


Centre for Doctoral Training in Controlled Quantum Dynamics
Department of Physics, Imperial College London

Probing qubit memory errors at the 10^{-5} level



James Edward Tarlton

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of Imperial College

Declaration and copyright

Except when otherwise acknowledged and referenced, the following thesis is my own original work that took place during my PhD at Imperial College. It has not been already submitted to satisfy any degree requirement at this or any other university.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

James Tarlton

February 2018

Abstract

Trapped atomic ions are a promising candidate system for developing a quantum computer. All elementary operations for quantum information processing have already been achieved in trapped-ion systems with errors below the “fault-tolerant threshold”, meaning that these errors are sufficiently small to be removed with quantum error correction. However, the preservation of quantum information over the short timescales that are relevant for quantum computing has not, to our knowledge, been measured before in trapped-ion qubits. Previous investigations have measured decoherence over long times and extrapolated an exponential decay model to short times.

We directly measure the qubit memory error rate over short timescales, which is made possible by the high-fidelity single-qubit gates and state preparation and measurement in our system. Our qubit is a hyperfine “atomic clock” transition in $^{43}\text{Ca}^+$, and gates are applied with near-field microwaves. We find that the assumption of exponential decay of qubit state coherence does not apply in our system. The data is fit well by a model of $1/f$ frequency noise with a low-frequency cutoff. The smallest memory error that we measure is $3(1)\times 10^{-6}$ over 2ms. The time at which our qubit memory error is 1×10^{-3} , an important threshold for viable quantum error correction, is $\approx 400\text{ms}$.

We also present a design study into a new ion trap electrode geometry for applying near-field microwave two-qubit gates. This design features an ‘S’-shaped meander electrode to passively null the microwave field at the ions’ locations. It has improved isolation between the meander and other electrodes compared to previous work on this, which should reduce the sensitivity of the microwave field distribution to the boundary conditions of these electrodes. We show that the trap would allow for ion chains to be trapped, transported and split with feasible DC and RF voltages.

Acknowledgements

I would like to start by thanking my supervisors, Prof. David Lucas, Prof. Richard Thompson, and Prof. Danny Segal.

The content of this thesis is in no small part down to David's guidance, which was always helpful and a big factor in defining my time as a PhD student. Through his insight, patience and enthusiasm in our discussions, he helped improve my understanding and my decisions. I am particularly grateful for his support in the interruption of my studies for an internship, which was an invaluable experience for me.

Though we were based in different cities and so had limited opportunities to meet, Richard was very supportive throughout my studies. He gave me a fresh perspective on things, and was always checking on my progress and on my well-being, lending me a sympathetic ear whenever I had concerns about anything.

Despite his illness, Danny insisted on supporting me in my studies to the best of his ability, even into his final months. He was very kind and enthusiastic, and is perhaps the most compelling argument I have encountered that a cure for cancer cannot come soon enough.

Dr Thomas Harty took me under his wing from the start, and was a mentor to me for my most formative time as a PhD student. Through his knowledge, passion and attention to detail, he taught me a great deal, with many of my technical abilities having come directly from him. In particular, his supervision through my trap design work kept me sane and targeted, and the work would not have been possible without him.

Dr Joseph Goodwin took me on after my internship, and helped me through the difficult process of reacclimatizing to life as an ion trapper. His ambitious targets ensured that we never lost momentum, and he encouraged me to take on more

responsibility than I had thought I was capable of, which has been great for my personal development.

Dr Martin Sepiol was an absolute pleasure to work with over the years. One of the funniest, kindest people I know, he made our lab an enjoyable place to work, and he was a key person in keeping it functioning at all. He set a challenging example of work ethic, and was always happy to keep an “ion” things for me.

I never worked on the same project as Dr Christopher Ballance, but that did not stop him from being very helpful at times, in particular with my work on programming microprocessors. He was always generous with his time, and he would happily explain things to me when I had got the wrong end of the stick or my approach to a problem was less than ideal.

There were many others who contributed significantly to this work. Dr Timothy Ballance was instrumental in bending ARTIQ to our will. David “D5” Nadlinger introduced me to composite gates, and did most of the coding behind their implementation. Amy Hughes developed key bits of code for our experiments. Benjamin Merkel did much of the developmental work for our magnetic field feedback system. Dr Diana Aude Craik helped me to understand the gory details of how HFSS works. Prof. Andrew Steane had several interesting discussions with us about our experiments. Dr Hugo Janacek taught me how our dark resonance cooling works. Prof. Derek Stacey instructed me in giving presentations on my research.

For the rest of our group, their most important contribution was their companionship over the years. Jochen Wolf was an excellent source of distraction, and very supportive with anything from computer issues to problems with my landlady. Keshav Thirumalai was always great company, appreciative of all but the worst of my jokes, and very generous with his floor space. Dr. Norbert Linke, Vera Schäfer, Sarah Woodrow, Laurent Stephenson, Clemens Löschnauer, Bethan Nichol, Maciej Malinowski, Mathieu Dos Santos, Rustin Nourshargh, Caleb Rich, Daniel Shipley, Jack Segal and Marius Weber made up the rest of our little community, some for only a few months and others for much longer. I cannot imagine a better bunch of people to lunch with every day for three years.

There were also countless support staff in the department and elsewhere in the university without whom none of this would have been possible. In particular, the team at Advanced Research Computing were enormously helpful with my trap

simulations, going way beyond the call of duty to get my various bits of code working on their machine. Also, Mathew Newport was a great teacher in the workshop, making sure that I did not lose a single finger.

I have greatly valued the friendship and support of the other students in Cohort 5, and the time we spent together at home, in London, and in many other, more exotic, locations.

I am also grateful to the people that I shared time with outside work, for distracting me from my studies and making my time as a student more enjoyable. Thank you to everyone at frisbee for many a happy hour finding out what it really means to be out of breath. Thank you to Michael for the badminton, making sure that I exercised even when at my busiest, and for letting me win occasionally. Thank you to John Selby, Joe Munns, Lisa, Sonia, Tim, Ritu, Hitesh, Rashid, John and Jill for making my every house a home.

Thank you to my new friends in Bristol for keeping me going through what would otherwise have been two very lonely months.

Thank you to Lieve and everyone else in Brussels for giving me a tantalizing glimpse of what awaits me beyond my PhD.

Finally, thank you to Jill and John, my parents and my closest friends, for their unwavering and invaluable support, both moral and practical, throughout my studies.

Dedicated to the many PhD students suffering from mental health problems*

*See, for example, ‘Work organization and mental health problems in PhD students’, Levecque *et al.*, Research Policy, 46:868-879, 2017

Contents

Abstract	4
Acknowledgements	5
List of abbreviations	12
List of figures	13
List of tables	21
1 Introduction	22
1.1 Why develop a quantum computer?	22
1.2 Why use trapped ions?	28
1.3 Why use near-field microwaves?	31
1.4 Progress so far	33
1.5 Thesis outline	36
2 Theory and qubit system	38
2.1 Randomized benchmarking	38
2.2 BB1 composite gates	41
2.3 Phase noise	42
2.4 $^{43}\text{Ca}^+$ ions as qubits	45
2.4.1 Choice of qubit transition	46
2.4.2 Ionization	50
2.4.3 Laser cooling	51
2.4.4 Qubit initialisation	57
2.4.5 Quantum gates	57
2.4.6 Qubit readout	63

3	Apparatus	68
3.1	Ion trap and electrode connections	68
3.1.1	Ion trap	68
3.1.2	DC supply	72
3.1.3	Trapping RF	72
3.1.4	Diplexer	73
3.2	Vacuum system	74
3.3	Magnetic field coils	75
3.3.1	Main coils	75
3.3.2	Trim coils	76
3.4	Lasers, optics and imaging	77
3.4.1	Photo-ionization lasers	77
3.4.2	Ca ⁺ control lasers	77
3.4.3	Trap optics	79
3.4.4	Imaging system	80
3.5	Microwave systems	80
3.5.1	Frequency reference	80
3.5.2	Qubit microwaves	81
3.5.3	State preparation and readout microwaves	86
3.6	ARTIQ control system	86
4	Magnetic field control	88
4.1	Feedback	89
4.2	Mains electricity feedforward	94
4.3	Fluxgate magnetic field feedforward	97
4.4	Ion magnetic field servo	98
4.5	Field-independent point measurement	99
5	Short-timescale qubit decoherence	101
5.1	Ramsey experiments	103
5.2	Randomized benchmarking	107
5.2.1	$\pi/2$ -gates	109
5.2.2	BB1 composite gates	111
5.2.3	Interleaving delays	114
5.3	Decoherence suppression	120

6	Next-generation trap design	125
6.1	Preliminary decisions	129
6.2	Choice of electrode geometry	133
6.3	Choice of electrode widths	141
6.4	Simulation accuracy	146
7	Conclusions	152
7.1	Short-timescale qubit decoherence	152
7.2	Next-generation trap design	155
	Bibliography	158
A	Ground level structure of $^{43}\text{Ca}^+$	168
B	Magnetic field control circuits	171

List of abbreviations

Abbreviation	Meaning
AOM	Acousto-optic modulator
ARTIQ	Advanced Real-Time Infrastructure for Quantum physics
COM	Centre-of-mass
DAC	Digital to analogue converter
DDS	Direct digital synthesizer
FDA	Food and Drug Administration
PMT	Photomultiplier tube
RBM	Randomized benchmarking
SPAM	State preparation and measurement

List of Figures

1.1	The number of new drugs approved by the US Food and Drug Administration (FDA) per billion inflation-adjusted US dollars spent on R&D between 1950 and 2010. The fit is to Eroom’s Law, where the R&D efficiency halves every nine years. Figure from [6], with data from [8].	24
2.1	An illustration on the Bloch sphere of the pulse area error suppression achieved with the BB1 composite gate sequence. The initial state is $ 0\rangle$ and the target rotation is $\pi/2$ about the y -axis. There is a fractional undershoot of 0.2 on the raw gate. Each of the three pulses in the “do nothing” sequence have the same fractional undershoot, with the effect of these pulses being to almost completely correct the error in the raw gate.	43
2.2	The levels of $^{43}\text{Ca}^+$ that are relevant for the work in this thesis, with the resonance wavelength of each dipole-allowed transition. . . .	46
2.3	The energy splitting of the different states in the $4S_{1/2}$ ground level as a function of magnetic field. The clock transitions at 146G and 288G discussed in the main text are indicated with vertical lines.	49
2.4	The hyperfine structure of the $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at a magnetic field of 146G. The labels A , B , C , D , E and Q correspond to the transitions that are relevant for the work in this thesis.	50
2.5	The levels of ^{43}Ca and laser wavelengths that are relevant for our photo-ionization process. See section 3.4.1 for the details of the lasers used to address these transitions.	51
2.6	The states of $^{43}\text{Ca}^+$ and the laser wavelengths that are relevant for our Doppler cooling process.	53
2.7	The states of $^{43}\text{Ca}^+$ and the laser wavelengths and polarizations for the dark resonance that we use in our dark resonance cooling process.	54
2.8	The simulated $4P_{1/2}$ population as a function of the 397nm laser beam detuning. The blue curve is for our laser linewidths, and the red for monochromatic lasers. The black line indicates the location of the main dark resonance. Figure adapted from [62].	55

2.9	Part of the hyperfine structure of the $4S_{1/2}$ level (see figure 2.4 for the full structure), showing the D and E transitions, which are driven during the procedure to pump all population into $ F = 4, M = +4\rangle$.	58
2.10	Part of the hyperfine structure of the $4S_{1/2}$ level (see figure 2.4 for the full structure), showing the A , B , C and Q transitions, which are driven during the procedures to initialize and readout the qubit state.	59
2.11	The states of $^{43}\text{Ca}^+$ and the laser wavelengths that are relevant for our qubit state shelving process.	65
3.1	A block diagram overview of the apparatus.	69
3.2	The ion trap used in our experiments. (a) Photograph of the trap mounted on the aluminium holder. (b) CAD model of the trap, with labels indicating the integrated microwave elements. (c) Schematic of the central trapping region, showing the different signals that can be applied to each electrode. The near-field microwaves produced by the central electrode are almost completely π -polarized, which, for our σ -polarized qubit, is useful for generating entanglement (via the field gradient) without driving the qubit transition. We do not apply microwaves to this electrode for the single-qubit work presented in this thesis. $\mathbf{B}_0$ indicates the direction of our 146G static magnetic field. Figure adapted from [69].	70
3.3	A block diagram of the diplexer filter network for each trap electrode input. For the central electrode, which does not carry RF signals, the RF input is grounded.	73
3.4	A schematic of our vacuum system as viewed from above, showing the directions of the laser beams, calcium atom beam, and the static magnetic field. The 850nm (π) laser beam enters through the top viewport at $\sim 45^\circ$ to the trap surface.	75
3.5	Schematic of the imaging system (not to scale). The beam splitter is removed from the beam path during most experiments to maximize the PMT detection efficiency and thus the ion readout fidelity. .	81

3.6	A schematic view of our qubit microwave frequency octupler, with the part number for each component. All components are made by Mini-Circuits except for the isolator, which is made by Semtech. Most components were bolted to a metal plate, with semi-rigid SMA cable connections. The attenuators were included to allow for fine-tuning of the microwave power. Different attenuators were used for Ramsey experiment memory error measurements (see section 5.1) to give a different output microwave power. The switch configuration for that experiment was also different, with both switches in the ‘forward’ direction (switching between two outputs rather than two inputs) and no detuned microwave source (see section 3.5.2 for an explanation of why the switch configuration was changed for our randomized benchmarking experiments). All other components were the same for all experiments presented in this thesis.	82
3.7	Power transients on our 3.2GHz qubit microwave source for different switch configurations. The three configurations are with both switches in the forward configuration (switching the output between two ports rather than the input), and with the second switch in the backward configuration with and without a detuned microwave source connected to the other input of this switch. The three plots are scaled as a fraction of the final power and offset from one another by 0.01dB for clarity. The measurements were performed by connecting the microwave output to our diplexer via an ATC ATc3-6 circulator and measuring the reflected power from the third circulator port with a Mini-Circuits ZX47-40LN+ power detector connected to an oscilloscope averaging over 1000 sequences. The microwave power was turned on for $10\mu\text{s}$ pulses, with DDS profile switching times of $12\mu\text{s}$. We applied sequences of 32 of these pulses (to be roughly similar to our composite gate sequences, as described in section 5.2.2) with inter-sequence delays of 1ms. The first $1.6\mu\text{s}$ of each pulse measurement contained a large upward transient, which is likely an artefact of the measurement and is omitted for clarity. (a) The transients over the entire sequence. For the measurements without the detuned microwave source (blue and red), the transients over each gate extend above the upper plot limit and are omitted for clarity. (b) A close-up of the final microwave pulse showing the entirety of the fast transients.	84
4.1	An overview of our control system for the static 146G magnetic field.	89
4.2	An illustration of our main coil current feedback loop.	90

4.3	The spectral density of the magnetic field noise caused by noise in the coil current, as deduced from measurements of this current, with the feedback loop open and closed. The measurement is of the voltage on the feedback circuit after the reference subtraction, so it includes noise generated by the circuit before this point in the loop. There is a significant reduction in the noise up to the feedback loop bandwidth of 3kHz. The spike at $\sim 17\text{kHz}$ should be suppressed by filters later in the feedback loop as explained in the main text. The noise at harmonics of 50Hz is further suppressed with a feedforward system as described in section 4.2.	93
4.4	The decay in Ramsey fringe contrast on the $ 4, +4\rangle$ to $ 3, +3\rangle$ transition in the ground level with increasing Ramsey delay. Each Ramsey experiment is triggered by the mains electricity phase. The fit is to the Gaussian function $e^{-(t/T)^2}$, giving a decay time constant of $T = 0.70(6)\text{ms}$	94
4.5	Measurement of coherent magnetic field fluctuations over the 20ms mains electricity cycle, with and without feedforward variation of the magnetic field coil current. The measurement consists of a series of Ramsey experiments on the $ 4, +4\rangle$ to $ 3, +3\rangle$ transition at different times after being triggered by the mains electricity phase. The magnetic field shift at each time is then determined by the frequency shift of the central Ramsey fringe. The RMS magnetic field variation of the two data sets is $1.30(5)\text{mG}$ and $0.14(6)\text{mG}$ without and with the feedforward current respectively.	96
4.6	An example period of time illustrating the performance of the fluxgate magnetic field feedforward. The ambient magnetic field (as measured by the fluxgate sensor) changes significantly with the movement of a nearby lift. This change is automatically cancelled by making corresponding changes to the main coil current. Over this period, we also measure the field directly with the ion, making additional corrections to the coil current where the value set by the fluxgate servo is not quite correct. These two plots are offset from one another by 1mG for clarity.	98
4.7	An example measurement of the qubit field-independent frequency point. The magnetic field is scanned over the point which results in this field-independence. At each field point, the qubit frequency is measured with a Ramsey experiment, scanning the relative phase of the two $\pi/2$ -pulses to measure the shift in the Ramsey fringe peak. The quadratic fit to the data gives a magnetic field offset of $-0.06(2)\text{mG}$	100
5.1	An example set of results from a Ramsey experiment, showing how the population in $ 0\rangle$ varies with the phase of the second $\pi/2$ -pulse relative to the first. The Ramsey delay between the two $\pi/2$ -pulses was 3s.	104

5.2	An illustration on the Bloch sphere of the qubit state evolution during an example Ramsey experiment with a non-zero frequency detuning error. The initial state is $ 1\rangle$, which is followed by two $\pi/2$ -pulses that are of equal phase and separated by some Ramsey delay. During the inter-pulse delay, the detuning error causes the qubit state to rotate in the x - y plane, causing the final state to differ from the target state $ 0\rangle$	105
5.3	The drift of the detuning between the qubit transition and the microwaves over an example period of time. The measurements were taken with Ramsey experiments, using a Ramsey delay of 3s. . .	106
5.4	The results of our Ramsey experiment measurement of the qubit memory error. (a) The raw data. (b) The SPAM error, which was measured simultaneously to the Ramsey fringe contrast. (c) The data after correction for the SPAM error.	108
5.5	The results of an example calibration of the $\pi/2$ -gate duration. Each sequence contains 2000 Clifford gates, with 50 different random sequences used in the calibration. The sequence fidelities measured for a range of gate durations are fit to a quadratic curve, with the fit giving a calibrated optimum gate duration of $9.615(3)\mu\text{s}$	110
5.6	The drift of the calibrated $\pi/2$ -gate duration over an example period of time. The data is obtained by performing calibrations like that presented in figure 5.5, finding the gate duration which corresponds to a $\pi/2$ -rotation.	111
5.7	The results of our gate error measurement. The fidelities are for RBM sequences with different numbers of Clifford gates, with SPAM measurements interleaved between each RBM sequence. 50 different random sequences of Clifford gates were used for each data point. The gate time was recalibrated every ~ 3 hours. Fitting the SPAM-corrected data to equation 2.3 with the SPAM error set to zero gives an average Clifford gate error of $1.65(9)\times 10^{-6}$	112
5.8	Fidelity measurements of RBM sequences with normal $\pi/2$ -gates and with BB1 composite gates, both with and without a 1% reduction in the microwave amplitude. The average gate errors r from fitting to equation 2.3 are given in the plot legend. The SPAM error is not measured independently but is a floating parameter in each fit. 50 different random sequences of Clifford gates are used for each data point.	114
5.9	Gate sequences for the RBM experiments with and without interleaved delays τ . The RBM sequence consists of m randomized Clifford gates C_i , with a final gate C_{m+1} to randomly rotate the qubit state into either $ 0\rangle$ or $ 1\rangle$. The sequences without interleaved delays instead include the same total delay at the end of the sequence, when the qubit state is insensitive to phase noise, to ensure that the amount of ion heating is the same.	115

5.10	Average memory errors for a range of delay lengths. Also plotted are several inferred systematic contributions to the measured errors, which are explained in the main text. The data is fit to a combination of white noise (see equation 2.14) and $1/f$ frequency noise with a low-frequency cutoff (see equation 2.17). There is a data point at 1ms with an inferred memory error that is negative, so it does not appear on this logarithmic plot, but it is included in the fit. The fit has a χ^2 p -value of 0.035.	116
5.11	An illustration of our vernier technique to improve the combined frequency resolution of the synthesizer and the DDS (not to scale). By changing the DDS frequency by one step in one direction and the synthesizer frequency by 57 steps in the other direction, we change the radiation frequency by 2mHz, which is a smaller amount than that which is achieved by changing only one of the two frequency settings.	117
5.12	An illustration on the Bloch sphere of alternating the rotation axis of spin echo pulses. The initial state, and the target rotation axis for the spin echo pulses, are on the x -axis. A slight error in the relative rotation angle results in the qubit state walking away from the initial state.	122
5.13	The results of our spin echo dynamical decoupling Ramsey experiments. The axis of rotation for the spin echo pulses was the same as the target qubit state in the x - y plane of the Bloch sphere. The data is fit to an exponential decay curve, giving an effective coherence time of 1000(500)s. The SPAM error was measured simultaneously to the Ramsey experiments and subtracted from the raw data.	123
6.1	An illustration of how three electrodes carrying electric currents can be used to generate a magnetic field distribution that, at a particular point, has zero field and a non-zero field gradient.	126
6.2	The temperature dependence of the thermal conductivities of several materials used in microfabrication. Figure adapted from [57], with data taken from [78].	130
6.3	A basic illustration of a generic surface Paul trap showing the electrode widths used in equation 6.2. The red electrodes carry RF voltages and the black electrodes are at RF ground, but they may carry DC or microwave signals. The black electrodes may be split into multiple electrodes with little effect on the RF field.	135

6.4	Schematics of the trap electrode geometries that we investigate with simulations in HFSS. The blue meander electrode carries the microwave currents, the red electrodes carry RF voltages, and the black electrodes are connected to the ground plane. The grey electrode is connected to the ground plane further to the left of the figure, 2mm from the trap centre, with a $5\mu\text{m}$ gap separating it from the ground plane in the trapping region. The meander and RF electrodes are connected to their sources to the left and right of the diagrams respectively. The design also includes electrodes further away for applying microwave fields for single-qubit gates and DC fields for trapping. (d) is similar to the geometry used in [47, 48] and (g) is the geometry chosen in this design study.	136
6.5	Schematic of the trap electrode geometry including the second microwave electrode and DC electrodes (not to scale). For simplicity, the two sets of DC electrodes are each combined into single electrodes. The blue electrode carries the microwave currents for two-qubit gates, the red electrodes carry RF voltages, the black electrodes are connected to the ground plane, the green electrode carries the microwave currents for single-qubit gates, and the magenta electrodes carry DC voltages. The grey electrode is connected to the ground plane further to the left of the figure, 2mm from the trap centre, with a $5\mu\text{m}$ gap separating it from the ground plane in the trapping region. The microwave and RF electrodes are connected to their sources to the left and right of the diagrams respectively, with the DC electrodes connected to the top and bottom.	139
6.6	(a) The central region of the ion trap design as implemented in HFSS. Each bend of the meander electrode is 1.5mm long. The electrode immediately above the upper bend of the meander is connected to the ground plane to the left of the area shown, 2mm away from the trap centre. For simplicity, the two sets of DC electrodes are each combined into single 1mm-wide electrodes. (b) A close up of the right end of the meander with each electrode labelled by the names by which they are referred to in this thesis. Also denoted are the quantities w_c , the distance between the two RF electrodes, and x_{RF} , the horizontal distance between the lower RF electrode and the ions' positions.	140
6.7	A diagram of the DC electrode segmentation to either side of the trapping region. There are nine $150\mu\text{m}$ -wide electrodes to allow for three trapping zones, with the central electrode further segmented into fractions $[1/6 \ 2/3 \ 1/6]$ to allow for the splitting of ion chains in the central trapping zone.	143

