relative to PSU ground, be sure to have no ground loop as it changes the measurement)
- Adding 2.2 uF (X5R) capacitors on both sides of FL1 on Kasli eliminates the ~ 900 kHz spur. (May be worth discussing for future Kasli revisions)


C.4. Kasli switch-mode noise (with common mode current chokes in Fastino outputs)



- Using the TTi EL302RT floated power supply to power Kasli eliminates the spurs at 450 kHz and below. (See summary plot below)
 - The higher frequency switch-mode spurs are likely not effected due to capacitive coupling between PSU and building ground.
- Filtering the PSU power to Kasli suppresses the switch-mode harmonics above 1 MHz.
 - The PSU was filtered using the 8119-RC common mode choke with a 1uF capacitor each side.
 - This filter should break the ground loop at higher frequencies.

The effect of the power supply modifications above are summarised in the plot below.

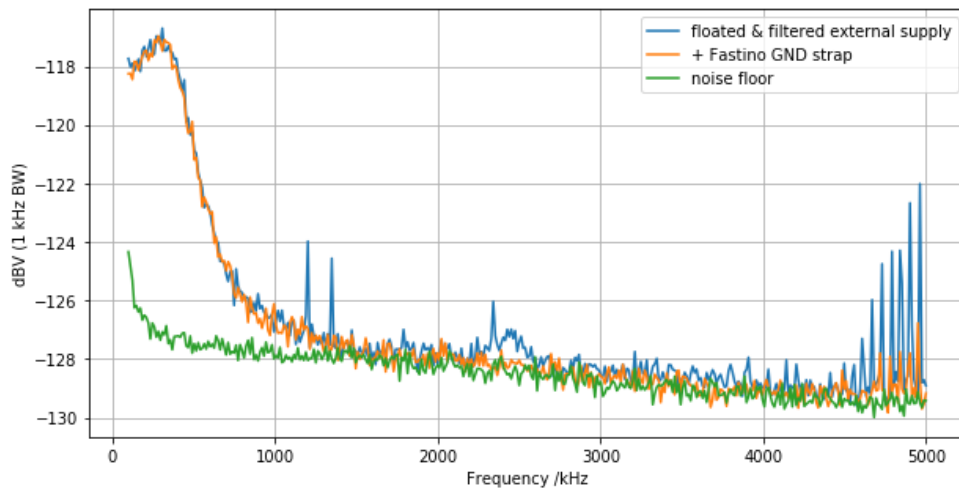
C.4. Kasli switch-mode noise (with common mode current chokes in Fastino outputs)



In the plot above, the 1.2 MHz spur is pickup from the environment.

- Adding a ground strap between the Fastino front-panel or Kasli ground and the Spectrum-analyser/building ground, suppresses all spurs except the ~900 kHz spur.

Combining this with a ground strap between the Fastino front panel and the spectrum analyser ground, the following spectrum is observed for channel 31:



Note that this spectrum goes from 100 kHz to 5MHz. Previous spectra were from 100 kHz to 1.5 MHz.

C.4.1 Modifications that didn't help reduce spurs

- Filtering the power lines and ground from Kasli to Fastino
- Connecting PSU ground to the analyser ground
- Ground line from Kasli ground to Fastino ground

C.5 Kasli switch-mode noise (without common mode current chokes in Fastino outputs)

- Spur free performance is achieved by using the Fastino fixes, a floated and filtered power-supply for Kasli and the capacitors on FL1 on Kasli.
 - This remains spur-free if the power-supply is grounded, but still filtered.
 - An unfiltered, grounded supply gives poor spur performance.
- Adding a ground connection from Kasli or Fastino ground to the analyser ground results in many spurs.
- Spurs intensities are very dependent on relative positioning on components.
- Spurs are not removed by cutting all digital lines and adding multiple filters onto the P3V3MP and P12V0 supplies.

C.6 Summary

- Low frequency noise on Fastino can be improved by applying the fixes for the issues IC3 (CC10-1212DF-E) housing not grounded #53 and N12V0A oscillating on some boards #54.