6.8	Plots of how the trap depth and the RF and DC voltages required for trapping, as well as the largest DC voltages required to split pairs of ions, vary with the distance w_c between the two RF electrodes and the horizontal location of the RF field null x_{RF} as measured from the upper edge of the lower RF electrode. The radial and axial mode frequencies, radial trap axis directions, and stability parameter q are as chosen in section 6.1.	144
6.9	(a) Plots of the mismatch between the microwave and RF field minima in the horizontal (x) and vertical (y) directions as a function of x_{RF}/w_c and the width of the centre bend of the meander electrode, with the lines and surfaces representing linear interpolations between the simulated values. (b) The lines in parameter space along which the horizontal (blue) and vertical (red) distances between the two field nulls equal zero. The lower and upper bends of the meander electrode are $10\mu\text{m}$ thinner and $20\mu\text{m}$ wider respectively than the centre bend (see the main text for an explanation of these choices).	147
6.10	Plots showing how the microwave and RF field null coordinates change with the number of adaptive passes completed. x is the horizontal position and y is the height (with the trap axis in z). The first two adaptive passes give field distributions of poor resolution and are excluded for clarity. These simulations were carried out with 64GB of RAM, with the computer unable to increase the precision any further than for the final points shown here.	148
6.11	Plots showing how the distance between the microwave and RF field nulls changes, both in x (top) and in y (bottom), with the width of the gaps between the electrodes (left) and the thickness of the electrodes (right). The distances are defined as the coordinates of the microwave null minus those of the RF null.	149
6.12	Plots showing how the distance between the microwave and RF field null changes, both in x and in y , with the temperature of the trap. The only change made to the design between these simulations is an adjustment of the conductivity of the electrodes corresponding to the change in temperature. The distances are defined as the coordinates of the microwave null minus those of the RF null.	150
6.13	The trap electrode geometry with a ‘U’-shaped microwave electrode. The blue electrode carries the microwave currents, the red electrodes carry RF voltages, and the black electrodes are connected to the ground plane. The meander and RF electrodes are connected to their sources to the left and right of the diagrams respectively. The design also includes electrodes further away for applying DC fields for trapping.	151

List of Tables

6.1	Table of the optimized electrode widths for the design.	146
6.2	Table of the microwave magnetic field parameters for the electrode widths in table 6.1. The field gradient is given in two directions, that in which it is maximal, and that of the radial trap mode in which it has the largest projection; the latter of these quantities is what determines the two-qubit gate speed. The input power is 1W. . .	146
7.1	Summary of the results achieved in our experimental setup with an intermediate-field atomic clock qubit in trapped $^{43}\text{Ca}^+$ ions. . . .	155
A.1	The resonance frequencies, frequency dependence on the static magnetic field B , and coupling strengths for all magnetic dipole transitions in the ground $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at 146.0942G, with symbols as defined in equation 2.19. The second order frequency dependence on magnetic field for the $ F = 4, M = 0\rangle - F = 3, M = +1\rangle$ transition is 2.416kHz G^{-2}	169
A.2	The resonance frequencies, frequency dependence on the static magnetic field B , and coupling strengths for all magnetic dipole transitions in the ground $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at 287.7827G, with symbols as defined in equation 2.19. The second order frequency dependence on magnetic field for the $ F = 4, M = +1\rangle - F = 3, M = +1\rangle$ transition is 2.4514kHz G^{-2}	170

Chapter 1

Introduction

1.1 Why develop a quantum computer?

The modern computer encodes problems in a series of physical bits, each of which is in one of two possible states, and applies logic gates to these bits to perform the desired calculations. These machines are known as “classical computers” because their bits are restricted to being in classical states. Over the last few decades, there has been increasing interest in developing “quantum computers”, which would not have this same restriction. The quantum bits, or “qubits”, that make up a quantum computer can be in quantum superposition states, and they can also be entangled with one another. We can use single-qubit gates to manipulate the states of individual qubits, and two-qubit gates to interact with pairs of qubits to generate entanglement between them. With these two types of gate, it is possible to produce arbitrary quantum states, which may have entanglement shared between many qubits.

However, given the wide range of capabilities of classical computers, many people may find it difficult to think of computational problems relevant to our lives that we cannot already solve quickly and easily. So when we physicists enthuse over quantum computers, it may seem that we are offering a solution to a non-existent

problem. Moreover, progress so far has been limited. Quantum computers were first proposed in the 1980s [1, 2, 3], but the current state-of-the-art systems have only a handful of qubits¹. These systems are generally the size of a large room, so the idea that a quantum computer could one day be competitive with our existing classical computers may seem implausible.

We argue in section 1.4 that, although a useful quantum computer has not yet been developed, substantial progress has been made towards this goal. To illustrate how a quantum computer may provide tangible benefits to society, we consider the development of medicinal drugs.

The process of developing drugs for medical applications can be characterized in the following way: identify candidate drugs, test these drugs, and then sell the drugs that pass the tests. The average cost of this process per drug that makes it to market is more than £1 billion [5, 6, 7], with most of this money going into R&D. As figure 1.1 shows, the development cost per new drug has been increasing exponentially over the last few decades. “Eroom’s Law” states that the number of drugs developed per billion UK pounds halves roughly every nine years. This downward trend means that individual patients and/or the taxpayer need to pay more and more for new medicines, and if they do not then progress in drug development could grind to a halt.

The most important determinant of R&D efficiency is the attrition rate in clinical trials (the probability of candidate drugs failing to reach the market) [5], which has been increasing over the last few decades [9, 10]. There are several possible reasons for this trend: drugs may be getting inherently more difficult to develop (i.e. existing drugs represent the “low-hanging fruit” of the industry), the increasing effectiveness

¹D-Wave Systems Inc. claim to have already developed a useful quantum computer with the approach of “quantum annealing”. However, attempts to detect a “quantum speedup” over classical computers have been inconclusive [4], so we exclude this machine from the term “quantum computer” as used in this thesis.




Figure 1.1: The number of new drugs approved by the US Food and Drug Administration (FDA) per billion inflation-adjusted US dollars spent on R&D between 1950 and 2010. The fit is to Eroom's Law, where the R&D efficiency halves every nine years. Figure from [6], with data from [8].

of drugs may be raising the bar on how effective new drugs must be to make it to market, or companies may have been overestimating the utility of advances in basic research and overusing new methods emerging from this research [6]. However, this trend could potentially be stopped or reversed by improving the process by which candidate drugs are selected. While this process is highly developed and sophisticated, to a large extent it still relies on educated guesswork, because it is not possible to exactly simulate the behaviour of large molecules, such as drugs and their interactions with proteins, with classical computers.

To understand why this is the case, consider first a classical system consisting of n binary elements (i.e. elements with two possible states). The total number of possible states of this system is 2^n , so it may be represented by n bits of a classical computer. Increasing the number of binary elements in the system by 1 requires only 1 additional bit in the computer: the computational requirements scale linearly with the system size, so it is feasible to simulate increasingly large systems by simply increasing the size of the computer.

Consider next a system with significant quantum mechanical effects, such as a drug or a protein. Each of the n binary elements can now be in quantum superposition states, and they can be entangled with other elements. To describe the state of such a system, $2^n - 1$ different complex numbers are required, each of which must be encoded into some number of computational bits to achieve the required level of precision. Increasing the number of binary elements to $n + 1$ would require an additional 2^n complex numbers: the computational requirements scale exponentially with the system size, making it infeasible to simulate large quantum systems.

To demonstrate this infeasibility, we can estimate the computational requirements for a full simulation of a single human protein (which is important for predicting the effect of drugs, as they work by interacting with proteins). The median protein length in humans is 375 amino acids [11], with the amino acids having an average of 19 atoms [12]. Even if only one binary element is necessary to describe each atom, then the total number of elements is still $n \sim 7,000$, which would require $\sim 10^{2,000}$ complex numbers. As a comparison, the number of atoms in the observable universe is $\sim 10^{80}$ [13].

Such a system could, in principle, be exactly simulated with only 7,000 qubits. Efficiently expressing a molecule as a collection of binary elements is not a trivial task, but molecules as large as beryllium hydride, which contains three atoms, have been simulated with small quantum computers [14]. While such small molecules can be simulated with classical computers, demonstrations of quantum simulations indicate the potential of this approach for simulations of molecules as large as drugs and human proteins. These simulations could then allow us to develop a better understanding of how a candidate drug would interact with the proteins in our bodies, enabling us to make better selections of drugs to test in clinical trials.

The applications of quantum computers are not, however, limited to simulations of

other quantum systems. As first shown in [15], quantum superpositions and entanglement allow for certain computational problems to be solved exponentially faster than is possible with a classical computer.

Perhaps the simplest example of this is the Deutsch-Jozsa algorithm [16]. Consider some function f , which takes n bits as an input and, depending on the values of these bits, gives 0 or 1 as an output. This function is known to be either constant (always returning 0 or always returning 1) or balanced (returning 0 for half of the possible input states, and 1 for the other half). The goal is to determine whether the function is constant or balanced.

A classical computer would solve this problem by repeatedly evaluating f for different input states. In the best case scenario, the first two outputs are different, and f is found to be balanced after just two evaluations. In the worst case scenario, the first 2^{n-1} outputs are the same, with $2^{n-1} + 1$ evaluations of f being necessary to determine whether it is constant or balanced.

With a quantum computer, the problem can be solved with just a single evaluation of f . We first prepare $n + 1$ qubits into the following quantum superposition state:

$$\frac{1}{\sqrt{2^{n+1}}} \sum_{x=0}^{2^n-1} |x\rangle (|0\rangle - |1\rangle). \quad (1.1)$$

Next we implement f as a quantum oracle, which changes the state $|x\rangle |y\rangle$ to $|x\rangle |y \oplus f(x)\rangle$, where $\oplus$ means addition modulo 2. Applying this oracle gives:

$$\begin{aligned} & \frac{1}{\sqrt{2^{n+1}}} \sum_{x=0}^{2^n-1} |x\rangle (|f(x)\rangle - |1 \oplus f(x)\rangle) \\ &= \frac{1}{\sqrt{2^{n+1}}} \sum_{x=0}^{2^n-1} (-1)^{f(x)} |x\rangle (|0\rangle - |1\rangle). \end{aligned} \quad (1.2)$$

We next apply the Hadamard transformation H to each qubit, which in matrix notation is:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (1.3)$$

This results in the following state, where we now ignore the final qubit:

$$\frac{1}{2^n} \sum_{y=0}^{2^{n-1}} \sum_{x=0}^{2^{n-1}} (-1)^{f(x)+x \cdot y} |y\rangle, \quad (1.4)$$

where $x \cdot y$ is the sum of the bitwise product of x and y .

The probability of measuring this n -qubit state to be $|0\rangle$ is:

$$\left| \frac{1}{2^n} \sum_{x=0}^{2^{n-1}} (-1)^{f(x)} \right|^2. \quad (1.5)$$

This probability is equal to 1 if f is constant and 0 if f is balanced. Therefore, we can distinguish between the two by measuring the final state and seeing whether or not it is $|0\rangle$.

The Deutsch-Jozsa algorithm is of little practical use, but there are several other algorithms that would hugely reduce the time taken to solve certain common computational problems. These problems include searching databases [17], evaluating discrete logarithms and factorizing integers [18], the latter of which would also enable a quantum computer to break the commonly used RSA cryptosystem. All of these tasks can be solved efficiently by classical computers for small numbers, but, as for simulations of quantum systems, quickly become unsolvable as the problem size increases.

One might think that, regardless of how useful a quantum computer might be in principle, due to the fragility of quantum states, any large system of qubits will accumulate errors that make accurate computations impossible. These errors could be caused by unwanted interactions with the environment, or by faulty single- or two-qubit gates, which result in qubit states that differ from the target states. However, by encoding each logical qubit in many physical qubits, it is possible to correct for errors in any one of those physical qubits, so long as the error rate is below the “fault-tolerant threshold” [19]. To develop a quantum computer, one must ensure that error rates are below this threshold to enable quantum error correction to work.

1.2 Why use trapped ions?

There are certain requirements on any quantum system that we may use in a quantum computer. One formulation of these requirements is the DiVincenzo criteria [20]:

1. A scalable physical system with well characterized qubits
2. The ability to initialize the state of the qubits to a simple fiducial state
3. Long relevant decoherence times, much longer than the gate operation time
4. A “universal” set of quantum gates
5. A qubit-specific measurement capability

Hydrogen-like ions are a promising candidate system. Different energy levels of the outermost electron can be used for the qubit states (providing there are at least two stable or metastable levels available for this purpose), and the ions’ states can be manipulated with electromagnetic radiation resonant with the different transitions. The qubit state can be initialized and measured with optical methods.

However, isolating individual ions from their environment to achieve long memory times is not a trivial task. It is possible to trap atoms by using a static magnetic field to interact with field-sensitive transitions, with laser beams maintaining a low atom temperature through Doppler cooling. However, this sort of trap, known as a magneto-optical trap, is quite shallow, and generating entanglement between atoms in separate traps would be difficult.

By ionizing the atoms, one can instead use electric fields to exert a force on them. This Coulomb interaction is relatively strong, so one can apply large forces with modest voltages on nearby electrodes. However, it is not possible to trap a charged particle with static electric fields alone. To understand why this is the case, consider Gauss's Law in free space:

$$\nabla \cdot \mathbf{E} = 0, \tag{1.6}$$

where $\mathbf{E}$ is the electric field. To have a restoring force F_i towards the origin in each dimension x_i , we require:

$$\frac{\partial F_i}{\partial x_i} < 0. \tag{1.7}$$

For $\mathbf{F} = q\mathbf{E}$, where q is the ion charge, this requirement is inconsistent with Gauss's Law. It is possible to apply a field which traps the ion in one or two dimensions, but not all three.

One solution to this problem is to use an oscillating electric field, which, in two of the three dimensions, switches between trapping and anti-trapping fields, so that, at any given time, the ion is being pulled towards the trap centre in one direction and pushed away in the other. At any fixed location in this field, the force averaged over each oscillation cycle is zero. However, the ion is not at a fixed location: the

oscillating field induces an oscillation in the ion position. This motion results in a net force towards the trap centre. The ion can then be confined in the third dimension with a DC electric field.

Traps which confine charged particles with the use of oscillating electric fields are known as Paul traps, or RF traps (the oscillation frequency is generally RF) [21]. Another approach is to instead use static electric and magnetic fields, with this type of trap known as a Penning trap. While Penning traps are used widely in other fields, for quantum computing applications the large magnets required to produce the trapping field would be a big challenge for scaling the system up to large numbers of qubits. Also, the ions orbit around the trap centre rather than being at rest, which makes it more difficult to interact with individual ions. Therefore, Paul traps are generally used instead.

For the typical spacing of $\sim 10\mu\text{m}$ between trapped ions, the direct interaction between the electrons of adjacent ions is too weak to be used for entanglement generation, but we can instead use the Coulomb interaction for this purpose [22]. If we cool the ions to very low temperatures (we generally aim for a mean vibrational occupation of $\bar{n} < 1$), we can use radiation with the frequency detuned slightly away from the transition resonance to stimulate vibration by a single quantum, conditional upon the electronic state, therefore coupling the electronic and vibrational states. If multiple ions are trapped within the same harmonic potential well, then they have shared vibrational modes. These modes can be used as a bus to couple their electronic states, generating entanglement between the ions.

Another desirable feature of trapped ions is that (within each species) they are always identical. One would never (in principle) need to adjust the properties of the coupling radiation for each new ion added to the system. This is not true for artificial quantum systems, with each new qubit having different properties to the

others, which could be a significant challenge when attempting to scale these systems up to large numbers of qubits.

1.3 Why use near-field microwaves?

Radiation pulses used to generate superpositions and entanglement (also known as single- and two-qubit gates respectively, analogous to classical logic gates) generally consist of laser light. This light couples to electric dipole transitions of the ions. Single-qubit gates are performed by choosing the pulse area (the product of the Rabi frequency and the pulse duration) to drive Rabi flops, rotating the qubit state around the surface of the Bloch sphere² in the desired direction and by the desired amount. The radiation can be made to interact with the ions' vibrational modes via the field gradient over the spatial extent of the ions' wavefunctions, enabling two-qubit gates to be performed via the Coulomb interaction.

The use of laser light to perform quantum gates in small systems of ions has been very successful, but there are some major challenges with scaling this approach to large numbers of ions. There are stringent requirements on the stability of the direction, frequency, phase, and intensity of each beam. While these issues can be addressed in a small number of laser systems, these systems generally require regular maintenance. Scaling the system up to large numbers of qubits would require a similarly large number of these high-maintenance laser systems, or a small number of systems with much higher beam power and complex optical control to share each beam between lots of ions.

In contrast to lasers, the generation and control of microwaves is a well-established area of technology. Radiation of sufficient stability for quantum computing appli-

²The Bloch sphere is a way of visualising single-qubit states, where the poles correspond to $|0\rangle$ and $|1\rangle$ and the qubit state is represented by a vector from the origin to a point on the surface of, or inside, the sphere.

cations can be produced using equipment that has already been developed in areas such as telecommunications. However, the long wavelength of microwaves means that the field gradient of free-space radiation is too small to generate entanglement sufficiently quickly for use in quantum computing. Also, with the small separation between ions in the same trap, it is not possible to focus the radiation onto individual ions to perform single-qubit gates.

There are two approaches to overcoming these problems. One is to apply a static magnetic field gradient, which, in combination with the microwave radiation, allows for entanglement to be generated [23]. The resulting gradient in the Zeeman shift on the qubit transition frequencies can also be used to selectively address ions by adjusting the microwave radiation frequency [24, 25]. The qubit transition frequency must be first-order dependent upon magnetic field, making it sensitive to fluctuations in this field, but decoherence can to some extent be suppressed by dressing the states with an additional microwave field [26, 27].

There are two major difficulties with this approach. Firstly, the generation of sufficiently large external magnetic field gradients is a significant technical challenge, which would be more difficult for traps with large numbers of ions. Secondly, the application of a field gradient to an array of field-sensitive ion-qubits results in different phases accumulated for each qubit state. In principle, this could be corrected for by either applying microwave pulses that reverse these rotations, or keeping track of the different phases and adding corresponding phases to the applied gates. However, making such a correction accurately, especially if ions are being moved between traps and therefore having Zeeman shifts that change over time, would be a difficult task.

Another approach is to use “near-field” microwave radiation, where an oscillating current close to the ions is used to produce magnetic fields that drive the qubit

transitions [28]. The large field gradient required for entanglement generation now comes from the sharp decay in field amplitude with the distance from the oscillating current. This also allows for individual ions to be addressed in an array of traps, as it is possible to apply a strong field at the location of one ion with much weaker fields experienced by ions elsewhere in the system. These weaker fields can be cancelled out with secondary microwave currents near the ions which we do not want the primary current to interact with [29].

A major challenge with the use of near-field microwaves is how close the microwave currents must be to the ions. This distance must be tens of micrometres at most, otherwise either the microwave currents or the gate durations become infeasibly large. To achieve such a short distance, the microwave currents can be in electrodes integrated into a surface trap [30, 31]. These traps have all electrodes for producing the RF and DC trapping fields in a plane, with the ions suspended above this surface. The achievable electrode sizes and ion trapping heights are compatible with the near-field microwave approach.

Another challenge is with the shape of the microwave field distribution. Single-qubit gates are driven with microwave fields, while two-qubit gates are driven with microwave field gradients. To shape the microwave field such that the ions experience a field gradient without a large field is not a trivial task, but this can be achieved with careful consideration of the electrode geometry in the surface ion trap. We discuss this in more detail in chapter 6.

1.4 Progress so far

In this section, we briefly summarize the progress that has been made so far towards developing a trapped-ion quantum computer. One of the main challenges is reducing qubit errors to below the fault-tolerant threshold. Depending on the nature of the

errors and the quantum error correction code used, this threshold can be as high as $\approx 1\%$ [32]. However, in order to keep the number of physical qubits per logical qubit at a manageable level, we must achieve error rates much lower than this theoretical limit.

The lowest reported single-qubit gate error rate is $1.0(3)\times 10^{-6}$, which was achieved with near-field microwave radiation in our ion trap system [33]. This measurement was carried out with a randomized benchmarking sequence of Clifford gates, each of which was composed of an average of two $\pi/2$ -gates.

The lowest error rate achieved so far for two-qubit gates is 1×10^{-3} [34, 35]. These experiments used lasers to drive Raman transitions. We have achieved a comparable error rate of $3(1)\times 10^{-3}$ with near-field microwave radiation [36], which is impressive given how new this approach is: the first demonstration of this method was only in 2011 [31], compared to 1998 for laser-driven two-qubit gates [37].

In any quantum algorithm, quantum information will need to be stored in certain qubits while gates are being performed elsewhere in the system. Errors in this memory due to quantum decoherence are generally characterized by preparing superposition states and observing their decay over periods of seconds or longer, fitting the measured decoherence to an exponential decay model. The longest measured time constant of this decay, also called the “coherence time”, is $50(10)$ s, which was measured in our ion trap system [33]. This is much longer than the typical duration of gates, which is on the timescale of milliseconds or below [33, 34, 35, 36].

Arguably the main challenge faced in quantum computing is in scaling systems up to large numbers of qubits. Simple algorithms have been implemented with ion traps containing as many as five qubits [38]. However, as the number of ions increases, so too does the complexity of the vibrational mode structure, making it increasingly

difficult to perform two-qubit gates. Therefore, to scale up trapped-ion quantum computing to lots of qubits, it is necessary to generate entanglement between ions in different traps, with the quantum computer consisting of many individual traps.

There are two main approaches to solving this problem. One is to physically move ions between traps, enabling ions in different traps to be entangled with one another. Surface ion traps are well-suited to this idea: the surface can be extended to contain multiple trapping regions within the same vacuum system, using changing DC potentials to shuttle the ions around the surface [28, 39]. Ion trap chips with as many as 150 trapping zones have been developed [40]. Ions have been transported by $\sim 300\mu\text{m}$ with qubit state errors as low as 1×10^{-5} [41] and motional mode heating as low as 0.1 quanta [42, 43].

Another approach is to use photons to generate entanglement between ions in different vacuum systems. Each of a pair of ions in different traps can be made to emit a photon that is entangled with the ion's state. These photons can then be interfered with one another in such a way that the ions' states become entangled [44]. This process is probabilistic, and the success probability of each attempt to entangle ions may be low, but it can be repeated many times until a successful attempt is heralded. Subsequent entanglement distillation techniques can be used to improve the purity of the entangled ion state. A prototype node for such a network has been developed, with entanglement generated between an ion and a photon, as well as between ions of different species, with the idea being that one species can be used for quantum gates within each node and the other can be used for photonic links between nodes [45].

Both of these approaches to the scalability problem face significant challenges, and it is unclear whether either of them will enable the development of a large trapped-ion quantum computer. However, progress is continually being made towards this goal,

and no fundamental barrier has yet been encountered.

1.5 Thesis outline

The main focus of this thesis is on reducing and characterizing memory errors in trapped $^{43}\text{Ca}^+$ qubits. The decoherence of trapped-ion qubits is generally only measured over long time periods (seconds or tens of seconds), where the memory errors are much larger than the fault-tolerant threshold. Timescales at which the memory error is below this threshold, and so can be tolerated in quantum algorithms, have not, to our knowledge, previously been probed directly in trapped-ion qubits, and are instead inferred by extrapolation of an exponential decay model fitted to data at long times. We directly measure memory errors as low as $3(1)\times 10^{-6}$ over millisecond timescales, and show that the decoherence in our system does not fit a simple exponential model.

In addition to this experimental work, we present a design study on possible future electrode geometries for improving the fidelity of near-field microwave two-qubit gates. This study builds on the work presented in [46, 47, 48] towards producing large microwave field gradients without also producing large fields, where this is achieved with minimal experimental overhead to enable the approach to be scaled up to large systems. We investigate a new electrode geometry which could address problems encountered in previous traps.

The thesis is structured in the following way. Chapter 2 gives theoretical background on randomized benchmarking, BB1 composite gates and phase noise, which is relevant for work presented later in the thesis, and describes our $^{43}\text{Ca}^+$ qubit system. In chapter 3, we present our experimental setup. Chapter 4 presents our work on suppressing the magnetic field noise and drifts in our system, which can cause qubit decoherence and errors in the preparation and readout of the qubit state.

In chapter 5, we present our results on probing qubit memory errors over short timescales, with analysis and discussion of these results. Chapter 6 gives the details of our next-generation near-field microwave trap design study. Finally, in chapter 7 we conclude the thesis, summarizing our results and giving an outlook on possible future work.

Chapter 2

Theory and qubit system

In this chapter, we cover some background and theory that is relevant to the work in this thesis. Section 2.1 covers the basics of randomized benchmarking, a way of measuring small errors in quantum gates. In section 2.2 we discuss the BB1 composite gate, which makes our gates robust to fluctuations in radiation intensity. Section 2.3 describes the basics of phase noise and a few different noise models that are relevant for the experimental results presented in this thesis. Finally, in section 2.4, we describe how we use $^{43}\text{Ca}^+$ ions as qubits, which is mostly as presented in [49, 50].

2.1 Randomized benchmarking

An important element of research towards developing a quantum computer is characterizing the errors of quantum gates. These errors must be significantly smaller than the fault-tolerant threshold in order for quantum error correction to work, which allows for error-free computations to be carried out. Gate errors can, in general, be evaluated by preparing the qubits in some known state, applying the gate, and then measuring the final qubit state, repeating this process to evaluate the probability of the qubits being in the target state.

This method works well under the assumption that the error in preparing and reading out qubit states is much lower than the gate error. However, in trapped-ion quantum computation, the error of single-qubit gates is generally much lower than the state preparation and measurement (SPAM) error¹. This means that we need to apply many gates between each qubit state preparation and readout to amplify the gate error and thus make it measurable.

One might think that we can simply repeat the target gate lots of times between each state preparation and measurement rather than just once, and then divide the measured error by the number of gates. To understand why this approach would not generally work, consider the common case of there being an error in the pulse area, i.e. the product of the Rabi frequency and the gate time. Viewing the situation in the Bloch sphere, the Bloch vector representing the qubit state will rotate by an angle that is either smaller or larger than the desired value θ by some fraction ϵ . The corresponding gate error is quadratic in ϵ :

$$\begin{aligned} 1 - F &= 1 - \cos^2\left(\frac{\epsilon\theta}{2}\right) \\ &\approx \frac{1}{4}(\epsilon\theta)^2, \end{aligned} \tag{2.1}$$

where we have chosen the initial Bloch vector to be perpendicular to the rotation axis. Repeating this gate n times means increasing the target angle θ to $n\theta$, giving the following total error:

$$1 - F \approx \frac{1}{4}(\epsilon n\theta)^2. \tag{2.2}$$

¹From a computational perspective, it makes sense that this should be the case, as each computation will generally include many single-qubit gates but only prepare and readout the qubit states once.

This expression scales quadratically with n rather than linearly, so dividing it by n does not give the error of a single gate. Simply repeating gates many times allows their errors to combine in a coherent fashion, meaning that the sequence error is dependent not only on the average gate error, but also on the nature of that error.

This problem can be solved by using randomized benchmarking (RBM). The idea is to randomly select gates from some set rather than simply repeat the same one. This set is generally chosen to be the Clifford set, which, along with certain states and measurements, is sufficient for universal quantum computing² [51]. In addition, the RBM protocol using the Clifford set can be scaled up to large numbers of qubits [52].

If the final target state is randomized between $|0\rangle$ and $|1\rangle$, then the fidelity of a sequence of m randomized Clifford gates is [52]:

$$F = \frac{1}{2}((1 - 2r_{SPAM})p^m + 1), \quad (2.3)$$

where r_{SPAM} is the SPAM error and the average gate error rate r over all Clifford gates is related to the depolarizing parameter p by:

$$r = \frac{1}{2}(1 - p). \quad (2.4)$$

It is also possible to extend this protocol to probe not just the average error of the entire Clifford set, but also of some specific Clifford gate C . This can be done by interleaving the RBM sequence with C , so repeatedly applying a random Clifford gate followed by C , making a new random choice of Clifford gate for each of m repetitions of this two-gate sequence. The average error rate of C can be estimated

²“Universal” in this context means that any unitary operation can be approximated to arbitrary accuracy.

by [53]:

$$r_C^{est} = \frac{1}{2} \left(1 - \frac{p_C}{p} \right), \quad (2.5)$$

where p_C and p are the depolarizing parameters with and without the interleaved gate C respectively. If we alternate between experiments with the two sequences, then drifts in the SPAM error have an equal effect on the two measured sequence fidelities, and the average SPAM error is constant between the experiments. In this case, r_C^{est} can be expressed in terms of the measured sequence fidelities F_C and F with and without the interleaved gate C respectively:

$$r_C^{est} = \frac{1}{2} \left(1 - \sqrt[m]{\frac{2F_C - 1}{2F - 1}} \right). \quad (2.6)$$

In section 5.2.3, we perform such an experiment with C being simply a delay (i.e. the identity operation) to measure the memory error over each delay.

2.2 BB1 composite gates

One of the main sources of error in quantum gates is slow variations in the pulse area, i.e. under- or over-rotations in the Bloch sphere by some fraction ϵ of the desired angle θ . As shown in equation 2.1, the corresponding gate error is quadratic in ϵ . Expanding this to sixth power in ϵ gives the following equation:

$$1 - F \approx \frac{1}{4}(\epsilon\theta)^2 - \frac{1}{48}(\epsilon\theta)^4 + \frac{1}{1440}(\epsilon\theta)^6, \quad (2.7)$$

where we have again chosen the initial Bloch vector to be perpendicular to the rotation axis, giving the maximum possible gate error. Fortunately, it is possible to almost completely remove this source of error by using composite gate sequences.

The BB1 sequence [54] includes a “do-nothing” pulse before or after the main gate pulse, which, in the absence of any pulse area error, has no effect on the qubit state. This pulse is the following sequence of rotations in the Bloch sphere:

$$U_{BB1} = \pi_{\phi_1} 2\pi_{\phi_2} \pi_{\phi_1}, \quad (2.8)$$

where $\phi_{1,2}$ are angles of the axis of rotations for the constituent pulses in the x - y plane. By careful choice of these angles, one can use the do-nothing pulse to cancel the second and fourth order terms of equation 2.7, meaning that the dominant term is proportional to the sixth power of ϵ . For a target rotation around the axis rotated by ϕ_0 from the x -axis in the x - y plane, the following values should be used:

$$\begin{aligned} \phi_1 &= \phi_0 \pm \cos^{-1} \left(-\frac{\theta}{4\pi} \right), \\ \phi_2 &= \phi_0 \pm 3 \cos^{-1} \left(-\frac{\theta}{4\pi} \right). \end{aligned} \quad (2.9)$$

The resulting gate error for the total BB1 gate sequence is:

$$1 - F \approx \frac{(\pi\epsilon)^6}{4608} \left(\frac{\theta}{\pi} \right)^2 \left(32 + 14 \left(\frac{\theta}{\pi} \right)^2 - \left(\frac{\theta}{\pi} \right)^4 \right). \quad (2.10)$$

Figure 2.1 illustrates this error suppression on the Bloch sphere. For an example pulse area error of $\epsilon = 0.01$ on a $\pi/2$ -gate ($\theta = \pi/2$), which would normally cause an error of 6.2×10^{-5} , the resulting error on the BB1 gate sequence is only 1.8×10^{-12} . Because of its effectiveness in suppressing pulse area errors, we use the BB1 sequence for our experiments on interleaving delay periods between gates (see section 5.2).

2.3 Phase noise

When a qubit is in a state other than $|0\rangle$ or $|1\rangle$, it is susceptible to phase noise, which rotates its state by some angle in the x - y plane of the Bloch sphere. When it is in




Figure 2.1: An illustration on the Bloch sphere of the pulse area error suppression achieved with the BB1 composite gate sequence. The initial state is $|0\rangle$ and the target rotation is $\pi/2$ about the y -axis. There is a fractional undershoot of 0.2 on the raw gate. Each of the three pulses in the “do nothing” sequence have the same fractional undershoot, with the effect of these pulses being to almost completely correct the error in the raw gate.

the x - y plane (i.e. its sensitivity to phase noise is maximal), an uncorrected phase noise distribution $\phi(\tau)$ accumulated over time τ leads to an average probability r of qubit error:

$$r = \frac{\langle \phi^2(\tau) \rangle}{4}. \quad (2.11)$$

When it is instead at one end of one of the three axes of the Bloch sphere, chosen randomly between these six states (as is the case when interleaving delays in ran-

domized benchmarking sequences), one must multiply this expression by a factor of $2/3$. This is the probability of the qubit state being in the x - y plane, with the qubit state otherwise being in $|0\rangle$ or $|1\rangle$, where it is insensitive to phase noise. In this case, the average probability of qubit error is instead:

$$r_{int} = \frac{\langle \phi^2(\tau) \rangle}{6}. \quad (2.12)$$

Phase noise is caused by noise in the qubit frequency ω relative to that of the radiation used to manipulate it. For frequency noise with single-sided spectral density $S(\omega)$, the average phase noise accumulated over some time τ is given by [55]:

$$\langle \phi^2(\tau) \rangle = \frac{\tau^2}{2\pi} \int_0^\infty S(\omega) \left(\frac{\sin(\omega\tau/2)}{\omega\tau/2} \right)^2 d\omega. \quad (2.13)$$

There are several different types of frequency noise, and here we present the ones that are relevant to the work in this thesis. One type that occurs in a wide range of physical systems is white noise, which has a constant spectral noise density:

$$\begin{aligned} S_w(\omega) &= S_w^0, \\ \langle \phi_w^2(\tau) \rangle &= \frac{S_w^0}{2} \tau = 2 \frac{\tau}{T_w}, \end{aligned} \quad (2.14)$$

where T_w is the white noise dephasing time. Qubit coherence is commonly assumed to decay exponentially, so this is our expected noise distribution for the memory error experiments presented in chapter 5.

Another common type of noise is that which is correlated over very long timescales, i.e. the frequency is constant over each measurement period τ but it may change between measurements:

$$\begin{aligned}
S_{corr}(\omega) &= 4\pi\sigma_\omega^2\delta(\omega), \\
\langle\phi_{corr}^2(\tau)\rangle &= \sigma_\omega^2\tau^2 = 2\left(\frac{\tau}{T_{corr}}\right)^2,
\end{aligned}
\tag{2.15}$$

where σ_ω is the standard deviation of the qubit frequency, T_{corr} is the correlated noise dephasing time, and $\delta(\omega)$ is the Dirac delta function.

A third type of noise is $1/f$ noise, which has a spectral noise density inversely proportional to the frequency:

$$S_{1/f}(\omega) = \frac{S_{1/f}^0}{\omega/2\pi}.$$
(2.16)

This distribution diverges for small frequencies, corresponding to indefinite drifts over long periods. In practice, these drifts are corrected for, either with the passive stability of feedback loops or with periodic measurement and adjustment of relevant experimental quantities. This correction can be accounted for by including a cutoff frequency ω_c :

$$S_c(\omega) = \frac{S_c^0}{(\omega + \omega_c)/2\pi}.$$
(2.17)

2.4 $^{43}\text{Ca}^+$ ions as qubits

In order for $^{43}\text{Ca}^+$ ions to be used as qubits in a quantum computer, there are several things that one must do. Initially, it is necessary to ionize neutral ^{43}Ca atoms. After trapping, the ions must be cooled down to prevent them from being lost from the trap, and for high-fidelity entanglement generation, the vibrational mode being

used to couple the internal states of different ions needs to be near the ground state. Next, each ion must be initialized into one of two states to be used as the qubit, and high-fidelity quantum gates must be performed on these qubits, transferring population between the states in a coherent way to generate superposition states and entanglement with other ion qubits. The final qubit states must then be measured, before cooling and preparing the ions for the next iteration of quantum gates (with ions reloaded as and when they are lost). Figure 2.2 shows the levels of $^{43}\text{Ca}^+$ that are relevant for this work, as well as the laser wavelengths required to address each transition.




Figure 2.2: The levels of $^{43}\text{Ca}^+$ that are relevant for the work in this thesis, with the resonance wavelength of each dipole-allowed transition.

2.4.1 Choice of qubit transition

$^{43}\text{Ca}^+$ ions have a non-zero nuclear spin ($I = 7/2$), which means that each level exhibits hyperfine structure. The application of a magnetic field splits this structure through the Zeeman effect into a series of different states within each hyperfine manifold. As discussed throughout this section, the hyperfine splitting makes the

manipulation of the ions' internal state much more challenging, but it also provides a great opportunity with our choice of qubit transition.

In general, there are two types of qubit transitions used with trapped ions: ground-level transitions and optical transitions. In both cases, it is important to ensure that the lifetime of each qubit state is much longer than the time it takes to perform each quantum gate, as otherwise it would be impossible to complete any long computational sequence before losing the stored quantum information. Optical transitions are between two different levels, which generally can be driven with visible light, and are chosen to be dipole-forbidden (e.g. quadrupole transitions), which means that any population in the upper level is stable over relatively long periods. Quantum gates are then performed with narrow-linewidth lasers. Qubit transitions within the ground level, on the other hand, have a very long lifetime because the states have very similar energies. These transitions can also typically be driven with microwave radiation, which, as discussed in chapter 1, has certain advantages over the use of laser radiation. For these reasons, we use a hyperfine transition within the $4S_{1/2}$ ground level for our qubit (for ions without hyperfine structure, different Zeeman states can still be used as a qubit).

It is desirable to choose one such transition whose frequency is first-order independent of magnetic field (an “atomic clock” transition). This means that the transition frequency is largely insensitive to fluctuations in the magnetic field, reducing the quantum decoherence caused by these fluctuations. We would also like to choose an atomic clock transition for which the field-independent frequency point occurs at a non-zero magnetic field. This means that the degeneracy between states within each hyperfine manifold is broken by the different energy shifts induced through the Zeeman effect.³ However, due to the experimental challenge associated with

³For zero-field atomic clock transitions, it is necessary to introduce a small magnetic field to break the degeneracy, which results in a non-zero first-order qubit frequency dependence on magnetic field fluctuations.

generating large magnetic fields that are homogeneous over the ion trapping region, we would like to choose a transition with its field-independent point at a relatively small magnetic field. The first such transition, as shown in figure 2.3, is between the $|F = 4, M = 0\rangle$ and $|F = 3, M = +1\rangle$ states at a magnetic field of 146.0942G, so we choose this as our qubit transition and use coils (see section 3.3) to generate the necessary field. The next transition is between $|F = 4, M = +1\rangle$ and $|F = 3, M = 0\rangle$ at 146.3314G, which has similar properties, but would require an additional microwave π -pulse to be initialized and read out (see sections 2.4.4 and 2.4.6); this would only be a minor inconvenience, but as there is no advantage to its use we do not consider it further.

It is worth noting that there would be two advantages with the use of the next atomic clock transition, which is between the $|F = 4, M = +1\rangle$ and $|F = 3, M = +1\rangle$ states and occurs at 287.7827G. The coupling of microwaves to the transition is a factor of 1.4 stronger than that for $|F = 4, M = 0\rangle \leftrightarrow |F = 3, M = +1\rangle$ (see appendix A). Also, it is π -polarized rather than being σ -polarized. This is an advantage for traps like ours (see section 3.1.1), which cannot generate near-field microwaves with σ^+ or σ^- polarization alone, only an equal combination of the two. This means that when driving σ -polarized transitions, half of the microwave power is wasted, whereas for π -polarized transitions, all power can be used to drive the transition. Combining these two advantages means that, for constant microwave power, the Rabi frequency for the 288G transition would be roughly double that for the 146G transition. This does not matter for the single-qubit work presented in this thesis as near-field microwave single-qubit gate times are very short, but for entanglement generation, this would boost the operation speed and fidelity significantly. Therefore, there is strong motivation to overcome the technical challenge of generating such a large magnetic field and use this transition as a qubit in future work.

At 146G, the states in the ground $4S_{1/2}$ level can be described by the quantum




Figure 2.3: The energy splitting of the different states in the $4S_{1/2}$ ground level as a function of magnetic field. The clock transitions at 146G and 288G discussed in the main text are indicated with vertical lines.

numbers F and M , with $F \in \{3,4\}$ and M being some integer between $-F$ and F . See figure 2.4 for an illustration of this splitting. The states in the $4P_{1/2}$ level can be described with equivalent quantum numbers. For the $3D_{3/2}$ and $4P_{3/2}$ levels, the Zeeman splittings at 146G scramble the hyperfine-manifold structure, so they can only be described by a single quantum number M , the projection of the total angular momentum. For the $3D_{5/2}$ level, the hyperfine splitting is sufficiently small that 146G takes it into the high-field regime, where states can be described by the quantum numbers M_I and M_J . See [56, 57, 33, 58] for details of how the hyperfine

state energies vary with magnetic field.




Figure 2.4: The hyperfine structure of the $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at a magnetic field of 146 G. The labels A , B , C , D , E and Q correspond to the transitions that are relevant for the work in this thesis.

2.4.2 Ionization

A common method of trapping ions is to produce an atomic beam through the trapping region, and to ionize the atoms in this location with laser light, a process which is called photo-ionization. This allows the trapping field to be applied continuously, as it does not interfere with the neutral atomic beam, only interacting with atoms once they have been ionized.

When the atomic beam only contains one isotope of the element to be ionized, there is no strict requirement on the photo-ionization laser frequency; it must simply be above a certain threshold to remove an electron from the atom. However, ^{43}Ca has quite low abundance – even in our isotopically enriched source it is only 12% of the calcium (the other 88% is ^{40}Ca). Therefore, to ensure that we load the desired

isotope, it is necessary to ionize in an isotopically selective way.

We achieve this isotope selectivity using the two-stage approach presented in [59, 60] and illustrated in figure 2.5. First, the $4s^2\ ^1S_0$ to $4s4p\ ^1P_1$ transition (423nm) is stimulated with laser light at the resonance frequency for ^{43}Ca . The detuning from the ^{40}Ca transition frequency is sufficiently large that this isotope is unaffected. We also apply 375nm light to achieve ionization once the outermost electron is in an excited state (the ionization threshold wavelength is 390nm). The 423nm laser light is frequency-stabilized to maintain isotope selectivity (see section 3.4.1 for details), but there is no strict frequency requirement on the 375nm laser light so this is not stabilized.




Figure 2.5: The levels of ^{43}Ca and laser wavelengths that are relevant for our photo-ionization process. See section 3.4.1 for the details of the lasers used to address these transitions.

2.4.3 Laser cooling

The technique of using laser light to cool atoms and ions with the Doppler effect is used in a wide range of scientific fields. The basic idea is to drive an electric-dipole

transition with laser light, where the momentum of each absorbed photon is transferred to the ion, giving it a small push in the direction of the laser beam. Detuning the laser frequency slightly to the red of resonance means that ions preferentially absorb photons if they are moving towards the beam source. After absorbing many photons in this way, the vibration of the ion in the trap is reduced, where millikelvin temperatures are achievable with this scheme. For ion traps, if the beam has a component along all three vibrational modes, then it is possible to cool the ion in all of these modes with just this one beam. The limitation on this cooling method is typically with the “recoil heating” caused by the emission of the absorbed light, which gives the ion a momentum kick in the opposite direction to the photon emission.

Doppler cooling is usually applied to atoms with only a small number of relevant transitions: a main cooling transition, and one or more transitions to recapture population that is lost from the cooling transition due to decay to other levels. However, our application of a 146G magnetic field splits the levels through the Zeeman effect by more than the transition linewidth, so the different states are spectroscopically resolvable. The three levels of $4S_{1/2}$, $4P_{1/2}$ and $3D_{3/2}$, all of which are important for cooling on the $4S_{1/2}$ to $4P_{1/2}$ transition (population may decay from $4P_{1/2}$ to $3D_{3/2}$), are split into 64 different states.

It may seem that this splitting means that ~ 100 different laser beams would be required to cool $^{43}\text{Ca}^+$ on the $4S_{1/2}$ to $4P_{1/2}$ transition and repump from $3D_{3/2}$ (to prevent the ion from getting trapped in this metastable level). However, through careful choice of beam frequencies, polarizations and intensities, it is possible to achieve Doppler cooling with only three beams, as illustrated in figure 2.6.

The $4S_{1/2}$ to $4P_{1/2}$ transition (with a wavelength of 397nm) is addressed with two frequencies, one close to the $4S_{1/2} |F = 3, M = +3\rangle$ to $4P_{1/2} |F = 4, M = +4\rangle$




Figure 2.6: The states of $^{43}\text{Ca}^+$ and the laser wavelengths that are relevant for our Doppler cooling process.

transition frequency and the other close to the $4\text{S}_{1/2} |F=4, M=+4\rangle$ to $4\text{P}_{1/2} |F=4, M=+4\rangle$ transition frequency. The polarization is predominantly σ^+ , with a small component of π . The reasoning behind these choices is that we want to pump population towards one end of the ground state manifold (chosen to be the end with positive M). The π component ensures that it is not trapped in $4\text{S}_{1/2} |F=4, M=+4\rangle$ (there is no $M=+5$ state in the $4\text{P}_{1/2}$ manifold to be reached with σ^+ radiation), with the $4\text{S}_{1/2} |F=4, M=+4\rangle$ to $4\text{P}_{1/2} |F=4, M=+4\rangle$ transition being the main cycling transition for Doppler cooling.

To repump from the metastable $3\text{D}_{3/2}$ level, we use a single 866nm beam. This is possible because the frequency spread within $3\text{D}_{3/2}$ is relatively small, enabling us to address all states by saturation broadening the transitions to $4\text{P}_{1/2}$ with a high laser intensity.

To further cool the ions, we make use of a Λ -type two-photon dark resonance [61] between $4\text{S}_{1/2} |F=4, M=+4\rangle$, $4\text{P}_{1/2} |F=4, M=+4\rangle$ and $3\text{D}_{3/2} |M=+5\rangle$ as il-

illustrated in figure 2.7. When the radiation coupling the $4P_{1/2}$ state to the $4S_{1/2}$ and $3D_{3/2}$ states is tuned near resonance, the absorption of light is suppressed, with population trapped in the lower states. This dark resonance provides the two conditions that are required for traditional Doppler cooling: a low absolute fluorescence when the Doppler shift is small, and a large increase in fluorescence when the Doppler shift increases. In addition, with careful choice of laser beam polarizations, this Λ -system can be prevented from coupled to any other states, making it stable. See figure 2.8 for the change of $4P_{1/2}$ population with the 397nm beam detuning. The $4P_{1/2}$ population, and therefore the fluorescence, increases sharply with the frequency shift from the dark resonance.




Figure 2.7: The states of $^{43}\text{Ca}^+$ and the laser wavelengths and polarizations for the dark resonance that we use in our dark resonance cooling process.

Because our lasers are not monochromatic, there is actually a slight leakage of population out of the dark resonance, some of which ends up in the $4S_{1/2} |F=3\rangle$ hyperfine manifold. To prevent population from getting trapped here, we use a second 397nm laser beam to pump population back to $4P_{1/2} |F=4\rangle$. For experimental simplicity, these two beams are the same ones used for our Doppler cooling process (see above), with the powers reduced when switching from Doppler cooling to dark resonance




Figure 2.8: The simulated $4P_{1/2}$ population as a function of the 397nm laser beam detuning. The blue curve is for our laser linewidths, and the red for monochromatic lasers. The black line indicates the location of the main dark resonance. Figure adapted from [62].

cooling.