- Spur-free performance requires the user to be cautious in their grounding approach and some minor Kasli modifications.
- There are many ways Fastino could be configured to be spur free.
- It may be worth recommending Kasli be used with a class II power supply.
- Common mode chokes are helpful in some and detrimental in other grounding configurations.
 - When driving a grounded load and there is a ground path from Kasli/Fastino ground, common mode chokes are desirable.
 - When there is no ground path from Fastino, common mode chokes are undesirable.
- Operating Fastino with common mode chokes and a ground strap achieves good performance and is not as sensitive to the configuration details.

Appendix D

Accounting for Trap Filters in Potential Well Evolution

D.1 Compensating DAC voltages for trap filters

From the SURF README (written by me):

D.1.1 On Trap Electrode Filters

Ion trap electrodes are regularly filtered to reduce ion heating. Commonly, these filters are low-pass filters (though more complex filters are also used). If the well evolution timescales approach those modified by the filter, this should be accounted for. One may desire to include the filter effect in the solver cost function. However, this approach leads to more complex cost function structures. This may significantly increase solver times. An alternative approach is to use SURF without knowledge of the trap filters. In this case, the voltages returned by SURF represent the voltage on the trap electrodes. However, the desired external drive voltages can readily be found by inverting the filter. In many cases the required drive voltage changes will be larger

than the electrode voltage changes. You may therefore wish to increase the cost-function voltage-step-weight. This way increased slew-rate demands can be compensated for in the solver.

D.1.2 Inverting the Filter

Trap filters can usually be represented by a transfer-function $G(w)$. $V(w)/D(w) = G(w)$ w is the angular frequency, $V(w)$ the trap-electrode-voltage complex-amplitude and $D(w)$ the driving-voltage complex amplitude. Re-arranging gives: $D(w) = V(w) / G(w)$ Inverse Fourier transforming, we attain $D(t)$ We can therefore derive $D(t)$ from $V(t)$ through discrete Fourier transforms: $D(t) = \text{Inverse_Fourier_transform}(\text{Fourier_transform}(V(t)) / G(w))$

D.1.3 Notes on Fourier Transform

The following advice assumes causal trap filters. More care should be taken for non-causal filters.

- $V(t)$ will need to be padded to attain good results.
 - The filter will have an associated settling time. The end of the desired $V(t)$ should be padded with $V(t_{\text{end}})$. The padding time must be long enough for the electrode to settle sufficiently. This additional sequence time must be included in the drive voltage playback.
 - Many DFT implementations assume the input data window to be periodic. This leads to a large voltage glitch at the start of the calculated drive sequence. The start of $V(t)$ should therefore be padded with $V(t_{\text{start}})$. This part of the calculated sequence should be discarded in the playback.

Appendix E

Rules of Thumb for Potential-Well Solver Parameters

Here we present rules for setting sensible weights in the static and dynamic cost functions described in chapter 5. The driving factors behind these rules are two-fold. Firstly, we wish to set relative weights between cost terms such that the cost function produces reasonable trade-offs between different potential well properties. Secondly, the cost function is represented using floating point numbers. It is best practice to keep the value of these near unity when performing optimisation.

1. When setting the weights for potential well derivatives, one should pick weights such that acceptable parameter deviations result in cost terms near unity.
2. The voltage penalising cost term r_0 is not strictly necessary as the solver already respects hardware voltage limits. However, railing electrode

voltages may be experimentally undesirable. In such cases r_0 may be scaled to avoid railing electrode voltages. Generally, a sensible value results in a cost term near unity for acceptable electrode voltages.

3. To set the voltage slew rate parameter r_1 one should establish what voltage difference corresponds to a significant deviation from the target potential (depends on parameter tolerances). The characteristic “bad” slew rate is calculated by multiplying this voltage-change with the trap filter frequency. We can now set r_1 such that this “bad” slew rate is around unity.
4. The voltage curvature penalising term r_2 may be used to avoid ringing for inductive filters. Suitable values of r_2 may be found analogously to r_1 .

When using the ion separation trajectories there is an additional quartic potential cost term. This should be set such that the cost function contribution of the maximum achievable quartic term is around unity.