A key parameter for determining the minimum temperature reached by dark resonance cooling in $^{43}\text{Ca}^+$ is the population in $4P_{1/2}$. This population determines the fluorescence, and therefore the amount of recoil heating that the ion experiences. If the population at the dark resonance is reduced, the effectiveness of the heating processes will fall off more rapidly than that of the cooling processes, resulting in an a lower final temperature. Therefore, to minimize the temperature we want the population to be as small as possible [58].

However, there is another consideration in our choice of $4P_{1/2}$ population. At the start of the cooling process, the range of Doppler frequency shifts induced by the ion's vibration is much larger than the width of the main dark resonance. The initial cooling is only made possible by the existence of other dark resonances in the system. While most of these resonances are unstable, and so are unsuitable for reaching very low temperatures, when the ion is relatively hot they can still collectively provide a substantial amount of cooling power.

As the $4P_{1/2}$ population is lowered, the heating processes of these other dark reso-

nances become stronger at a greater rate than the cooling processes; the resonances cannot cool the ion as effectively. If the population is lowered too far, these resonances will lose all of their ability to cool, and will instead heat the ion. Therefore, should we decrease the population below a certain point, the initial cooling power would become insufficient to lower the ion temperature to the point at which the main dark resonance can take the ion to the minimum temperature. This trade-off is the limiting factor in the temperature that can be reached in the cooling process.

Keeping this consideration in mind, we optimized our dark resonance cooling scheme with a theoretical and experimental investigation of the temperature achieved with different intensities, frequencies and polarizations of the laser beams. Our theoretical calculations used the optical Bloch equations, which govern the evolution of the system under different radiation fields [58]. Guided by these calculations, we scanned the different beam parameters to minimize the ion temperature, reaching a value of 0.3mK, which was measured using sideband thermometry [62, 50].

For high-fidelity entanglement generation, which uses one of the trap vibrational modes to couple the internal states of different ions, it is necessary to cool this vibrational mode into the ground state. This can be achieved with resolved sideband cooling, which drives a low-linewidth transition (where the transition sidebands are resolved) in such a way that the ions lose quanta of vibrational motion. In our work on two-qubit gates [36, 50], we used the Raman transition from $4S_{1/2} |F = 4, M = +4, n\rangle$ to $4S_{1/2} |F = 4, M = +3, n - 1\rangle$ via $4P_{1/2}$, where n is the initial vibrational quantum number on the vibrational mode used for entanglement generation. Additional beams at 397nm and 866nm pumped population from $4S_{1/2} |F = 4, M = +3, n - 1\rangle$ to $4S_{1/2} |F = 4, M = +4, n - 1\rangle$, closing the cooling cycle. However, the temperature reached by dark resonance cooling is sufficiently low for the single-qubit work presented in this thesis, so we do not apply sideband cooling here.

2.4.4 Qubit initialisation

Once the ion has been cooled, the next step is to prepare its internal state in one of the qubit states. We start by driving population into $4S_{1/2} |F = 4, M = +4\rangle$. This process begins with a pulse of σ^+ -polarized 397nm light, with two frequency components to address states in both the $|F = 3\rangle$ and $|F = 4\rangle$ hyperfine manifolds of the ground level. However, imperfections in the polarization of this beam result in there being a small population in states with $M = 2$ and $M = 3$. Most of this population can be recovered with another pulse of σ^+ -polarized 397nm light, but with only the frequency component addressing the $|F = 3\rangle$ manifold, driving all population into $|F = 4\rangle$, preferentially to states with increased M . Next, we use one of the trap microwave electrodes (see section 3.1.1 for more information about the trap) to apply microwave π -pulses resonant with the $|F = 3, M = +3\rangle \leftrightarrow |F = 4, M = +3\rangle$ and $|F = 3, M = +2\rangle \leftrightarrow |F = 4, M = +2\rangle$ transitions, which we name D and E respectively (see figure 2.9), followed by another pulse of single-frequency 397nm σ^+ light. This sequence of microwave D and E π -pulses and a 397nm σ^+ laser pulse is repeated twice, at which point most of the lost population is reclaimed.

Our qubit transition Q is between the states $|F = 4, M = 0\rangle$ and $|F = 3, M = +1\rangle$. To transfer population into the qubit manifold from $|F = 4, M = +4\rangle$, we apply three microwave π -pulses A , B and C to transfer population to $|F = 3, M = +3\rangle$, $|F = 4, M = +2\rangle$ and then $|F = 3, M = +1\rangle$ respectively (see figure 2.10).

2.4.5 Quantum gates

To drive transitions in the ground level of $^{43}\text{Ca}^+$, we use magnetic fields oscillating at microwave frequencies. These fields are generated by passing currents through electrodes in the ion trap. The field $\mathbf{B}(t)$ interacts with the magnetic dipole moment $\boldsymbol{\mu}$ of the ion. The Rabi frequency Ω of these “magnetic dipole transitions” between $|F = 4, M = M_m\rangle$ and $|F = 3, M = M_n\rangle$ is defined by the following equation:




Figure 2.9: Part of the hyperfine structure of the $4S_{1/2}$ level (see figure 2.4 for the full structure), showing the D and E transitions, which are driven during the procedure to pump all population into $|F = 4, M = +4\rangle$.

$$\Omega = -\frac{1}{\hbar} \langle F = 3, M = M_n | \boldsymbol{\mu} \cdot \mathbf{B}(t) | F = 4, M = M_m \rangle \quad (2.18)$$

$$= \frac{1}{\hbar} (-1)^{q+1} \langle F = 3, M = M_n | \mu_q | F = 4, M = M_m \rangle B_{-q}(t), \quad (2.19)$$

which is written in the spherical basis, with $q = M_n - M_m$ and only $q \in \{0, \pm 1\}$ corresponding to magnetic dipole allowed transitions. The transition frequencies, dependence on magnetic field, and coupling elements are derived for fields of 146G and 288G in [57, 49, 50] and given in tables A.1 and A.2.

For generating entanglement, large field gradients are needed to induce coupling between ions via their shared motional modes in the trap. This field gradient is traditionally produced with far-field, visible laser radiation, which has a sufficiently short wavelength (and, therefore, produces a sufficiently large field gradient) to generate entanglement over reasonable timescales.




Figure 2.10: Part of the hyperfine structure of the $4S_{1/2}$ level (see figure 2.4 for the full structure), showing the A , B , C and Q transitions, which are driven during the procedures to initialize and readout the qubit state.

For microwaves, the wavelength is too long to use the field gradient of far-field radiation to generate entanglement in the same way as for visible light. By instead using near-field microwaves, we obtain a sufficiently high field gradient from the fact that the field amplitude decays very quickly with distance from the current-carrying electrodes. We now formalize this argument by deriving expressions for the coupling strengths between the ions' internal and motional states, showing how this coupling depends on the magnetic field and field gradient.

We consider a Paul trap with much weaker confinement in the z direction than in the x or y directions. When sufficiently cold, ions confined in such a trap align along the z direction. They will have some residual oscillatory motion, which, providing the anharmonicity of the trapping potential and the mutual Coulomb interactions are small, can be described as a set of independent harmonic oscillators [63]. Therefore, the deviation of each of N ions from their equilibrium positions can be described as that due to a set of normal modes along each of the three principal trap axes $\hat{q}_l$.

The position of the n th ion in the string is:

$$\mathbf{r}_n = \mathbf{r}_n^0 + \sum_{l=1}^3 \sum_{j=1}^N \hat{\mathbf{q}}_l b_n^{l,j} q_{l,j}^0 (a_{l,j} + a_{l,j}^\dagger), \quad (2.20)$$

where $\mathbf{r}_n^0$ is the equilibrium position of the ion, $b_n^{l,j}$ is each normal mode's eigenvector, $q_{l,j}^0 = \sqrt{\hbar/2m\omega_{l,j}}$ is the spatial extent of each ground state wavefunction, $\omega_{l,j}$ is each frequency, and $a_{l,j}$ and $a_{l,j}^\dagger$ is each pair of ladder operators.

The next step in the derivation is to insert this equation into the expression for the magnetic field. Given that the deviation of each ion from its equilibrium position is small, we can describe the magnetic field at its location with the first two terms of a Taylor series:

$$\mathbf{B}(\mathbf{r}_n) \approx \mathbf{B}(\mathbf{r}_n^0) + \frac{d\mathbf{B}}{d\mathbf{q}}(\mathbf{r}_n^0)(\mathbf{r}_n - \mathbf{r}_n^0), \quad (2.21)$$

where $d\mathbf{B}/d\mathbf{q}$ is the Jacobian of the field gradient:

$$\frac{d\mathbf{B}}{d\mathbf{q}} = \begin{pmatrix} \frac{\partial B_{+1}}{\partial q_1} & \frac{\partial B_{+1}}{\partial q_2} & \frac{\partial B_{+1}}{\partial q_3} \\ \frac{\partial B_{-1}}{\partial q_1} & \frac{\partial B_{-1}}{\partial q_2} & \frac{\partial B_{-1}}{\partial q_3} \\ \frac{\partial B_z}{\partial q_1} & \frac{\partial B_z}{\partial q_2} & \frac{\partial B_z}{\partial q_3} \end{pmatrix}, \quad (2.22)$$

and $\mathbf{B}$ is written in the spherical basis with unit vectors $\{\hat{\mathbf{e}}_{+1}, \hat{\mathbf{e}}_{-1}, \hat{\mathbf{e}}_0 \equiv \hat{\mathbf{e}}_z\}$ and the derivative is taken in the directions of the principal trap axes.

Writing this out as a summation and substituting in equation 2.20:

$$\mathbf{B}(\mathbf{r}_n) = \sum_{Q=0,\pm 1} (-1)^Q B_Q \hat{\mathbf{e}}_{-Q} \quad (2.23)$$

$$= \sum_{Q=0,\pm 1} (-1)^Q \left(B_Q^0 + \sum_{l=1}^3 \sum_{j=1}^N \frac{\partial B_Q^0}{\partial q_l} b_n^{l,j} q_{l,j}^0 (a_{l,j} + a_{l,j}^\dagger) \right) \hat{\mathbf{e}}_{-Q}, \quad (2.24)$$

where we have introduced the abbreviations $\partial B_Q^0 / \partial q_l \equiv \partial B_Q(\mathbf{r}_n^0) / \partial q_l$ and $B_Q^0 \equiv B_Q(\mathbf{r}_n^0)$.

Next, we work out the interaction Hamiltonian $H_{I,n}(t)$ resulting from the magnetic field interacting with each ion's magnetic dipole moment:

$$H_{I,n}(t) = -\boldsymbol{\mu} \cdot \mathbf{B}(\mathbf{r}_n) \cos(\omega t - \phi), \quad (2.25)$$

where we have inserted the extra term $\cos(\omega t - \phi)$ for a field oscillating at frequency ω with phase ϕ .

Writing this out in the qubit subspace $\{|0\rangle, |1\rangle\}$:

$$H_{I,n}(t) \equiv \begin{pmatrix} H_{00}(t) & H_{01}(t) \\ H_{10}(t) & H_{11}(t) \end{pmatrix} = H_{01}(t) \sigma_{x,n} + \frac{1}{2} (H_{11}(t) - H_{00}(t)) \sigma_{z,n}, \quad (2.26)$$

where we have dropped a constant additive energy offset $(1/2)(H_{00}(t) + H_{11}(t))\mathbb{I}$ and used the fact that $H_{I,n}(t)$ is Hermitian and real to set $H_{10}(t) = H_{01}(t)$.

The interaction Hamiltonian for the n th ion can now be written as:

$$H_{I,n}(t) = -(\bar{\mu}_Q B_{-Q} \sigma_{x,n} + \bar{\mu}_z B_z \sigma_{z,n}) \cos(\omega t - \phi), \quad (2.27)$$

where $Q \equiv M_1 - M_0$ is the qubit transition polarization and we use the definitions:

$$\bar{\mu}_Q \equiv (-1)^Q \langle 1 | \mu_Q | 0 \rangle, \quad (2.28)$$

$$\bar{\mu}_z \equiv \frac{1}{2} (\langle 1 | \mu_z | 1 \rangle - \langle 0 | \mu_z | 0 \rangle). \quad (2.29)$$

Substituting in equation 2.24:

$$\begin{aligned} H_{I,n}(t) = & -\cos(\omega t - \phi) \left(\bar{\mu}_Q B_{-Q}^0 \sigma_{x,n} + \bar{\mu}_z B_z^0 \sigma_{z,n} \right. \\ & \left. + \sum_{l=1}^3 \sum_{j=1}^N \left(\bar{\mu}_Q \frac{\partial B_{-Q}^0}{\partial q_l} \sigma_{x,n} + \bar{\mu}_z \frac{\partial B_z^0}{\partial q_l} \sigma_{z,n} \right) b_n^{l,j} q_{l,j}^0 (a_{l,j} + a_{l,j}^\dagger) \right). \end{aligned} \quad (2.30)$$

We now move into the interaction picture with respect to the ions' internal and motional energy:

$$H_0 = \frac{1}{2} \hbar \omega_0 \sigma_z + \sum_{l=1}^3 \hbar \omega_{l,n} \left(\frac{1}{2} + a_{l,n}^\dagger a_{l,n} \right), \quad (2.31)$$

where ω_0 is the qubit frequency. We also make a rotating wave approximation, neglecting all terms rotating at frequency $2\omega_0$ to reach the equation:

$$\begin{aligned} \tilde{H}_{I,n}(t) = & \frac{1}{2} \hbar \Omega_Q e^{-i(\delta t - \phi)} \sigma_{+,n} + \frac{1}{2} \hbar \Omega_z e^{-i(\omega t - \phi)} \sigma_{z,n} \\ & + \sum_{l=1}^3 \sum_{j=1}^N \left(\frac{1}{2} \hbar \Omega_{Q,l,j,n} e^{-i(\delta t - \phi)} \sigma_{+,n} + \frac{1}{2} \hbar \Omega_{z,l,j,n} e^{-i(\omega t - \phi)} \sigma_{z,n} \right) \\ & \times \left(a_{l,j} e^{-i\omega_{l,j} t} + a_{l,j}^\dagger e^{i\omega_{l,j} t} \right), \end{aligned} \quad (2.32)$$

where we have defined the magnetic field detuning $\delta \equiv \omega - \omega_0$ and the following Rabi frequencies:

$$\Omega_Q \equiv -\frac{1}{\hbar} \bar{\mu}_Q B_{-Q}^0, \quad (2.33)$$

$$\Omega_z \equiv -\frac{1}{\hbar} \bar{\mu}_z B_z^0, \quad (2.34)$$

$$\Omega_{Q,l,j,n} \equiv -\frac{1}{\hbar} \bar{\mu}_Q \frac{\partial B_{-Q}^0}{\partial q_l} b_n^{l,j} q_{l,j}^0, \quad (2.35)$$

$$\Omega_{z,l,j,n} \equiv -\frac{1}{\hbar} \bar{\mu}_z \frac{\partial B_z^0}{\partial q_l} b_n^{l,j} q_{l,j}^0. \quad (2.36)$$

The first and second terms act only on the internal state of each ion, driving single-qubit gates. The third and fourth term couple the internal and motional states, allowing for the creation of entanglement between different ions in the chain. The strength of the single-qubit gates is proportional to the magnetic field, whereas the strength of the entanglement generation is proportional to the field gradient.

2.4.6 Qubit readout

We read out our qubit state using state-dependent fluorescence detection, otherwise known as electron shelving [64]. This involves transferring the population in one qubit state into a metastable level, and then driving the population in the other state on some transition that does not couple to the level containing the population of the first qubit state. The two states are then distinguished from one another by observing whether the ion fluoresces at the frequency of the driven transition.

We start this process by applying the microwave π -pulses C , B , A , and then Q , which map $|F=3, M=+1\rangle$ to $|F=4, M=+4\rangle$ and then $|F=4, M=0\rangle$ to $|F=3, M=+1\rangle$ (see figure 2.10). At this point, we use the technique originally described in [65] to optically pump the population in $|F=4, M=+4\rangle$ to the $3D_{5/2}$ level (our “shelf”), while not interacting with the population in $|F=3, M=+1\rangle$. This process starts with the application of a pulse of σ^+ -polarized 393nm laser light resonant with the $4S_{1/2} |F=4, M=+4\rangle$ to $4P_{3/2} |M=+5\rangle$ transition. The large

detuning from $4S_{1/2} |F = 3, M = +1\rangle$ means that the population in this state is unaffected.

Once in $4P_{3/2} |M = +5\rangle$, the population can decay one of several ways. It could decay to the metastable $3D_{5/2}$ level, which is our chosen shelf level and so is where we want it to end up eventually. It could also decay to $3D_{3/2} |M = +4, +5\rangle^4$; to reclaim this population, we must apply pulses of 850nm laser light. There are two possible states with $M = +4$ and one with $M = +5$ that the population could decay to. We apply all three different frequencies for the transitions from these three states back to $4P_{3/2} |M = +5\rangle$. We start with a σ^+ -polarized pulse, which reclaims the population in $3D_{3/2} |M = +4\rangle$, and then apply a π -polarized pulse for the population in $3D_{3/2} |M = +5\rangle$. It is important to apply the polarizations in this order, as π -polarized light could transfer population from $3D_{3/2} |M = +4\rangle$ into $4P_{3/2} |M = +4\rangle$, which could then decay to the $4S_{1/2} |F = 3\rangle$ hyperfine manifold, making this population indistinguishable from that in the other qubit state.

The dominant decay path from $4P_{3/2} |M = +5\rangle$ is back to $4S_{1/2} |F = 4, M = +4\rangle$. To repeatedly clear population from this state, as well as the relevant states in $3D_{3/2}$, we apply the above pulse sequence of 393nm and 850nm laser light eight times, each time transferring some population into the $3D_{5/2}$ shelf level, with a final 393nm pulse at the end of this sequence. The number of repetitions was chosen as the minimum necessary to ensure that all population is transferred [49]. See figure 2.11 for an illustration of this process.

Once the population that was originally in the $4S_{1/2} |F = 3, M = +1\rangle$ qubit state has been shelved in the $3D_{5/2}$ level, we implement the same high-fluorescence Doppler cooling from section 2.4.3. This process does not interact with $3D_{5/2}$, so if the ion is in this state then no fluorescence will be observed. However, if the ion is still

⁴As explained in section 2.4, F is not a good quantum number in this level.




Figure 2.11: The states of $^{43}\text{Ca}^+$ and the laser wavelengths that are relevant for our qubit state shelving process.

in $4S_{1/2}$, it will emit a large number of 397nm photons during the Doppler cooling. These photons are emitted in random directions, with 0.23% detected with an optical system and a photomultiplier tube (see section 3.4.4). If such photons are detected, then the ion should have originally been in the $4S_{1/2} |F=4, M=0\rangle$ qubit state; if not, the $4S_{1/2} |F=3, M=+1\rangle$ state.

In practice, when in the ‘dark’ state, there can still be some non-zero number of photons detected. This can be due to the detector ‘measuring’ a photon even when there was not one collected (the detector “dark count”), due to collection of light from sources other than the ion, or due to the ion decaying from the metastable $3D_{5/2}$ shelf level into $4S_{1/2}$. In our previous work [50], we counted the number of photons measured over a 3ms period, chose some threshold value to distinguish between the ‘bright’ and ‘dark’ states, and deduced the qubit state from whether the number of counts was above or below this threshold.

We now use a time-resolved maximum likelihood method for our readout. By simply setting a threshold on the number of photons measured over a single period, the

time of each photon detection is not used in the state estimation. If the ion starts in the ‘dark’ state and decays some time later to the ‘bright’ state, then the rate of photon detection would start off being low, and then undergo a step change at the decay time to a higher number. As first suggested in [66], this temporal information about the photon detections can be used to increase the readout fidelity.

We use the method presented in [65, 67], which involves recording the numbers of photons detected in each of some number (we use 16) of detection periods (we use $t_D = 2^{17}\text{ns} = 131\mu\text{s}$ for each period). These values were chosen as a rough compromise between minimizing the total readout time and the number of detection periods, while also minimizing the readout error.

This process produces a list of the numbers of photons n_i measured over each of the N detection periods t_d . The next step is to calculate the probabilities p_B and p_D of such a list occurring were the ion to have started in the bright and dark states respectively. For p_B , this calculation is fairly straightforward – there is a negligible probability of the population leaving the collection of states involved in the readout fluorescence process, which means that the expression is simply:

$$p_B = \prod_{i=1}^N B(n_i), \quad (2.37)$$

where $B(n_i)$ is the probability of measuring n_i photons over any one of the detection periods. $B(n_i)$ is approximately Poissonian, so it can be calculated for any n_i from just the mean of the distribution, which we measure on a regular basis, and is generally ~ 3.6 for our system and choice of t_d .

The calculation of p_D is more complicated. The ion may remain in the dark state throughout the readout process, in which case the distribution is similar to equation 2.37, with a Poissonian distribution $D(n_i)$ of mean typically ~ 1.2 . There is

also the possibility of the ion decaying during one of the detection periods, in which case the photon emission probability distribution would be governed by $D(n_i)$ up to the decay time and then by $B(n_i)$ until the end of the readout process.

These considerations lead to the equation (see [67] for a derivation):

$$p_D = \left(1 - \frac{Nt_d}{\tau}\right)M_N + \left(\frac{t_d}{\tau}\right)S_N, \quad (2.38)$$

where τ is the lifetime of the $3D_{5/2}$ shelf level and we have the following recursion relations:

$$M_0 = 1, \quad M_k = M_{k-1}D(n_k), \quad (2.39)$$

$$S_0 = 0, \quad S_k = (S_{k-1} + M_{k-1})B(n_k). \quad (2.40)$$

Once the qubit state has been measured, any population remaining in the $3D_{5/2}$ shelf level is deshelled with a pulse of 854nm laser light.

Chapter 3

Apparatus

In this chapter, we give a description of the experimental setup. Much of the apparatus was developed in previous work and is as described in [57, 49, 50]. Two parts that have undergone significant changes are the qubit microwave drive system and the overall experimental control system. The control of our static magnetic field has also developed substantially, and is described separately in chapter 4. Figure 3.1 gives an overview of the apparatus in block diagram form, showing how each of the main components interact with one another.

3.1 Ion trap and electrode connections

3.1.1 Ion trap

The ion trap used for the experiments in this thesis is a “5-wire” [68] surface trap, depicted in figure 3.2. It consists of gold electrodes lithographically fabricated on a sapphire substrate (see [57] for details of our fabrication process), and it operates at room temperature, with ions trapped $75\mu\text{m}$ above the surface.

The trapping field is produced with a pair of RF electrodes (carrying signals at 38.3MHz with an amplitude of 72V) and six DC electrodes (with voltages of up




Figure 3.1: A block diagram overview of the apparatus.

to 6.1V). These signals result in secular frequencies of $\sim 3\text{MHz}$ in the radial and $\sim 500\text{kHz}$ in the axial directions. The trap depth is $\sim 59\text{meV}$.

The trap was designed to perform two-qubit gates with near-field microwaves. For this application, it is desirable to be able to null the microwave field at the ions' locations (a non-zero field generally results in single-qubit rotations) while having a non-zero field gradient (which can be used to couple to the ions' motion, for use in entangling gates), so three microwave electrodes were included, the minimum number required to achieve this field distribution [30]. The size of the microwave magnetic field gradient that can be produced is critical, as one of the main limitations of the two-qubit gate fidelity is the achievable gate speed [50]. This motivates us to trap the ions as close to the electrodes as possible.

However, if the ions get too close to the trap surface then “anomalous heating” of the ions becomes the dominating source of gate infidelity [28]. Anomalous heating




Figure 3.2: The ion trap used in our experiments. (a) Photograph of the trap mounted on the aluminium holder. (b) CAD model of the trap, with labels indicating the integrated microwave elements. (c) Schematic of the central trapping region, showing the different signals that can be applied to each electrode. The near-field microwaves produced by the central electrode are almost completely π -polarized, which, for our σ -polarized qubit, is useful for generating entanglement (via the field gradient) without driving the qubit transition. We do not apply microwaves to this electrode for the single-qubit work presented in this thesis. $\mathbf{B}_0$ indicates the direction of our 146G static magnetic field. Figure adapted from [69].

is the heating of ions by electric field noise in the trap electrodes, where the source of this noise is unknown and is an active area of research in the field of ion trapping. The magnitude of this effect increases with decreasing ion height d , scaling as $\sim 1/d^4$

[70, 71]. The resulting increase in the ions' vibrational level interferes with two-qubit gates, which use the vibrational state as a bus for entanglement generation. As a compromise between the two competing factors of gate speed and anomalous heating, an ion height of $75\mu\text{m}$ was chosen.

The need to maximize the microwave field gradient led to the decision to use the RF electrodes to carry microwave currents as well as RF voltages. Two alternatives were to have multiple microwave electrodes between the RF electrodes (as in [30]), or to have one or more microwave electrodes beyond the RF electrodes (so further away from the trapping region, which is between the RF electrodes). These possibilities were investigated with the Ansys HFSS software, and found to result in smaller field gradients at the ions' locations [57].

We also couple DC voltages onto the microwave electrodes in order to produce a static quadrupole field in the radial plane. This rotates the radial trap axes by $\sim 40^\circ$, allowing both radial motional modes to be cooled with laser beams parallel to the trap surface [72]. If our 397nm laser needed to be at an angle to the surface, then, due to its high frequency, it would cause charging of the trap substrate through the photoelectric effect. In addition, the scattered blue photons would reduce our readout fidelity (which relies on the absence of blue light when the ion is in the 'dark' state). The DC voltages on the microwave electrodes also break the degeneracy of the two radial modes – this was important for our entanglement generation, which used one of these modes. The coupling of different frequency signals onto the three central electrodes requires the signals to be combined with a diplexer before being sent into the vacuum chamber (see section 3.1.4).

Another feature of the trap is integrated microwave circuit elements. These elements were included to maximize the microwave current in the electrodes for a given input power of 3.2GHz microwave signals. Quarter-wavelength coupling elements

impedance-match the 50Ω input transmission lines with half-wavelength resonators, the centres of which are under the trapping region. The resonators have $Q \sim 5$, with open circuits at one end to maximize the microwave currents at the centres of the resonators and thus maximize the microwave field gradient at the ions' locations.

The gold trap electrodes were electroplated onto the sapphire substrate following the evaporative deposition of thin layers of titanium (to ensure good adhesion to the substrate) and silver (to seed the gold electroplating process). Sapphire was chosen for this application due to its high thermal conductivity, which was important for the dissipation of the large microwave powers necessary for applying two-qubit gates. The substrate was glued onto an aluminium holder with epoxy¹. The three central electrodes were glued with conductive epoxy² to launch pins connected to modified 2.92mm SMA connectors. The six DC control electrodes were connected to the trap holder with copper wire bonds.

3.1.2 DC supply

The DC trap voltages are generated by a multi-channel DAC³ on an evaluation board⁴. For the six DC electrodes to either side of the trapping region, the DAC outputs are connected via feedthroughs in the vacuum can base flange. For the three axial electrodes, the DACs are connected via the diplexer (see section 3.1.4). All channels are filtered with low-pass filters before connection to the trap.

3.1.3 Trapping RF

The 38.3MHz trapping RF voltage is generated by a frequency synthesizer⁵ and amplitude-stabilized with a custom-made circuit which uses a DC voltage source

¹Epoxy Technology, Inc., EPO-TEK 353ND

²Epoxy Technology, Inc., EPO-TEK H20E

³Analog Devices, Inc., AD5372

⁴Analog Devices, Inc., EVAL-ADUM4160

⁵HP/Agilent/Keysight, 8656B

to set the amplitude of a square-wave signal [33]. A low-pass filter then removes the higher-frequency components of this signal. The stability of the RF power is important for experiments described in this thesis, as the RF currents in the trap result in a significant AC Zeeman shift ($\sim 5\text{Hz}$) on our qubit transition, so fluctuations in RF power can cause qubit decoherence. The output of this circuit is passed through an amplifier⁶, stepped up in voltage with a transformer, and then passed into our diplexer (see section 3.1.4), where the signal is combined with microwave and DC signals.

3.1.4 Diplexer

To apply DC, RF and microwave signals to the same electrodes, we combine the signals with a diplexer. The RF signal is first stepped up by a factor of 9.6 in voltage with a transformer⁷, which has a Q factor of 65 and a resonance frequency of 38.7MHz. The three frequencies are combined with a custom-made filter network (the “diplexer”) and then connected to the trap chip via SMA feedthroughs. See figure 3.3 for a block diagram of the filter network.




Figure 3.3: A block diagram of the diplexer filter network for each trap electrode input. For the central electrode, which does not carry RF signals, the RF input is grounded.

⁶Spectrum Microwave/API Technologies Corp, TM3310

⁷Micrometals, T80-10

3.2 Vacuum system

Figure 3.4 gives a schematic view of our vacuum system. The main chamber is a “spherical octagon”⁸, which has eight 2.75” flanges on the sides as well as 6” flanges on the top and bottom, with the ion trap at the centre of the chamber. Six of the side flanges have viewports⁹ to allow for laser access to the trap. One of the side entrances has SMA feedthroughs for the signals to be applied to the axial trap electrodes, and the other entrance is connected to the vacuum pumps, valves and pressure gauge. The bottom flange has DC feedthroughs for the six DC electrodes, feedthroughs for the calcium oven, and tapped holes for mounting the aluminium trap holder. The top flange has another viewport¹⁰, this one used for ion fluorescence collection for state readout, as well as for one of the infrared laser beams.

We use an ion pump¹¹ and a getter¹² to maintain a vacuum pressure of 10^{-12} - 10^{-11} mbar. Also connected to this entrance of the vacuum chamber are two valves connected in series, which were used for pumping the system down from atmospheric pressure with a turbo pump. The valves were closed and the pump was disconnected after this initial pumpdown. Two valves were used instead of one in case one of them leaked, which would otherwise have necessitated a repeat of our baking process (taking several weeks) rather than simply closing the second valve. We monitor the vacuum pressure with an ionization gauge¹³.

Our calcium source is a resistively heated oven, consisting of a stainless steel tube containing granules of calcium, where the isotope mixture is 12% ^{43}Ca and 88% ^{40}Ca . This oven is typically heated with a 5.5A current, with each loading of a $^{43}\text{Ca}^+$ ion

⁸Kimball Physics, MCF600-SphOct-F2C8

⁹Torr Scientific, VPZ38QBBar-LN, fused silica, 3mm thick, broadband antireflection coating on both surfaces, specified reflectance per surface of 0.5% at 397nm and 9.5% at 866nm

¹⁰Torr Scientific, VPZ100-ITO/AR-LN, Kodial glass, antireflection coating on air side and indium tin oxide layer on vacuum side. Net transmission for 397nm of 76% at normal incidence

¹¹Varian/Agilent Technologies, Diode VacIon Plus 20

¹²SAES Getters, GP100 MK5

¹³Varian/Agilent Technologies, UHV-24P




Figure 3.4: A schematic of our vacuum system as viewed from above, showing the directions of the laser beams, calcium atom beam, and the static magnetic field. The 850nm (π) laser beam enters through the top viewport at $\sim 45^\circ$ to the trap surface.

(which can be done automatically when an ion is lost) taking a few minutes (the typical single-ion lifetime is ~ 1 day).

3.3 Magnetic field coils

3.3.1 Main coils

An important feature of our qubit transition is that it is an “atomic clock” transition at 146G, which means that its frequency is independent of magnetic field to first order. This means that fluctuations in the ambient magnetic field have little effect on the qubit frequency, suppressing the decoherence caused by these fluctu-

ations. The fact that this field independence occurs at a non-zero magnetic field is also desirable, as one can distinguish the different Zeeman states in the ground level; for zero-field hyperfine clock transitions, a small field is required to break the degeneracy, introducing a first-order frequency dependence on magnetic field.

To generate a magnetic field of 146G we have a pair of magnetic coils in a Helmholtz-like configuration around the vacuum chamber. The rectangular coil shapes were designed such that they do not impede laser beams propagating through the centres of any of the six side viewports. Each coil consists of hollow copper wire¹⁴ wrapped 16 times around an aluminium former and set with epoxy resin¹⁵. The coils are powered by a pair of power supplies¹⁶ operating in series, together generating 135A with 10V across the coils, which are also connected in series. The coils are cooled to $\sim 40^{\circ}\text{C}$ by circulating liquid coolant¹⁷, which is cooled with a chiller¹⁸, through the centres of the wires. Coolant is used instead of water to prevent biofouling, which can reduce the flow rate over time. The coils are connected in parallel for the coolant circulation.

In order to ensure that the magnetic field is constant, a small amount of current is picked off from the power supplies; the amount of current picked off is adjusted in response to fluctuations in both the supply current and the ambient field. The control of this current is described in chapter 4.

3.3.2 Trim coils

In order for the laser beams and microwaves to have the correct polarizations with respect to the quantization axis, it is important for the static magnetic field to be in

¹⁴Wolverine Tube Europe OFHC copper 3.2mm x 3.2mm square cross section with 2mm diameter bore, Kapton wrapped by S & W Wire Company

¹⁵Araldite 2011 epoxy

¹⁶Keysight Technologies/Agilent, 6671A

¹⁷Applied Thermal Control Ltd, Hexid heat transfer fluid

¹⁸Parker Hiross Hyperchill Plus ICEP007 (50Hz)

the correct direction. To correct for any slight misalignment in the main coils, as well as any ambient field in directions other than the quantization direction, we use three trim coils powered by separate supplies¹⁹: two around the top and bottom flanges to produce a field in the vertical direction and one around the microwave feedthrough flange to produce a field in the horizontal plane at 90° to the quantization axis.

3.4 Lasers, optics and imaging

3.4.1 Photo-ionization lasers

As the ^{43}Ca isotope used in our experiments has relatively low abundance (even in our isotopically-enriched source, which has an abundance of only $\approx 12\%$), it is important that our ionization is isotope-selective [60]. We achieve this with a two stage photo-ionization scheme using two laser frequencies, one of which excites only ^{43}Ca neutral atoms, and one of which provides photons of sufficient energy for ionization to occur in these excited atoms (see section 2.4.2 for the details of our photo-ionization scheme). The first of these laser beams is at 423nm, produced by a diode laser²⁰ and frequency stabilized by feeding back on the diode voltage from the output of a wavelength meter²¹ every 10s. The second beam is at 375nm, produced by a free-running diode²² (without an external cavity). This laser is not frequency locked as there are no strict requirements on its frequency stability. The two beams are switched on and off with a mechanical shutter.

3.4.2 Ca^+ control lasers

The five Ca^+ control laser beams are generated by diode lasers²³ and frequency stabilized to reference etalons²⁴. The lengths of these etalons, and therefore the laser

¹⁹Aim-TTi - Thurlby Thandar Instruments Limited, QL series power supplies

²⁰TOPTICA Photonics AG, DL pro

²¹HighFinesse/Ångstrom, WS7-60

²²TOPTICA Photonics AG, iBEAM-SMART 375-S_12274

²³TOPTICA Photonics AG, DL100

²⁴NPL “Low Drift Etalons”, specified drift rate of less than 0.5MHz/hour

frequencies, are controlled with piezo actuators. For the 397nm, 393nm and 866nm lasers, the frequency is stabilized with Pound-Drever-Hall locks. For the 850nm laser, the frequency stability is less critical, so we use the simpler method of a side-of-fringe lock. The frequency stability of the 854nm laser is so unimportant that we do not implement a lock, but rather monitor the frequency using the wavelength meter and manually tune the etalon piezo voltage around once per day to correct for long-term drifts. We monitor the output frequencies and spectra of the lasers by feeding picked-off samples of the beams to our wavelength meter and a pair of optical spectrum analysers, one for the blue lasers²⁵ and one for the red²⁶.

There are optical isolators immediately after the output of each laser to protect them from retro-reflected light, which could potentially result in optical feedback. The 397nm, 866nm and 850nm beams are filtered with blazed diffraction gratings to remove any amplified spontaneous emission, which could result in unwanted ion shelving or deshelling. Each beam is switched on and off with double-pass acousto-optic modulators (AOMs). These AOMs are also used to stabilize the beam intensities at the ion as described in section 3.4.3.

While the AOMs provide good extinction of the laser beams when switched off and with the RF signal frequencies detuned from the values used for beam transmission, there is still some leakage. For the 397nm beam, this leakage is sufficiently large to interfere with our longer experiments, where there can be periods of tens of seconds between state preparation and readout. To prevent leakage in these experiments, the 397nm beam is not only blocked with an AOM, but also with a mechanical shutter. After the AOM and shutter, the beam goes through an electro-optic modulator²⁷ operated at 2.93GHz, which produces sidebands of 85% intensity relative to the carrier. This allows for repumping of both hyperfine manifolds of the $4S_{1/2}$

²⁵TOPTICA Photonics AG, FPI100

²⁶Technical Optics, SA-7.5

²⁷New Focus, 4431M resonant EOM

level when Doppler cooling and detecting fluorescence.

The two frequencies of 866nm light, one for high-fluorescence Doppler cooling and one for dark-resonance cooling, are each produced by a separate double-pass AOM, set to produce a relative offset of 360MHz. These AOMs are then used to switch between the two frequencies.

3.4.3 Trap optics

The laser systems as described above are on a separate optical table to the ion trap vacuum system, with optical fibres carrying the beams to the main experiment table. These fibres introduce noise into the beams in two ways. Firstly, any small variation in the beam pointing at the input causes a large change in the coupling efficiency, and thus the output power. Secondly, changes in the fibre temperature and/or strain change the birefringence, causing variations in the output beam polarization. The output of each of these fibres is passed through a polarizing beam splitter to ensure that the beams' polarizations are well-defined, turning polarization variations into power variations.

To stabilize the power of each beam, a sample is picked off after each polarizing beam splitter and sent to a photo-diode²⁸. The signal from each photo-diode is fed to an FPGA board²⁹, which implements a PID loop to stabilize the beam intensity by varying the amplitude of the corresponding AOM RF drive earlier in the beam path (before the fibre input). After the photo-diode pickoff, each beam is focussed onto the ion with a 200mm achromatic doublet lens.

²⁸Thorlabs, PDA-36A

²⁹National Instruments Corporation, PCIe-7852R

3.4.4 Imaging system

The ion fluorescence is collected through the top viewport of our vacuum chamber into an optical system mounted above the chamber, which focusses the light onto either a photomultiplier tube³⁰ (PMT) for qubit state detection or a camera³¹ for diagnostic purposes. The focussing is achieved using a custom-made five-element lens, which has a specified transmission at 397nm of 95% and a numerical aperture of 0.3. There is an adjustable aperture in the focal plane of the beam to remove most background light. There are also spectral filters to further reduce the number of background counts on the PMT. The total detection efficiency of the PMT system (i.e. the probability of detecting any particular photon emitted by the ion) is 0.23%. See figure 3.5 for a schematic of the imaging system.

3.5 Microwave systems

3.5.1 Frequency reference

The frequency reference for our microwave signals is an atomic rubidium frequency standard³². The 10MHz frequency output of this reference is specified to drift by $< 5 \times 10^{-11}$ per month³³, and be stable to $< 1 \times 10^{-10}$ for temperature changes between 10°C and 40°C, and $< 2 \times 10^{-10}$ for a 1G magnetic field reversal. The specified initial accuracy is 5×10^{-11} and the annual drift is specified to be $< 5 \times 10^{-10}$, so given that the standard was purchased in 2005, at the time of writing we expect a frequency error of $< 7 \times 10^{-9}$. By beating the signal against an oscillator locked to GPS time, we find that the standard frequency is slightly higher than specified, with an error of $\sim +7 \times 10^{-10}$, which is well within the specifications.

³⁰ET Enterprises Limited, P25PC-12

³¹Andor Technology Ltd, Luca DL-406M-OEM

³²Stanford Research Systems, FS725

³³All specifications stated here are as a fraction of the 10MHz target frequency.




Figure 3.5: Schematic of the imaging system (not to scale). The beam splitter is removed from the beam path during most experiments to maximize the PMT detection efficiency and thus the ion readout fidelity.

The output of the rubidium standard is used to clock a signal generator³⁴, which steps the frequency up to $\sim 1\text{GHz}$ (we made small changes to the nominal frequency around once per week to correct for frequency drifts). This generator has an output resolution of 10mHz .

3.5.2 Qubit microwaves

The 1GHz output of our signal generator clocks a direct digital synthesizer (DDS) source³⁵. There are eight output profiles on the DDS, each with its own frequency, phase and amplitude settings – switching between profiles (with different phases, for

³⁴Rohde & Schwarz, SMA100A

³⁵Analog Devices, Inc., AD9910

example) allows us to perform rotations around different axes on the Bloch sphere. The output of the DDS is frequency-octupled with the chain of components depicted in figure 3.6, similar to the configuration in [49]. The microwaves are turned off by changing the DDS profile to one with zero output power, and closing two microwave switches operating in series after our frequency octupler.




Figure 3.6: A schematic view of our qubit microwave frequency octupler, with the part number for each component. All components are made by Mini-Circuits except for the isolator, which is made by Semtech. Most components were bolted to a metal plate, with semi-rigid SMA cable connections. The attenuators were included to allow for fine-tuning of the microwave power. Different attenuators were used for Ramsey experiment memory error measurements (see section 5.1) to give a different output microwave power. The switch configuration for that experiment was also different, with both switches in the ‘forward’ direction (switching between two outputs rather than two inputs) and no detuned microwave source (see section 3.5.2 for an explanation of why the switch configuration was changed for our randomized benchmarking experiments). All other components were the same for all experiments presented in this thesis.

A significant change to the previous setup (which is detailed in [50]) is the switch configuration. We found that there are transients on the microwave power being sent to the trap, both during each microwave pulse (which are separated by $12\mu\text{s}$ in our experiments to allow the DDS sufficient time to reliably switch between profiles), and over the first ~ 10 pulses following a longer period without application of microwaves. See figure 3.7 for plots of the power transients over these two timescales.

Simple randomized benchmarking of gates (see section 2.1) is only sensitive to the faster of these two transients. For an experiment involving thousands of gates with short inter-gate delays, the slower transient only affects the fidelities of the first few gates, so when dividing the total sequence error rate by the number of gates this contribution becomes insignificant. However, when measuring the error of a delay with interleaved randomized benchmarking (see section 5.2.3), each Clifford gate is preceded by a significant delay. This means that, when comparing the average error of a Clifford gate to that of the combination of a delay and a Clifford gate, the calculated error of the delay will also include that due to the microwave power transient, introducing a systematic inaccuracy into the measurement.

In principle, one could use direct measurements of the microwave power transient to estimate the resulting extra gate infidelity and subtract this from the difference in error measured between interleaved and non-interleaved experiments. However, a more reliable approach is to reduce the transients to the point at which their effect is negligible. We found that, for our switches³⁶, the size of the slow transient could be reduced by a factor ~ 2 by operating the second switch in reverse, i.e. with the input switching between two ports rather than the output.

In addition to reducing the slow power transient, having the second switch operating in reverse allows us to keep its output port consistently warm (reducing thermal

³⁶Mini-Circuits, ZASWA-2-50DR+




Figure 3.7: Power transients on our 3.2GHz qubit microwave source for different switch configurations. The three configurations are with both switches in the forward configuration (switching the output between two ports rather than the input), and with the second switch in the backward configuration with and without a detuned microwave source connected to the other input of this switch. The three plots are scaled as a fraction of the final power and offset from one another by 0.01dB for clarity. The measurements were performed by connecting the microwave output to our diplexer via an ATC ATc3-6 circulator and measuring the reflected power from the third circulator port with a Mini-Circuits ZX47-40LN+ power detector connected to an oscilloscope averaging over 1000 sequences. The microwave power was turned on for $10\mu\text{s}$ pulses, with DDS profile switching times of $12\mu\text{s}$. We applied sequences of 32 of these pulses (to be roughly similar to our composite gate sequences, as described in section 5.2.2) with inter-sequence delays of 1ms. The first $1.6\mu\text{s}$ of each pulse measurement contained a large upward transient, which is likely an artefact of the measurement and is omitted for clarity. (a) The transients over the entire sequence. For the measurements without the detuned microwave source (blue and red), the transients over each gate extend above the upper plot limit and are omitted for clarity. (b) A close-up of the final microwave pulse showing the entirety of the fast transients.

effects) by connecting another microwave source to the second input of the switch. We find that this removes the fast transients occurring over each microwave pulse (see figure 3.7). The second microwave source is a signal generator³⁷ with the power chosen to be similar to that of the input to the other switch port.

For the choice of frequency of this second microwave source, there are competing considerations. We would like it to be similar to our qubit transition frequency of 3.2GHz, as this then causes the thermal behaviour of the switch, diplexer and trap to be the same as for when we apply the output of our frequency octupler. However, it must be sufficiently detuned that it does not drive transitions in our ion. In addition, the AC Zeeman shift induced onto our qubit transition by the second microwave source should be fairly small, as any variations in this shift due to fluctuations in power would otherwise result in qubit decoherence (if the shift were perfectly constant then it would be corrected for when calibrating the qubit frequency with Ramsey scans). We choose a frequency of 3.8GHz, which results in an AC Zeeman shift on our qubit transition of $\sim 100\text{mHz}$.

Another method that we use to reduce the slower power transients after periods without applying microwaves is to turn the DDS on some time before turning the switches on. This allows us to combine the power transients of the switches and the DDS in a more destructive way (i.e. with one partially cancelling the effect of the other), reducing the overall transient. The optimal time interval depends upon the delay length, so we calibrate it for each delay length for which we measure our qubit memory error. Another consideration in the choice of time interval is that a significant microwave power reaches the trap electrode when the DDS is on even when the switches are off, resulting in a small extra qubit rotation. Our choices of time intervals are discussed further in section 5.2.3.

³⁷Keysight Technologies, N5181B

3.5.3 State preparation and readout microwaves

For the microwave pulses used in state preparation and readout (i.e. all microwave pulses addressing transitions other than the qubit transition), we use the same oc-tupler setup as presented in [49, 50]. This is slightly different to the setup used for our qubit microwaves, with the switches both operated in the forwards direction, different attenuators and an extra amplifier.

3.6 ARTIQ control system

The majority of our experimental control is done with ARTIQ (Advanced Real-Time Infrastructure for Quantum physics). This is a control system designed for use in quantum information experiments, initiated by the Ion Storage Group at the National Institute of Standards and Technology and developed by M-Labs, with our group and others also contributing. The high-level programming language is based on Python, where experiments are carried out by running Python scripts which interface with the various different pieces of hardware. ARTIQ achieves nanosecond timing resolution by compiling and executing time-critical code on FPGA hardware. We use a KC705 development board from Xilinx as our hardware for the control system.

The experiment Python scripts execute tasks either in series or in parallel (we often need to perform multiple tasks at the same time, e.g. simultaneous pulses of several laser beams). This allows us to modularize the experiments, with common tasks (such as initializing and reading out the qubit state) being called as functions within higher-level scripts, with the code for these tasks rarely changing after being written and tested the first time.

We automatically monitor parameters of interest by querying these parameters in the for/while loops of the experimental code, performing other tasks when neces-

sitated by the parameter values. For example, our experiment scripts generally include regular checks of the ion fluorescence rate, and if ion loss causes this to drop below a certain value then a function is called to load a new ion. This function again monitors the fluorescence and returns to the main script when a new ion has been loaded. We can also repeat certain tasks independently of the main experiment script. A separate control script for our fluxgate magnetic field feedforward system (see section 4.3 for details of this system) regularly queries the measurement of our fluxgate sensor and makes corresponding adjustments to the current going through our main coils, all independent of and in parallel with any experiments.

We start and monitor experiments via a graphical interface, which can be simultaneously used by multiple computers to interact with the same experimental system. The interface allows us to easily adjust different parameters each time experiments are performed. Multiple experiments can be scheduled, with the interface proceeding to the next experiment when the current one has finished, and with priority codes allowing us to move important experiments to the front of the queue. There are also different databases displayed in the graphical interface, with the database values updated by commands in the experiment Python script, allowing us to monitor experiments in real time. We can also plot these databases on graphs to give live visualizations of experimental results.

Chapter 4

Magnetic field control

In this chapter, we describe our control of the static 146G ($1\text{G} = 100\mu\text{T}$) magnetic field needed to make our qubit transition frequency independent of magnetic field to first-order (an “atomic clock” transition). This control is important to maximize the qubit coherence time, and to minimize the state preparation and measurement errors resulting from variations in the frequencies of the transitions used to prepare and read out the qubit state (which are sensitive to magnetic field).

Our control consists of five parts. In section 4.1, we describe our feedback system on the coil current, which suppresses fast noise in the output of the power supplies. Section 4.2 details our work on a feedforward current applied to cancel out ambient magnetic field noise at the mains electricity frequency of 50Hz and harmonics. These two control methods are implemented with the same set of circuit boards, for which much of the design work, including the notch and bi-quadratic filters, was carried out by Benjamin Merkel. The circuit diagrams are in appendix B. As well as designing and building the circuit boards and the rest of the hardware, our development work involved programming a single-board computer and a microcontroller, as well as testing the system.

In section 4.3, we describe our use of a magnetic field fluxgate sensor to measure and automatically correct for slow drifts in the ambient magnetic field. Section 4.4 details how we directly use the ion to occasionally correct for residual drifts that are not dealt with by the fluxgate feedforward. Finally, in section 4.5 we describe how we determine the correct set-point for the ion servo to ensure that we are aiming for the right magnetic field. Figure 4.1 gives an overview of the combined control system.




Figure 4.1: An overview of our control system for the static 146G magnetic field.

4.1 Feedback

Our main coil power supplies were chosen for their relatively quiet output, but given the high current of 135A required to generate our magnetic field of 146G, the absolute noise on this current is significant for our experiments. We suppress this noise by measuring the current with a current transducer and a measuring resistor, subtracting a reference signal, and feeding back on fluctuations in the residual signal, shunting away a small amount of current corresponding to these fluctuations. See figure 4.2 for an illustration of the feedback loop. The bandwidth of the loop is 3kHz. The different electronic components were carefully chosen to minimize the

noise induced onto the measured signal, as this noise is propagated through the feedback loop and induces additional noise on the coil current.




Figure 4.2: An illustration of our main coil current feedback loop.

The current passing through our main magnetic field coils is first passed through a current transducer¹. This works by inducing a magnetic flux around a loop of magnetic material, and generating a current in a secondary wire wound around the material to produce an equal and opposite magnetic flux. The residual magnetic flux is measured by applying a square voltage wave onto another wire, also wound around the magnetic material, at a frequency of $\sim 17\text{kHz}$ and measuring any asymmetry in the voltage signal across the measuring resistor connected to the secondary current (which would indicate a non-zero average magnetic flux). We use a measuring resistor² with a Kelvin connection, which means that two pairs of terminals are connected directly to the resistive material rather than through some shared conductor, avoiding the fluctuations in resistance that this can cause.

The voltages on either side of the measuring resistor are passed through a low-pass filter with common-mode and differential bandwidths of 2MHz and 100kHz respectively to remove the very fastest noise (which is beyond the loop bandwidth of 3kHz). This fast noise could otherwise be rectified into DC voltage shifts by

¹LEM, IT 200-S ULTRASTAB

²Vishay Precision Group, VCS1625Z

nonlinearities in the op amps. There is then a differential amplifier, which transforms the differential signal into a single voltage relative to the circuit board ground.

A voltage corresponding to our desired coil current is subtracted from the measured signal, resulting in a voltage which represents the surplus current in the system (i.e. resulting in a magnetic field greater than the target value of 146G). This surplus is chosen to always be positive, as the system is only capable of shunting away current, not providing extra. The target voltage is generated by a pair of 16-bit Digital to Analogue Converters (DACs)³ where the output of one is amplified by a factor of 100 relative to the other, with the combination therefore having an effective resolution of 22 bits. Both DACs produce voltages which are some fraction of that provided by a voltage reference⁴.

The feedback DAC output voltages are controlled via a Beaglebone Black single-board computer. This computer receives commands to change the DAC voltages from the main experimental control computer via an ethernet cable, and it then reprograms the DAC values via a Serial Peripheral Interface.

After the reference voltage subtraction, there is a notch filter at 17kHz and a sharp low-pass filter, both of the design in [73]. The purpose of these filters is to suppress the 17kHz oscillation generated by the current transducer. This suppression is, however, not perfect, and the feedback loop bandwidth is chosen to be 3kHz in order to ensure that the residual noise at 17kHz does not increase the noise in the coil current.

Another consideration is the large output capacitance of the coil power supplies. The purpose of this capacitance is to suppress high-frequency noise on the output, for which the capacitance is of low impedance and so is a favourable dissipation

³Analog Devices, AD5541A

⁴Linear Technology, LTC6655 (2.5V)

route. However, for the same high frequencies that our system attempts to shunt away from the coils, whose high inductance means that they have high impedance at high frequencies, most of the current is taken from the supply output capacitance rather than the coils. To ensure that a larger current is shunted away to account for this effect, we include a bi-quadratic filter of the design in [74].

Next there is a type II loop filter, which provides gain for the feedback loop and defines the loop bandwidth. This is followed by a current shunt connected across the coils, which dissipates a significant amount of heat and so is on a separate circuit board to improve the temperature stability of the main board (for some parts of the feedback circuit, temperature fluctuations result in uncompensated fluctuations in the coil current). This shunt is achieved with a pair of transistors; while this could in theory be done with a single transistor, the high gain required for our shunt makes it more convenient to instead use two transistors in combination. Figure 4.2 shows each of the components in the feedback loop.

Figure 4.3 shows the power spectral density of the main coil current with the feedback loop open and closed. There is a significant noise suppression at frequencies below our loop bandwidth of 3kHz. The 50Hz mains electricity causes large spikes for both spectra at harmonics of this frequency. There are large spikes at $\sim 17\text{kHz}$ generated by the current transducer; as explained above, this noise should be largely filtered out before the signal reaches the current shunt circuit.

To measure the magnetic field noise experienced by the ion, we use the transition between the $|4, +4\rangle$ and $|3, +3\rangle$ states in the $4S_{1/2}$ ground level, as the frequency of this transition is first-order sensitive to changes in the magnetic field ($df/dB = -2.36\text{kHz mG}^{-1}$). We perform a series of Ramsey phase scans, measuring how the Ramsey fringe contrast decreases with increasing Ramsey delay. Figure 4.4 shows the result of this measurement, which was triggered by the mains electricity phase




Figure 4.3: The spectral density of the magnetic field noise caused by noise in the coil current, as deduced from measurements of this current, with the feedback loop open and closed. The measurement is of the voltage on the feedback circuit after the reference subtraction, so it includes noise generated by the circuit before this point in the loop. There is a significant reduction in the noise up to the feedback loop bandwidth of 3kHz. The spike at ~ 17 kHz should be suppressed by filters later in the feedback loop as explained in the main text. The noise at harmonics of 50Hz is further suppressed with a feedforward system as described in section 4.2.

to minimize the contribution of coherent mains noise (see section 4.2 for details of our efforts to suppress this noise). We find that the contrast decays with a Gaussian shape, fitting the model of noise that is correlated over timescales much longer than the delay period (see equation 2.15). The fit gives a decay time constant of 0.70(6)ms, which corresponds to a magnetic field noise of 0.14(1)mG RMS.




Figure 4.4: The decay in Ramsey fringe contrast on the $|4, +4\rangle$ to $|3, +3\rangle$ transition in the ground level with increasing Ramsey delay. Each Ramsey experiment is triggered by the mains electricity phase. The fit is to the Gaussian function $e^{-(t/T)^2}$, giving a decay time constant of $T = 0.70(6)$ ms.

4.2 Mains electricity feedforward

The magnetic field coil power supplies, as well as other equipment in the laboratory, produce a significant amount of noise at the oscillation frequency of the mains electricity, 50Hz, and at harmonics of this frequency. A large part of this noise is coherent over each 20ms period, which means that we can measure the noise over this cycle (triggering the measurement from the mains electricity phase) and then shape the coil current with the opposite variation to cancel the noise.

To achieve this cancellation, we use very similar circuitry to that presented in section 4.1 (most of which shares the same circuit board). Another DAC and voltage reference provide a voltage signal corresponding to the current that we want to shunt away from the coils. This signal is then applied to a second current shunt circuit,

with a pair of transistors to shunt away the current.

A significant difference to the feedback system is that rather than making occasional adjustments to the DAC output voltage, it is now necessary to change it many times over each 20ms mains electricity cycle. To achieve this, we use a microcontroller⁵ to program the DAC with new values at a high repetition rate. This microcontroller calculates the voltages required for the combination of sinusoids at 50Hz, 100Hz, 150Hz, 200Hz and 250Hz with preprogrammed amplitudes and phases, where the overall output signal is triggered by the mains electricity phase. The programmed amplitudes and phases are occasionally updated via the same single-board computer used to update the feedback DAC voltages. They are calibrated by measuring the magnetic field on the ion's field-sensitive $|4, +4\rangle$ to $|3, +3\rangle$ transition over the mains electricity cycle. Each harmonic of the feedforward system is then applied in turn with a certain amplitude and phase, and the measured magnetic field variation for each harmonic is used to calculate the optimum combination of harmonics to cancel the ambient field fluctuation.

Figure 4.5 shows our measurement of the magnetic field fluctuation over the 20ms mains electricity cycle, before and after applying our feedforward cancellation. The measurement consists of a series of Ramsey experiments on the field-sensitive transition, using the shift in the central Ramsey fringe to determine the change in the magnetic field. The graph shows a significant reduction in the magnitude of the magnetic field fluctuation.

We also repeat the Ramsey fringe contrast experiment described in section 4.1 without triggering off the mains electricity, making the experiment sensitive to coherent field fluctuations over each mains cycle. Without the coil current feedforward, the fitted RMS magnetic field noise is 1.34(8)mG, which is in agreement with the value

⁵Microchip Technology Inc., PIC32MX250F128B




Figure 4.5: Measurement of coherent magnetic field fluctuations over the 20ms mains electricity cycle, with and without feedforward variation of the magnetic field coil current. The measurement consists of a series of Ramsey experiments on the $|4, +4\rangle$ to $|3, +3\rangle$ transition at different times after being triggered by the mains electricity phase. The magnetic field shift at each time is then determined by the frequency shift of the central Ramsey fringe. The RMS magnetic field variation of the two data sets is $1.30(5)\text{mG}$ and $0.14(6)\text{mG}$ without and with the feedforward current respectively.

of $1.30(5)\text{mG}$ determined by measuring the magnetic field over the 20ms mains cycle (see figure 4.5). After applying the feedforward, the Ramsey fringe contrast decays much more slowly, giving a fitted RMS field noise of $0.30(2)\text{mG}$. This value is in rough agreement with the combination of the RMS noise measured with Ramsey experiments triggered from the mains ($0.14(1)\text{mG}$, see section 4.1) and the residual coherent noise measured over the mains cycle ($0.14(6)\text{mG}$, see figure 4.5), which sum in quadrature to give a value of $0.20(6)\text{mG}$.

It should be noted that this system only provides cancellation in the direction of the main 146G magnetic field, as only the current through the main field coils is varied. This means that any variation in the magnetic field perpendicular to this direction

results in small rotations of the overall magnetic field direction, causing variations in the polarizations of the laser and microwave radiation (as these polarizations are defined relative to the magnetic field direction). However, we find that this effect does not cause any problems for our experiments.

4.3 Fluxgate magnetic field feedforward

As well as the fast magnetic field noise addressed in sections 4.1 and 4.2, it is also necessary for us to correct for slow drifts. This correction can be done by iteratively using the ion to measure the field and then applying a correction to the coil current (see section 4.4), but this process takes ~ 10 s, time which would need to come out of whatever experiment was being done, increasing the total experimental time. An alternative approach is to measure the magnetic field with a separate sensor near the ion trap, which is not quite as accurate as using the ion itself, but this can be done in parallel with experiments. We use a fluxgate sensor⁶ for this purpose, reading the magnetic field automatically every 2s. These measurements are then used to adjust the reference voltage in the feedback circuit (see section 4.1) to change the amount of current being shunted away from the main field coils accordingly.

Figure 4.6 illustrates the performance of the fluxgate feedforward over an example period of time. The ambient magnetic field changes significantly due to the movement of a nearby lift, and the change measured with the fluxgate sensor is used to automatically update the main coil current accordingly. Over this period, we also regularly servo the field with measurements on the ion as described in section 4.4, making corrections to the current set by the fluxgate feedforward. The steps in field due to the lift operation are entirely cancelled out by the fluxgate feedforward, with no significant correction necessary with the ion servo.

⁶Bartington Instruments, Mag670-I-1000




Figure 4.6: An example period of time illustrating the performance of the fluxgate magnetic field feedforward. The ambient magnetic field (as measured by the fluxgate sensor) changes significantly with the movement of a nearby lift. This change is automatically cancelled by making corresponding changes to the main coil current. Over this period, we also measure the field directly with the ion, making additional corrections to the coil current where the value set by the fluxgate servo is not quite correct. These two plots are offset from one another by 1mG for clarity.

4.4 Ion magnetic field servo

To exactly correct for magnetic field drifts experienced by the ion, it is necessary to use the ion itself to measure the field. We do this by driving single π -pulses on the $|4, +4\rangle$ to $|3, +3\rangle$ transition, each time detuning the microwave frequency to one side of what the resonance frequency should be at the target magnetic field. These detunings are chosen to result in an excitation probability of 0.5 (i.e. at one FWHM from the line centre to either side of the resonance peak). Any imbalance in the populations measured for each detuning corresponds to an offset in the magnetic field. After each such measurement, the magnetic field is corrected by changing the reference voltage in the feedback circuit (see section 4.1) to adjust the amount of

current that is being shunted away from the main field coils.

This process is repeated until the measured imbalance passes below some threshold, at which point the microwave amplitude is halved and the iterative process repeated. By halving the microwave amplitude (i.e. doubling the pulse π -time), we halve the width of the frequency resonance peak and therefore double the sensitivity to frequency offsets. We go through a series of such microwave amplitude reductions until we reach a π -time of $400\mu\text{s}$. For this microwave power, we can achieve a magnetic field accuracy of $\sim 100\mu\text{G}$. This entire process is automated and takes $\sim 10\text{s}$ each time.

4.5 Field-independent point measurement

To ensure that the target frequency of the $|4, +4\rangle$ to $|3, +3\rangle$ transition in the ion servo corresponds to the magnetic field that makes the qubit frequency first-order independent of fluctuations in the field, we measure the quadratic variation of the qubit frequency with magnetic field deviation and program the control system to aim for the centre of this curve. Figure 4.7 shows an example of such a measurement. The frequency offset for each field value is determined by performing a Ramsey experiment, scanning the phase of the second $\pi/2$ -pulse to measure the phase at the peak of the Ramsey fringe. We are able to achieve a precision of $\sim 20\mu\text{G}$ with a field scan taking ~ 10 hours. We repeat the measurement every few months to check for drifts.




Figure 4.7: An example measurement of the qubit field-independent frequency point. The magnetic field is scanned over the point which results in this field-independence. At each field point, the qubit frequency is measured with a Ramsey experiment, scanning the relative phase of the two $\pi/2$ -pulses to measure the shift in the Ramsey fringe peak. The quadratic fit to the data gives a magnetic field offset of $-0.06(2)\text{mG}$.

Chapter 5

Short-timescale qubit decoherence

Qubit decoherence, or memory error, is one of the challenges faced in the development of a large-scale quantum computer. The accuracy of the results of a quantum computation relies on the preservation of the quantum information stored in qubits over the duration of the computation. Quantum error correction can help with this preservation, but the necessary number of physical qubits scales unfavourably with the size of the errors that must be corrected, resulting in a strong motivation to minimize these errors.

The standard method for probing quantum decoherence is to perform Ramsey experiments with a variety of delay periods t . After correcting for the SPAM error, the decay of the Ramsey fringe contrast C is then fit to an exponential function:

$$C = e^{-t/T_2^*}, \tag{5.1}$$

where the coherence time T_2^* characterizes the memory lifetime.

In general, these experiments are performed with long Ramsey delays (seconds or tens of seconds for trapped ions), where the corresponding contrast decay is 10^{-1} or larger, and memory errors over shorter periods are inferred by extrapolation of equation 5.1. The main reason for this is that SPAM error is generally $\sim 10^{-2}$, so directly probing smaller errors requires an increasingly large amount of data, placing limitations on the experiments. However, this estimation of memory errors over short time periods relies on the assumption that equation 5.1 is a good model in this regime, an assumption that has not, to our knowledge, previously been tested for trapped-ion qubits.

In section 5.1, we present our results on pushing this approach to probe memory errors at the 10^{-3} level. This is made possible by our relatively low SPAM error of $\sim 10^{-3}$.

To probe even smaller memory errors without a reduction in SPAM error, it is necessary to remove this limitation on experiments. This can be achieved by interleaving many short delay periods between randomized quantum gates [53, 75], which amplifies the memory error relative to the SPAM error. The limiting factor is then the size of the single-qubit gate error, but this is generally much smaller than the SPAM error.

In section 5.2, we present the results of such an experiment, probing memory errors below 10^{-5} . We find that the common assumption of exponential decay does not hold at these short delay periods, with memory errors much lower than would be predicted by fitting data at long periods to an exponential decay function. We find that the decoherence could instead be caused by $1/f$ noise.

Finally, in section 5.3 we present the results of an experiment on suppressing our qubit decoherence with a dynamical decoupling sequence. Fitting an exponential

decay model to delay times between 10s and 60s gives an effective coherence time of 1000(500)s.

5.1 Ramsey experiments

Ramsey experiments are a common experimental tool in quantum physics. They involve preparing the qubit in $|0\rangle$ or $|1\rangle$, applying a $\pi/2$ -pulse, waiting for some delay time, applying a second $\pi/2$ -pulse with some fixed phase relative to the first, and measuring the final state, repeating this sequence many times for a range of phase differences between the $\pi/2$ -pulses. The population P in $|0\rangle$ varies sinusoidally with the pulse phase difference ϕ :

$$P = \frac{C}{2} \cos(\phi - \phi_0) + A, \quad (5.2)$$

where C is called the Ramsey fringe contrast, ϕ_0 is some phase offset, and $A \approx 0.5$ is the mean value of P over the fringe. The contrast is ideally equal to 1, but in the presence of qubit decoherence and SPAM errors, it is generally significantly smaller than 1. See figure 5.1 for an example set of results from a Ramsey experiment.

Based on the assumption that the target qubit state in any given computation has equal probability of being $|0\rangle$ or $|1\rangle$, the average qubit memory error is:

$$r_{mem} = \frac{1 - C}{2}. \quad (5.3)$$

To correct for the SPAM error to isolate the error caused by decoherence, one simply needs to measure the SPAM error separately and subtract it from r_{mem} .

In investigations of qubit decoherence, the Ramsey fringe contrast is generally determined by performing Ramsey experiments with a range of phases and fitting the




Figure 5.1: An example set of results from a Ramsey experiment, showing how the population in $|0\rangle$ varies with the phase of the second $\pi/2$ -pulse relative to the first. The Ramsey delay between the two $\pi/2$ -pulses was 3s.

data to a sinusoid, floating the amplitude and phase of this sinusoid. However, the contrast extracted in this way does not generally give an accurate value for the memory error in the context of quantum computing. For example, if there is some slow frequency drift that is unknown and uncorrected, this would not affect the Ramsey fringe contrast, only the phase offset. However, this unknown phase offset would increase the qubit memory error in a quantum computation. This means that to accurately determine the memory error one must predict the phases of the peak and trough of the Ramsey fringes, measuring the qubit state at these points. Any error in the estimation of these phases is then included in the measurement, as would be the case in a quantum computation. Figure 5.2 illustrates this effect on the Bloch sphere.

A major challenge with this approach is that any detuning between the qubit transition and the radiation used to address the qubit (in our case, near-field microwave radiation) moves the phase of the Ramsey fringe peak. Any uncorrected drift in this




Figure 5.2: An illustration on the Bloch sphere of the qubit state evolution during an example Ramsey experiment with a non-zero frequency detuning error. The initial state is $|1\rangle$, which is followed by two $\pi/2$ -pulses that are of equal phase and separated by some Ramsey delay. During the inter-pulse delay, the detuning error causes the qubit state to rotate in the x - y plane, causing the final state to differ from the target state $|0\rangle$.

detuning increases the measured memory error. Fortunately, we find that the detuning in our system drifts by $\sim 10\text{mHz}$ over a period of a few days, as illustrated in figure 5.3. This means that this source of error only becomes significant for Ramsey delay times longer than $\sim 1\text{s}$.

For calibrating the microwave frequency, we measure the detuning with Ramsey experiments and make adjustments to the frequency on either the synthesizer (which is clocked with an atomic Rubidium clock frequency reference) or the qubit microwave




Figure 5.3: The drift of the detuning between the qubit transition and the microwaves over an example period of time. The measurements were taken with Ramsey experiments, using a Ramsey delay of 3s.

DDS (which is clocked with the synthesizer). The synthesizer allows for higher precision in setting the qubit frequency¹, but this is still larger than the offset caused by drifts. We correct for any offset caused by this limited resolution by calculating the corresponding phase of the Ramsey peak for each Ramsey delay time, adding this phase to that of the second $\pi/2$ -pulse (which would be 0° and 180° for zero detuning).

To correct for SPAM errors, we follow each Ramsey experiment (consisting of a single state preparation, Ramsey pulses and delay, and state readout sequence) with a SPAM measurement (the same, but without the Ramsey pulses). We alternate the final state of the SPAM measurement with an additional π -pulse to follow the target state of the Ramsey experiment, which alternates between the peak and the trough of the Ramsey fringe. This interleaving means that our SPAM error measurements

¹The synthesizer frequency resolution of 10mHz on the output of 1GHz (which is increased to 3.2GHz by the DDS and frequency octupler) gives a qubit microwave frequency resolution of 32mHz.


are robust to drifts in the SPAM error over timescales longer than that of a single experimental sequence, ensuring that we are using the right values for our SPAM error correction. The delay between state preparation and readout for each SPAM error measurement is set to that of the corresponding Ramsey experiment, which accounts for errors caused by heating of the ion over this period. The SPAM errors are then subtracted from the raw memory errors to obtain corrected values.

We present the results of this experiment in figure 5.4. This direct measurement establishes an important limit on the memory error. We constrain the memory error to be $< 10^{-3}$ for delays $\lesssim 100\text{ms}$. This timescale is much longer than typical gate and SPAM times ($\lesssim 1\text{ms}$).


The range of Ramsey delays for which the memory error is larger than the measurement uncertainty is quite short, with delays less than $\sim 200\text{ms}$ having measured memory errors consistent with zero. This uncertainty could be reduced by collecting more data, but this approach would quickly become infeasible; the data presented here required several days of continuous experiments, and the uncertainty only decreases with the square root of the amount of data. Fortunately, it is possible to circumvent the SPAM error limitation with the approach explained in the following section.

5.2 Randomized benchmarking


To get around the limitation imposed on measurements by the SPAM error, we can include many delays between each state preparation and measurement rather than just one, amplifying the associated memory error relative to the SPAM error. To prevent the decoherence over each of these delays from combining coherently into that of a single, longer delay, we interleave the delays between random Clifford gates in an RBM sequence (see section 2.1). The fidelity of such an interleaved sequence is



(a)



(b)



(c)

Figure 5.4: The results of our Ramsey experiment measurement of the qubit memory error. (a) The raw data. (b) The SPAM error, which was measured simultaneously to the Ramsey fringe contrast. (c) The data after correction for the SPAM error.

compared to that of an equivalent sequence without interleaved delays to determine the average error caused by each delay.

While this approach is insensitive to SPAM error, it is now important to have a small gate error, and to ensure that this gate error does not change when interleaving delays – such a change in gate error would be incorporated into the inferred memory error, affecting the accuracy of the results. This section starts by giving a characterization of our gates, both the raw $\pi/2$ -gates and the composite gates used in the final experiments, and then moves on to present the results and analysis of our memory benchmarking experiments.

5.2.1 $\pi/2$ -gates

We use $\pi/2$ -gates around the $\pm x$ and $\pm y$ axes of the Bloch sphere to build our Clifford set. The optimal such composition (i.e. minimizing the number of $\pi/2$ -gates) has an average of 2.17 $\pi/2$ -gates per Clifford gate. Our RBM sequence randomly applies some number of gates from the Clifford set, with an additional Clifford gate at the end of the sequence to (in the absence of errors) move the qubit state into $|0\rangle$ or $|1\rangle$, with the final target state randomly chosen between these two possibilities.

To calibrate the duration of each $\pi/2$ -gate, we perform an RBM sequence with some number of Clifford gates, scanning the gate duration and measuring the change in sequence fidelity. Figure 5.5 shows the results of an example such calibration.

We find that the microwave power, and therefore the gate duration which corresponds to a $\pi/2$ -rotation, drifts significantly over a period of a few hours. Figure 5.6 shows this drift over an example period of time.

We determine the average error of each Clifford gate from the sequence fidelity using




Figure 5.5: The results of an example calibration of the $\pi/2$ -gate duration. Each sequence contains 2000 Clifford gates, with 50 different random sequences used in the calibration. The sequence fidelities measured for a range of gate durations are fit to a quadratic curve, with the fit giving a calibrated optimum gate duration of $9.615(3)\mu\text{s}$.

equation 2.3. We measure the SPAM error directly with measurements interleaved between each RBM experiment, with a delay instead of the Clifford gates to keep the ion heating the same and (when necessary) a single π -gate to give the same target state as for the corresponding RBM sequence (in a similar way to the experiments described in section 5.1).

Figure 5.7 shows the results of our fidelity measurement. The gate time is recalibrated every ~ 3 hours to correct for drifts in the microwave power. By fitting to equation 2.3 (with the SPAM error set to zero for the SPAM-corrected data), we calculate the average Clifford gate error to be $1.65(9)\times 10^{-6}$.




Figure 5.6: The drift of the calibrated $\pi/2$ -gate duration over an example period of time. The data is obtained by performing calibrations like that presented in figure 5.5, finding the gate duration which corresponds to a $\pi/2$ -rotation.

5.2.2 BB1 composite gates

The high gate fidelity from the measurement presented in figure 5.7 demonstrates the feasibility of overcoming our drifts in microwave power with frequent recalibration of the gate duration. However, the microwave power also changes significantly when interleaving delays between gates. This is because the duty cycle of the microwaves changes, causing temperature changes in all microwave components after the DDS.

There are two possible ways for correcting for this variation in microwave power. We could, in principle, recalibrate the gate time for each delay period in a similar way as presented in figure 5.5. However, the interleaved delays make these experiments take much longer to perform, so much so that this approach would be infeasible for many of the delay lengths for which we want to measure the memory error. We could instead measure the change in microwave power with a power-meter rather




Figure 5.7: The results of our gate error measurement. The fidelities are for RBM sequences with different numbers of Clifford gates, with SPAM measurements interleaved between each RBM sequence. 50 different random sequences of Clifford gates were used for each data point. The gate time was recalibrated every ~ 3 hours. Fitting the SPAM-corrected data to equation 2.3 with the SPAM error set to zero gives an average Clifford gate error of $1.65(9) \times 10^{-6}$.

than the ion, and convert these power measurements into corresponding $\pi/2$ -times. However, this approach would only give us an indirect measure of the microwave field at the ion, and so there is potential for some error here, which would lead to an over-estimate of the memory error.

Instead of attempting to correct for these changes in microwave power, we use BB1 composite gates in our RBM sequences. The theory behind these gates is described in section 2.2. They are unaffected by small changes in Rabi frequency, meaning that the gate fidelity is the same with and without delays interleaved between the gates. This robustness also has the advantage that there is no need to correct for the drifts in the microwave power illustrated in figure 5.6; we find that there is no

reduction in gate fidelity even several weeks after calibrating the duration of each constituent gate in the BB1 sequence.

For each logical gate, the BB1 sequence uses physical gates with three different phases, meaning that three different microwave profiles on our DDS are required. Our DDS only has eight profiles in total, which means that we can no longer apply $\pi/2$ -gates around the $\pm x$ and $\pm y$ axes in the Bloch sphere, but are instead limited to just the $+x$ and $+y$ axes. We use the same Clifford gate composition as in section 5.2.1 but with each $\pi/2$ -gate around $-x$ replaced by three $\pi/2$ -gates around $+x$, and the equivalent in y .

We also adjust our Clifford sequences to reduce the variation in length and thus result in a more consistent Clifford gate error and a more stable microwave duty cycle. We do this by adding sets of four logical $\pi/2$ -gates around $+x$ to some of the Clifford gates. The average number of logical $\pi/2$ -gates in each Clifford gate is now 3.58, with each of these consisting of several physical $\pi/2$ -, π - and 2π -gates as described in section 2.2.

Calibrating the duration of each BB1 gate now involves not just calibrating the duration of each $\pi/2$ -gate as shown in figure 5.5, but also calibrating the π - and 2π -gate durations. We perform the calibration of these two durations by implementing an RBM sequence with BB1 gates, scanning either the duration of each π -gate or that of each 2π -gate, setting it to the value that maximizes the sequence fidelity. We then iterate several times between calibrations of these two gate durations.

Figure 5.8 shows the results of a fidelity measurement with RBM sequences of BB1 gates, with a fit to the data giving an average BB1 gate error of $6(1)\times 10^{-6}$. We also measured the fidelity with normal $\pi/2$ -gates but the same logical construction of each Clifford gate as for the BB1 sequences (so an average of 3.58 $\pi/2$ -gates

per Clifford gate) to allow for a direct comparison. To demonstrate the robustness of the BB1 gates to variations in microwave amplitude, we repeated each fidelity measurement with a 1% reduction in microwave amplitude. While the average error of the normal $\pi/2$ -gates increased by more than two orders of magnitude, there was no significant increase in the BB1 gate error.




Figure 5.8: Fidelity measurements of RBM sequences with normal $\pi/2$ -gates and with BB1 composite gates, both with and without a 1% reduction in the microwave amplitude. The average gate errors r from fitting to equation 2.3 are given in the plot legend. The SPAM error is not measured independently but is a floating parameter in each fit. 50 different random sequences of Clifford gates are used for each data point.

5.2.3 Interleaving delays

To measure the average qubit memory error over delays of different lengths, we interleaved these delays in our RBM sequence, using BB1 composite gates, comparing the measured sequence fidelities to those without interleaved delays, as described in section 2.1 and illustrated in figure 5.9. For the sequences without interleaved delays, we included an additional delay just before state readout (when the qubit

state was $|0\rangle$ or $|1\rangle$, and so was insensitive to phase noise) equal in length to the total delay in each interleaved sequence; this was to ensure that the amount of ion heating before readout was the same as for sequences with interleaved delays. The experiments for each interleaved delay period were performed with some number m of Clifford gates, where the value of m was chosen for each delay length to give a total interleaved sequence infidelity of ~ 0.1 . This infidelity is much larger than that due to SPAM and Clifford gate errors and still much smaller than 0.5, at which point the qubit would completely decohere and the measurement would become insensitive to the precise value of the memory error.




Figure 5.9: Gate sequences for the RBM experiments with and without interleaved delays τ . The RBM sequence consists of m randomized Clifford gates C_i , with a final gate C_{m+1} to randomly rotate the qubit state into either $|0\rangle$ or $|1\rangle$. The sequences without interleaved delays instead include the same total delay at the end of the sequence, when the qubit state is insensitive to phase noise, to ensure that the amount of ion heating is the same.

Figure 5.10 shows the average memory errors for a range of delay lengths. The measured errors go down to $3(1) \times 10^{-6}$ at 2ms, at which point we hit the noise floor of the measurement due to the gate error. The graph also includes plots of several systematic sources of error and a fit to the data, which we now explain in more detail.

Before and after each ~ 4 hour period of data-taking, we measured the detuning between the qubit transition and the microwave radiation using the method described




Figure 5.10: Average memory errors for a range of delay lengths. Also plotted are several inferred systematic contributions to the measured errors, which are explained in the main text. The data is fit to a combination of white noise (see equation 2.14) and $1/f$ frequency noise with a low-frequency cutoff (see equation 2.17). There is a data point at 1ms with an inferred memory error that is negative, so it does not appear on this logarithmic plot, but it is included in the fit. The fit has a χ^2 p -value of 0.035.

in section 5.1. Based on these measurements, we made occasional adjustments (generally every few days) to the microwave frequency to correct for drifts. For delay periods of 1s and below, this approach enabled us to keep the detuning error much smaller than the memory error. We estimate the average detuning error for each data point by assuming a linear drift in detuning over the course of each period of data-taking. For measured detunings of δ_1 and δ_2 before and after the memory error measurement respectively, the average detuning error for delay τ is:

$$r_{det} = \frac{1}{2} \left(1 - \frac{1}{2\pi\tau(\delta_2 - \delta_1)} \left(\sin(2\pi\tau\delta_2) - \sin(2\pi\tau\delta_1) \right) \right). \quad (5.4)$$

For delay periods above 1s, the frequency resolution of the synthesizer (32.32mHz on the frequency of the radiation applied to the ion) is insufficient to keep the de-

tuning error much smaller than the memory error. To allow us to set the frequency with greater precision, we use it as a vernier with the DDS, which has a resolution of 1.844Hz on the radiation frequency. As illustrated in figure 5.11, this method enables us to set the radiation frequency with a resolution of 2mHz. The frequency drifts over the course of each period of data-taking, but the error caused by this drift is much smaller than the memory error (see figure 5.10).




Figure 5.11: An illustration of our vernier technique to improve the combined frequency resolution of the synthesizer and the DDS (not to scale). By changing the DDS frequency by one step in one direction and the synthesizer frequency by 57 steps in the other direction, we change the radiation frequency by 2mHz, which is a smaller amount than that which is achieved by changing only one of the two frequency settings.

As explained in section 3.5.2, the interleaving of delays between Clifford gates results in significant transients in microwave power over each of these gates. We estimate the resulting infidelity by measuring the fractional change in microwave power over the gates for each delay period, as explained in the caption for figure 3.7. These measurements are then used in numerical simulations of BB1 gates to estimate the resulting infidelity.

For most of the range of delays, the infidelity due to microwave power transients is insignificant compared to the memory error, but for delays of 2ms and below the power transient error is comparable to the memory error (see figure 5.10). We suppress the power transients for these delays by turning on the microwaves at the DDS some time before the end of each delay, with the microwaves prevented from reaching the trap solely by the switches (see figure 3.6 for an illustration of the different microwave components).

This method of reducing the microwave power transients introduces a new systematic error source. For the period of time when the DDS is on but the switches are off, there is significant microwave leakage through the switches. We consider the errors caused by both microwave leakage and power transients in our choice of times between turning the DDS and switches on, roughly minimizing the combined systematic error.

As shown in figure 5.10, the systematic error sources of frequency detuning, microwave power transients and microwave leakage are, for delays longer than $\sim 2\text{ms}$, much smaller than the measured memory error. The data shows that the memory error does not decay with a simple exponential shape of white noise (see equation 2.14), as is commonly assumed in trapped-ion quantum computing research.

One possible noise distribution is a combination of white noise and $1/f$ noise with some cutoff frequency (see equation 2.17). This is the distribution used in the fit in figure 5.10, with the fit having a χ^2 p -value of 0.035.

We found in section 4.1 that our magnetic field noise is dominated by low-frequency ($f \lesssim 1\text{kHz}$) noise, so this could have a $1/f$ distribution. At our target static magnetic field of 146.0942G , the qubit transition frequency is first-order independent of magnetic field, but an offset in the field would increase the field sensitivity. The variance of the qubit detuning can be equated to that of the qubit phase by the following equation [55]:

$$\langle \delta^2(\tau) \rangle = \frac{\langle \phi^2(\tau) \rangle}{\tau^2}. \quad (5.5)$$

From the precision of our field-independent point measurements (see section 4.7), we can be confident that the field offset is no larger than $100\mu\text{G}$. Taking the $1/f$

component of the noise model from the fit of figure 5.10 with τ set to $700\mu\text{s}$, the time constant of our field-sensitive transition Ramsey experiment in figure 4.4, and with a magnetic field offset of $100\mu\text{G}$, corresponds to a RMS magnetic field noise of 130mG , much larger than the measured value of $140\mu\text{G}$. The approximation that the transition frequency changes linearly with the field, with the proportionality constant evaluated at some constant field offset, is a poor one, and so the result of this calculation has limited significance. However, the fact that the calculated field noise required to explain the results of figure 5.10 is three orders of magnitude larger than the measured value suggests that magnetic field noise is not the dominant source of qubit memory error.

It is unclear what the dominant source of memory error is in our system. It could be noise introduced in the microwave chain, with key components being the rubidium frequency reference, the synthesizer and the DDS. It could be noise on the qubit transition frequency, which could be caused by noise in the AC Zeeman shift induced by RF currents in the microwave electrodes. What is clear, however, is that the commonly-assumed model of exponential decay does not apply in our system, with the error at short times being much smaller than would be inferred by fitting data at long times to an exponential model and extrapolating this model.

To assess whether the qubit memory errors measured in our system are likely to be a limiting factor in the development of a quantum computer, we can compare our measurements to the errors in two-qubit gates that have been demonstrated in ion trap experiments. The lowest such errors achieved to date are $\sim 1 \times 10^{-3}$ [34, 35, 36], with gate times $\sim 100\mu\text{s}$ with lasers and 3.3ms in our experiments with near-field microwaves. Assuming that delays between gates in computational sequences are similar to the two-qubit gate time, it is the memory error over these times that is of relevance for quantum computing. Our measured memory errors on the timescale of milliseconds is $< 10^{-5}$, indicating that memory error is not currently a limiting

factor for trapped-ion qubits of this type.

While the precise values of qubit memory error presented here are specific to our system, their deviation from the commonly-assumed model of exponential decay indicates the importance of directly probing memory errors at the millisecond timescales that are relevant to quantum computing. We have demonstrated that the simple metric of “coherence time”, which is commonly presented in research papers, can be highly dependent upon the range of times over which decoherence is measured, and therefore it is not generally a good figure of merit when assessing the performance of any given experimental system.

5.3 Decoherence suppression

We have demonstrated that the coherence time of a qubit extracted from decoherence measurements over long delays is of limited use when assessing the suitability of the qubit for quantum computing applications. However, experiments to date have generally used only long delays, largely due to relatively large SPAM errors as explained in previous sections. Therefore, to compare our system with others, it is necessary to also perform experiments over long delays.

In this chapter, we have already presented measurements of qubit decoherence over time periods as long as 20s. It is not possible to go much further than this as measurements over longer periods would quickly show complete loss of coherence.

It is, however, possible to preserve qubit coherence over longer periods with dynamical decoupling pulses. Such pulses are designed to counter random rotations of the qubit state around the Bloch sphere. By applying gates at intervals much shorter than the time over which the random rotations change in speed and/or rotation angle, it is possible to average the individual rotations to zero.

Perhaps the simplest form of dynamical decoupling is the insertion of “spin echo” pulses. These pulses rotate the qubit state by 180° about some axis in the x - y plane of the Bloch sphere. By inserting such a pulse in the middle of some delay period, where there is some unknown, but constant, rate of rotation around the z -axis, the rotations before and after the spin echo pulse cancel out to give zero net rotation. This method can be extended to a varying rotation rate by simply inserting more spin echo pulses.

One issue with such sequences is that errors in the spin echo pulses result in additional qubit decoherence. While our single-qubit gates have a very low average error immediately after calibrating the gate time (see figure 5.7), drifts in the microwave power (see figure 5.6) result in some under- or over-rotation angle θ in each gate. In a spin echo sequence consisting of a series of n π -pulses in the same direction, the resulting qubit state infidelity scales quadratically with the number of gates:

$$1 - F = \sin^2\left(\frac{n\theta}{2}\right) \approx \frac{n^2\theta^2}{4}. \quad (5.6)$$

To remove this source of error, we can simply alternate the direction of the spin echo rotations. This means that under- or over-rotations in the pulses that vary over a timescale much slower than the period between each pulse are completely cancelled out.

However, this introduces another source of error. Physically speaking, reversing the rotation direction of a gate does not involve rotating around the same axis in the opposite direction. It involves rotating in the same direction, but around an axis that is rotated by 180° about the z -axis. If there is some error ψ in the relative angle of the two rotation axes, then with each pair of spin echo pulses the qubit state will

walk away from its initial state, with the following state infidelity after n pulses:

$$1 - F = \sin^2 \left(\frac{n\psi}{2} \right) \approx \frac{n^2\psi^2}{4}. \quad (5.7)$$

Figure 5.12 illustrates this effect on the Bloch sphere.




Figure 5.12: An illustration on the Bloch sphere of alternating the rotation axis of spin echo pulses. The initial state, and the target rotation axis for the spin echo pulses, are on the x -axis. A slight error in the relative rotation angle results in the qubit state walking away from the initial state.

To remove this source of error, we instead perform two π -pulses about one axis, followed by two π -pulses about a second axis that is nominally rotated by 180° from the first. Our spin echo sequence consists of repeating this set of four π -pulses.

The delay between spin echo pulses should be as small as possible, as this allows for faster phase noise to be cancelled out by the spin echo sequence. However, the delay should not be so short that the error in each pulse becomes the dominant source of qubit decoherence. Guided by our results from section 5.2.1 and section 5.2.3, we choose a delay of 10ms.

We apply our spin echo sequence in Ramsey experiments performed as in section 5.1. Figure 5.13 shows the result of these experiments for total delay times ranging from 10s to 60s. Fitting the data to an exponential coherence decay curve gives an effective coherence time of 1000(500)s.




Figure 5.13: The results of our spin echo dynamical decoupling Ramsey experiments. The axis of rotation for the spin echo pulses was the same as the target qubit state in the x - y plane of the Bloch sphere. The data is fit to an exponential decay curve, giving an effective coherence time of 1000(500)s. The SPAM error was measured simultaneously to the Ramsey experiments and subtracted from the raw data.

It should be highlighted that in these experiments, the rotation axis for the spin echo pulses was the same as the target qubit state. This meant that the change in qubit state induced by each pulse was minimal. If the qubit state were elsewhere in the x - y plane of the Bloch sphere, then there would be an increased gate error, but this could be avoided by simply changing the spin echo pulse rotation axes to match the target qubit state. However, if the qubit state were not in the x - y plane then such an adjustment would not be possible, and, while the error caused by phase noise would be smaller, the error caused by gate error would likely be larger.

Chapter 6

Next-generation trap design

In this chapter, we present a design study of a next-generation ion trap for near-field microwave quantum gates. The rest of the apparatus could be largely the same as for the experiments presented in this thesis, but the new trap electrode geometry should allow for faster and higher fidelity two-qubit gates than we have been able to achieve with our existing trap [36].

A key goal for high fidelity near-field microwave two-qubit gates is to maximize the ratio of the microwave magnetic field gradient to the microwave magnetic field. The field gradient is responsible for coupling the internal and motional states of the ions (the precursor to generating entanglement between ions) and the field drives single-qubit operations. Interactions with each individual qubit interfere with two-qubit gates; ideally, one should be able to apply fields or field gradients alone to perform single- or two-qubit gates.

The first experiment on near-field microwave two-qubit gates [31] achieved this by driving three separate linear electrodes with microwave currents. By careful control of the relative amplitudes and phases of these currents, the fields produced by each electrode (which were all in the same plane, as the electrodes were along the same

axis) were made to roughly cancel at the ions' locations, as illustrated in figure 6.1. However, to implement such control is challenging – drifts of the relative amplitudes or phases of the three sources need to be corrected for on a regular basis.




Figure 6.1: An illustration of how three electrodes carrying electric currents can be used to generate a magnetic field distribution that, at a particular point, has zero field and a non-zero field gradient.

In our previous work [36], we instead opted to drive a single linear electrode which produces fields of a different polarization to our qubit transition. Due to the symmetry of our trap, the central electrode produces a field that is almost entirely π -polarized, which does not directly couple to our σ -polarized qubit transition. However, the field does cause off-resonant excitation of other transitions which are π -polarized, resulting in a significant gate error.

Another option is to combine the advantages of these two approaches: have a single electrode to avoid any strict requirements on microwave source amplitude and phase stability, but with multiple linear elements to cancel the field at the ions' locations. As first proposed in [46], this can be achieved with a meander electrode which turns

back on itself twice to make an ‘S’ shape (see figure 6.4 for an illustration of this). For any microwave amplitude or phase, the resulting microwave field is roughly zero at the same fixed location in the x - y plane (or along the same line in the z -direction, where multiple ions can be positioned along this line).

The challenge with this approach is in getting the microwave field cancellation to occur in the right place. While a Paul trap can confine ions in a range of locations by adjustment of the DC electrode voltages, the shape of the RF field is fixed by the electrode geometry¹, which produces a null in the RF field at a single location in the x - y plane. Using the DC field to shift ions away from this null results in RF micromotion, where the non-zero RF field causes ions to oscillate at the RF frequency. This oscillation generates sidebands in radiation (as experienced by the ions), causing significant errors in operations. For ions to experience neither RF micromotion nor microwave fields, it is necessary to ensure that the nulls of the RF and microwave fields are in the same place.

Due to their low frequency, RF voltages can be simulated in the same way as DC voltages, meaning that it is easy to accurately predict where the RF field null will be for any given electrode geometry. The same cannot be said of microwave currents. The higher frequency means that the current crowds towards the edges of the electrodes (the “skin effect”) in a way that depends upon the exact electrode shapes. While simple electrode structures, such as coplanar waveguides, can be treated analytically, the sort of meander shape necessary for our application must be modelled numerically. We use the simulator HFSS from Ansys, which calculates the fields around electrode structures using the finite-element method. Most of our simulations are performed on a cluster computer run by Advanced Research Computing at the University of Oxford².

¹One could, in principle, drive three or more RF electrodes with sources of different powers and phases, but this would impose strict requirements on the relative amplitude and phase stability of those sources.

²<http://www.arc.ox.ac.uk/>

While microwave fields can be simulated with software such as HFSS, due to the simulation complexity it is not trivial to ensure that the simulation results are accurate while keeping the computational requirements at a manageable level. In addition, the results can be heavily dependent on the boundary conditions of each electrode. Microwave currents couple from the input electrode onto others in the trap structure, and, depending on what is attached to the input of those electrodes, the microwave currents reflect back into the trap with some amplitude and phase, generating a secondary microwave field that can cause a significant shift in the location of the field minimum.

A meander-electrode ion trap has been designed and tested in another group [47, 48]. The simulations done during the design work for this trap predicted a separation between the RF and microwave field nulls of $0.9\mu\text{m}$, and, once the trap was fabricated and tested, the separation was measured to be $3.8\mu\text{m}$.

In this chapter, we work towards designing a new meander-electrode trap with a reduced distance between the RF and microwave field nulls to allow for high-fidelity entanglement generation. We present a new geometry which should reduce the microwave coupling onto other electrodes, which was one of the most significant problems in previous work [47]. We also develop a design process which provides electrode widths that correspond to an overlap of the RF and microwave field null locations in the simulations. In addition to the consideration of these field null locations, we also consider the RF and DC voltages necessary for trapping, transporting and splitting ion chains, as well as the trap depth, choosing electrode widths that result in reasonable values for each of these parameters. We investigate the accuracy of our simulations by looking at how the simulated field null locations change with increased model resolution, up to the maximum that we can achieve with available computational resources. We also consider how errors in the trap fabrication and

changes in trap temperature would change the field distributions.

6.1 Preliminary decisions

Before giving further consideration to the shapes of the trap electrodes, we make some preliminary decisions about the trap. One of the limiting factors in previous work on near-field microwave entanglement generation is the long gate time, as ion heating during the gate results in a significant gate error. This time is limited by the size of the magnetic field gradient at the ions' locations, which can be increased by bringing the ions closer to the trap surface, scaling with the ion height d as $\sim 1/d^2$. However, the ion heating rate would then also increase, scaling as $\sim 1/d^4$ [70, 71]. We plan for the trap to be cooled with a liquid helium cryostat to $\sim 20\text{K}$, which should reduce the ion heating rate by more than an order of magnitude [76, 77]. This would allow the ion to be brought closer to the trap surface. We choose an ion height of $40\mu\text{m}$, around half of that for our current trap ($75\mu\text{m}$).

We choose gold as the electrode material, which is widely used in surface ion traps. We choose sapphire for the trap substrate material because of its high thermal conductivity at 20K – see figure 6.2 for a comparison with other relevant materials. This is important for our application as a limitation on two-qubit gate speed (and, therefore, gate fidelity) is the maximum microwave power that can be dissipated in the trap before thermal transients have a significant effect on experiments. We simulate for a substrate thickness of $430\mu\text{m}$, as this is a standard thickness for high-quality sapphire wafers. Sapphire is a uniaxial crystal with one relative permittivity parallel to the c -axis and another perpendicular to it. We choose the c -axis to be in the vertical (y) direction, as this means that the substrate orientation is unimportant, simplifying the fabrication process.

For the electrode thickness, there are two competing factors. Thicker electrodes




Figure 6.2: The temperature dependence of the thermal conductivities of several materials used in microfabrication. Figure adapted from [57], with data taken from [78].

spread the microwave currents out over a greater surface area³, resulting in lower electrode temperatures and therefore lower losses. However, fabrication becomes more challenging with increased electrode thickness. We choose $8\mu\text{m}$ for the thickness, which should be manageable with existing fabrication techniques.

The widths of the electrodes are the main parameters that we vary to achieve our design goals, but we must initially choose a minimum possible width, as fabrication becomes more challenging for smaller electrode widths. We choose a minimum width of $10\mu\text{m}$, which has been achieved in several previous surface traps without difficulty.

³The skin effect means that most of the microwave current is confined to the electrode surfaces.

There are several considerations for our choice of the gap width between electrodes. Exposed areas of substrate can collect charged particles over time, and, due to the high electrical resistivity of the substrate, this charge remains on the surface for long periods, causing drifts in the DC field. To minimize this effect, gaps between electrodes should be as thin as possible. However, if the gaps around the RF electrodes are too thin, electrical breakdown can occur, causing damage to the trap. Additionally, due to imperfections in fabrication, the actual gap may be thinner than the designed width, resulting in RF breakdown or even shorts between electrodes. In line with previous work, we choose an electrode gap width of $5\mu\text{m}$. In areas of the trap further away from the ions, where exposed dielectric has less effect on the trapping field, the electrode gaps could be made wider to improve the fabrication yield.

To ensure that our simulations are performed with the correct microwave frequency, we also need to choose the qubit transition. We stick with $^{43}\text{Ca}^+$ as the ion, using the $|F = 4, M = +1\rangle$ to $|F = 3, M = +1\rangle$ transition in the $4S_{1/2}$ ground level, operating at a magnetic field of 288G (see section 2.4.1 for an explanation of the reasoning behind this choice). This means that our microwave frequency for simulations is 3.12GHz (see table A.2).

Next, we consider some parameters of the trapping field. These could be changed after trap fabrication by changing the DC and RF voltages. However, the electrode geometry may mean that the desired trapping parameters could only be achieved with voltages that are too large to produce with standard equipment. Therefore, it is important to think about them during the design process.

It is desirable to have the Doppler cooling laser beams in the plane of the trap because blue light (such as the beams at 397nm) would cause charging of exposed areas of dielectric material. Also, as we use the Doppler beams for readout, scatter of

397nm light from the trap surface would increase the background count of the PMT, reducing the fidelity of our state-dependent fluorescence detection readout scheme (see section 2.4.6). To achieve effective cooling of both radial modes with beams parallel to the trap plane, each radial trap axis must have a sufficiently large projection in this direction. This can be achieved with careful choices of voltages on the DC control electrodes⁴. We choose 15° to be the minimum angle between the radial trap axes and the trap surface to achieve effective cooling of all motional modes.

In our previous work on two-qubit gates [36], we used one of the radial modes for entanglement generation, as it is not possible to generate a microwave field gradient along the trap axis with microwave electrodes that lie along this axis. We plan to also use a radial mode with the new trap for the same reason. An important quantity for entanglement generation is the initial temperature of the utilized mode, measured as a mean number of vibrational quanta, which decreases with increasing mode frequency. Our existing setup has a frequency of 3.3MHz; to enable us to substantially reduce the gate error compared to our previous experiments, we want the new trap to have a significantly larger mode frequency. We choose 8MHz as the minimum frequency.

For each radial trap axis, a pair of trapped ions has two vibrational modes: centre-of-mass (COM) and rocking. For the COM mode, the ions oscillate in phase with one another; for the rocking mode, out of phase. The rocking mode heating rates are generally much lower than the COM mode heating rates. This is because ion heating is mostly caused by electric field noise that is correlated between ions, i.e. they generally get pushed in the same direction, not the opposite direction. For this reason, we used the rocking mode for entanglement generation in our previous work, and we plan to do the same with our next trap.

⁴In our current trap, we achieve this by applying asymmetric DC voltages to the microwave electrodes (see section 3.1.1). For a meander electrode geometry, where the several electrodes in the trapping region are replaced by a single one that bends around, this approach is not possible.

The radial mode rocking mode frequency $\omega_{r,rock}$ is related to the radial COM mode frequency $\omega_{r,COM}$ and the axial COM mode frequency $\omega_{a,COM}$ by the following equation:

$$\omega_{r,rock} = \sqrt{\omega_{r,COM}^2 - \omega_{a,COM}^2}. \quad (6.1)$$

To use one of the radial rocking modes to generate entanglement without interacting with the COM mode, there must be sufficient splitting between the two mode frequencies. As equation 6.1 shows, this splitting increases with the axial COM mode frequency. With this in mind, we choose a minimum axial COM mode frequency of 1MHz.

Finally, we set the stability parameter q of the trap to be 0.3, which is the value used in our current trap.

6.2 Choice of electrode geometry

We have already specified that the trap design developed in this design study will feature an ‘S’-shaped meander electrode for carrying microwave currents, but this allows for many different geometries for the RF and ground electrodes. In this section, we consider several such geometries and decide upon one for this trap design.

We follow the vast majority of previous work on Paul traps in having two RF electrodes, the minimum number required for trapping ions. It is desirable to have these two electrodes connected to one another as close to the trapping region as possible. This is because any difference in electrical length to this region would result in a phase difference between the RF voltages on each electrode, causing uncompensatable micromotion – a phase difference of just 1° could result in micromotion with an

equivalent temperature of $\sim 1\text{K}$ [79]. We make the decision that the RF electrodes will be connected at the end of the meander electrode.

In our current trap, we generate a high microwave current in the trapping region by extending the microwave electrodes by a quarter-wavelength and then leaving them open at the end. This means that the trap must be relatively long, but allows DC and RF voltages to also be applied to the microwave electrodes. Since we will only be applying microwave signals to the meander electrode, we opt instead for shorting the electrode to ground at the end of the meander. The design is therefore restricted in that there must be a connection from the end of the meander to the trap ground plane.

The location in the x - y plane of the RF field null (with x being horizontal and y being vertical) is approximately governed by the House equations [80]:

$$x_{RF} = \frac{w_c w_{RF,l}}{w_{RF,u} + w_{RF,l}}, \quad y_{RF} = \frac{\sqrt{w_c w_{RF,u} w_{RF,l} (w_c + w_{RF,u} + w_{RF,l})}}{w_{RF,u} + w_{RF,l}}, \quad (6.2)$$

where w_c is the distance between the two RF electrodes, $w_{RF,u}$ is the width of the RF electrode in positive x , and $w_{RF,l}$ of that in negative x , as illustrated in figure 6.3. The origin is defined to be between the central electrode and the RF electrode in negative x , on the top surface of the electrodes. The “gapless plane” approximation has been made, which means that the gaps between electrodes have been ignored in the definitions of the electrode edges. We make this approximation from now on, incorporating the $5\mu\text{m}$ gaps into the stated values of electrode widths (so we now say, for example, that each electrode must have a minimum width of $15\mu\text{m}$).

Solving the House equations for our specified ion height of $40\mu\text{m}$ and minimum electrode width of $15\mu\text{m}$, we find that the design is restricted to have a maximum of




Figure 6.3: A basic illustration of a generic surface Paul trap showing the electrode widths used in equation 6.2. The red electrodes carry RF voltages and the black electrodes are at RF ground, but they may carry DC or microwave signals. The black electrodes may be split into multiple electrodes with little effect on the RF field.

five electrodes between the two RF electrodes. Any more than this would require at least one of the RF electrodes to have a width below $15\mu\text{m}$. However, having five electrodes would impose restrictions on the possible locations for the RF field null, as well as giving very little flexibility on the electrode widths. Therefore, we restrict our geometry to have at most four electrodes between the RF electrodes.

Figure 6.4 shows the electrode geometries that satisfy our design restrictions and are investigated with full electromagnetic simulations in HFSS. We now consider each in turn and justify our choice of geometry (g).

Geometries (a) and (b) have the entire meander electrode enclosed between the two RF electrodes, with an extra electrode on one side to allow currents to return to




Figure 6.4: Schematics of the trap electrode geometries that we investigate with simulations in HFSS. The blue meander electrode carries the microwave currents, the red electrodes carry RF voltages, and the black electrodes are connected to the ground plane. The grey electrode is connected to the ground plane further to the left of the figure, 2mm from the trap centre, with a $5\mu\text{m}$ gap separating it from the ground plane in the trapping region. The meander and RF electrodes are connected to their sources to the left and right of the diagrams respectively. The design also includes electrodes further away for applying microwave fields for single-qubit gates and DC fields for trapping. (d) is similar to the geometry used in [47, 48] and (g) is the geometry chosen in this design study.

the ground plane. This means that the widths of the RF electrodes could be varied with little effect on the microwave field, potentially making it easier to overlap the microwave and RF field nulls. However, we find that such geometries would lead to the microwave field minimum being at a height of $\sim 10\mu\text{m}$, much smaller than our

target of $40\mu\text{m}$, so we rule them out.

Geometry (c) would allow return currents to flow on either side of the meander electrode. This would reduce the sensitivity of the field distribution to the properties of what the RF electrode is connected to off the trap chip. However, we find that it is not possible to achieve overlap of the RF and microwave field nulls with this geometry for our design rules.

Geometry (d), which is similar to the geometry used in [47, 48], has one of the RF electrodes between two of the meander segments, with a ground electrode on the opposite side. We find that it is not possible to align the RF and microwave field nulls by varying the electrode widths of this geometry, with the microwave field null too far in the positive x direction, so we rule it out. One difference of the geometry in [48] that may have enabled their rough alignment of the RF and microwave field nulls is the properties of the electrodes on either side of the meander; one, as for geometry (d), was an RF electrode, while the other was a second microwave electrode (which must be included somewhere for applying fields for single-qubit operations). The coupling of microwave currents to this electrode would have been different than for the ground electrode in geometry (d), changing the microwave field distribution. We could also have our second microwave electrode right next to the meander in an attempt to obtain similar behaviour, but we expect that the resulting strong coupling between the electrodes, and the corresponding uncertainty in the location of the microwave field minimum, would harm the performance of the trap.

To achieve overlap of the RF and microwave field nulls, the geometry must be changed to pull the microwave null in the negative x direction. One such change that achieves this is to remove the ground electrode between the lower meander bends, as in geometry (e). This change makes the lower bend more important for the microwave field in the trapping region, bringing the microwave field minimum

closer to it. However, we find that this change is still insufficient to overlap the two field nulls.

In an attempt to increase this effect, geometry (f) has an additional ground electrode between the upper RF electrode and the upper meander bend. However, this change actually increases the distance between the RF and microwave field nulls.

Finally, in geometry (g), we introduce a cut in the ground plane above the upper meander bend. This cut consists of a $5\mu\text{m}$ gap in the ground plane, extending $15\mu\text{m}$ in the positive x direction from the end of the meander electrode, and then in the negative z direction to a point 2mm away from the trap centre (we find that the exact dimensions of this cut do not have a significant effect on the fields, so they are fixed at these values). At this point, the cut ends, with the electrode to either side being a single, unbroken ground plane. However, the cut suppresses the microwave return currents immediately above the upper meander bend, so they instead flow around the cut, further away from the trapping region. This change has the effect of partially restoring the symmetry of the geometry in the x direction, moving the microwave field minimum in the negative x direction, as the microwave return currents immediately below the meander electrode are likewise suppressed by the presence of the lower RF electrode. We find that this change is sufficient to allow for the RF and microwave field nulls to be overlapped, and geometry (g) is used for the trap design.

Figure 6.5 shows this geometry with the second microwave electrode and DC electrodes, and figure 6.6 shows it as implemented in HFSS. The electrodes for producing DC trapping fields and microwave fields for single-qubit gates are relatively far away from the trapping region and separated from the other electrodes by areas of ground plane, so the details of their shapes have little impact on the location of the microwave and RF field nulls and they are only included in a crude fashion. All DC

electrodes on either side are combined together into single 1mm-wide electrodes. The single-qubit microwave electrode is set to be adjacent to the upper DC electrode, $15\mu\text{m}$ wide and $75\mu\text{m}$ away from the cut in the ground plane near the meander electrode. There is a $15\mu\text{m}$ -wide ground electrode between the lower RF and DC electrodes.




Figure 6.5: Schematic of the trap electrode geometry including the second microwave electrode and DC electrodes (not to scale). For simplicity, the two sets of DC electrodes are each combined into single electrodes. The blue electrode carries the microwave currents for two-qubit gates, the red electrodes carry RF voltages, the black electrodes are connected to the ground plane, the green electrode carries the microwave currents for single-qubit gates, and the magenta electrodes carry DC voltages. The grey electrode is connected to the ground plane further to the left of the figure, 2mm from the trap centre, with a $5\mu\text{m}$ gap separating it from the ground plane in the trapping region. The microwave and RF electrodes are connected to their sources to the left and right of the diagrams respectively, with the DC electrodes connected to the top and bottom.

There are two factors to consider in the choice of meander length. We will want to fabricate several trap chips, as there will be some probability of fabrication errors rendering any particular trap unusable (for example, different electrodes could be shorted together), and it would be useful to have one or more extra traps for performing preliminary tests before assembling the final trap apparatus. Therefore, they will be fabricated in batches of small chips on larger sapphire wafers. To fit as




Figure 6.6: (a) The central region of the ion trap design as implemented in HFSS. Each bend of the meander electrode is 1.5mm long. The electrode immediately above the upper bend of the meander is connected to the ground plane to the left of the area shown, 2mm away from the trap centre. For simplicity, the two sets of DC electrodes are each combined into single 1mm-wide electrodes. (b) A close up of the right end of the meander with each electrode labelled by the names by which they are referred to in this thesis. Also denoted are the quantities w_c , the distance between the two RF electrodes, and x_{RF} , the horizontal distance between the lower RF electrode and the ions' positions.

many chips as possible on the wafer, the chip, and therefore the meander electrode, should be as small as possible. However, we want to have three trapping regions to allow complex operations to be performed by transporting ions and performing single- and two-qubit gates on ions in any of the trapping regions. The meander and RF electrodes should be sufficiently short to ensure that end effects are reason-

ably small in the outer trapping regions. We choose 1.5mm for the length of each meander bend as a compromise between these competing factors.

6.3 Choice of electrode widths

For the chosen electrode geometry, there are five electrodes whose width have a significant effect on the microwave and/or RF fields: each of the three bends of the meander, and each of the two RF electrodes. There is also the constraint that the RF field null must be $40\mu\text{m}$ above the electrode surface; we use equation 6.2 to choose RF electrode widths that fulfil this constraint. For each set of electrode widths, we then make a more accurate determination of the RF field null location using HFSS, which deviates from the approximate location by up to a few hundred nanometres.

We scan the different electrode widths to find sets of widths that result in overlapping RF and microwave field nulls. We parametrize the four degrees of freedom with the following four quantities: the width of the centre meander bend, how much wider each of the upper and lower meander bends are than the centre, and the position of the RF field null in the horizontal (x) direction as calculated with equation 6.2. This last parameter, along with the widths of the centre and lower meander bends, fully constrains the widths of the two RF electrodes for the chosen ion height of $40\mu\text{m}$. We express this parameter as a fraction of the combined widths of the centre and lower meander bends (the RF field null must be between the two RF electrodes).

Overlapping the RF and microwave field nulls is the primary objective of the parameter scan, but there are several other factors that will impact the performance of the trap. For the microwave magnetic field, the amplitudes of the field and the field gradient are important, but we find that these quantities do not vary significantly when scanning the different electrode widths.

Other important quantities, which vary significantly with the electrode widths, are the DC and RF voltages required to trap ions, the DC voltages needed to split chains of ions, and the depth of the trapping field. When considering these quantities, which we evaluate analytically under the gapless-plane approximation, we apply our design rules as presented in section 6.1. For the DC voltages, we choose the set of voltages that minimizes the largest voltage of the set, and only consider this largest voltage in our choice of electrode widths.

The largest DC voltage is strongly dependent on the distance from the ion to the lower DC electrodes, so it is important to minimize this distance. We include a single ground electrode between the lower RF and DC electrodes, with the minimum possible width of $15\mu\text{m}$, to allow for microwave return currents to flow here. The DC voltages are less sensitive to the exact distance between the upper RF and DC electrodes, which includes the upper meander bend, two ground electrodes and the single-qubit microwave electrode; this is set to $100\mu\text{m}$ for these calculations.

Splitting a pair of ions involves using both quadratic and quartic potentials along the trap axis to raise a potential barrier between the ions, with each ending up in its own potential well at the end of the process. The difficult part of this process is producing a sufficiently large quartic potential to minimize both the resulting amount of ion heating and the splitting time [81]. Palmero *et al.* [82] have designed a protocol for achieving fast separation with little ion heating, and provide a table of values for the maximum required quartic potential coefficient and the corresponding axial trap frequency that can be maintained. The values for the largest DC voltage that we give here are those required to maintain an axial trap frequency of 0.8MHz; a lower frequency would result in a longer required splitting time to avoid significant ion heating. As illustrated in figure 6.7, we set the width of each DC electrode to be $150\mu\text{m}$, with the central electrode in each of the upper and lower sets split into three

thinner electrodes with fractions $[1/6 \ 2/3 \ 1/6]$ of the total width – these values are chosen to roughly minimize the largest DC voltage required for ion chain splitting.




Figure 6.7: A diagram of the DC electrode segmentation to either side of the trapping region. There are nine $150\mu\text{m}$ -wide electrodes to allow for three trapping zones, with the central electrode further segmented into fractions $[1/6 \ 2/3 \ 1/6]$ to allow for the splitting of ion chains in the central trapping zone.

For the transport of ions, the second order potential derivatives that can be generated along the trap axis must be significant at every position along that axis in order to achieve transport with minimal heating [80]. We find that this is achievable for our chosen DC electrode width of $150\mu\text{m}$ with DC voltages similar to those required for trapping, so we do not consider it any further, simply noting that if our requirements on the trapping field are met then so are our requirements on the fields needed for transport.

Figure 6.8 shows how different quantities of interest vary with the distance w_c between the two RF electrodes and the horizontal location of the RF field null x_{RF} as a fraction of w_c (measured from the upper edge of the lower RF electrode). The range of values for w_c is restricted by our design rules. x_{RF}/w_c cannot be larger than ~ 0.7 without the voltages required for splitting ion chains becoming infeasibly large. This is because as the location of the field minimum moves towards the upper RF electrode, the lower RF electrode increases in size, pushing the lower DC electrodes away from the ion. We find that x_{RF}/w_c cannot be smaller than ~ 0.6

with our restriction on the minimum electrode width.




Figure 6.8: Plots of how the trap depth and the RF and DC voltages required for trapping, as well as the largest DC voltages required to split pairs of ions, vary with the distance w_c between the two RF electrodes and the horizontal location of the RF field null x_{RF} as measured from the upper edge of the lower RF electrode. The radial and axial mode frequencies, radial trap axis directions, and stability parameter q are as chosen in section 6.1.

As figure 6.8 indicates, to minimize the required DC voltages the ion should be as close as possible to the lower RF electrode, as this reduces the required width of the lower RF electrode, bringing the lower DC electrodes closer to the ions. We find that this can be achieved by maximizing the difference between the widths of the centre and lower bends of the meander electrode, with the centre being wider

than the lower. However, this difference cannot be much more than $10\mu\text{m}$ before it becomes impossible to overlap the microwave and RF field nulls. Therefore, we set the difference to be $10\mu\text{m}$, which gives $x_{RF}/w_c \sim 0.6$.

Figure 6.8 shows that, for the range of parameters considered here, the RF voltage and trap depth change significantly only with w_c , not x_{RF}/w_c , and that a smaller such distance results in a better trap. However, to maintain an ion height of $40\mu\text{m}$ for a smaller w_c , the RF electrodes must be made wider, so much so that decreasing w_c results in the total distance between the two sets of DC electrodes increasing. This means that higher DC voltages are required for trapping ions and splitting ion chains. For this reason, there is a trade-off between the different parameters for our choice of w_c .

We choose $w_c = 50\mu\text{m}$ to compromise between the different factors. The required RF voltage amplitude is around 90V , which is comparable to the value of 72V for our current trap. The largest required DC trapping voltage is well below 10V , and so should be easy to produce with standard DACs. For splitting a pair of ions whilst maintaining an axial mode frequency of 0.8MHz , the largest required DC voltage is around 15V , which would be somewhat inconvenient to produce, so it may be necessary to allow the mode frequency to drop below 0.8MHz and have a correspondingly longer splitting time [82].

We find that the distance between the RF electrodes for which the microwave and RF field nulls are aligned can be adjusted by changing the difference in width between the upper and centre meander electrode bends, where a larger upper width results in a smaller distance between the RF electrodes, again scanning the centre electrode width to align the two field nulls. Our choice of a distance of $50\mu\text{m}$ between the two RF electrodes means that the upper meander bend should be $20\mu\text{m}$ wider than the centre bend.

Figure 6.9 shows the results of the parameter scan for our choices of the widths of the upper and lower bends of the meander electrode. The two remaining parameters are the width of the centre bend and x_{RF} . A linear interpolation of the parameters between the simulation data points gives the electrode widths in table 6.1. The corresponding magnetic field parameters are given in table 6.2.

Electrode	Width (μm)
Lower bend of meander	19.36
Centre bend of meander	29.36
Upper bend of meander	49.36
Lower RF	56.45
Upper RF	34.41

Table 6.1: Table of the optimized electrode widths for the design.

Parameter	Value
Maximum gradient	14.9Tm^{-1}
Projected gradient	11.5Tm^{-1}
Residual field	$10.3\mu\text{T}$

Table 6.2: Table of the microwave magnetic field parameters for the electrode widths in table 6.1. The field gradient is given in two directions, that in which it is maximal, and that of the radial trap mode in which it has the largest projection; the latter of these quantities is what determines the two-qubit gate speed. The input power is 1W.

6.4 Simulation accuracy

In this section, we consider possible sources of inaccuracy in the results of our simulations. These could be due to computational resources being insufficient to reach the required level of precision, or due to the HFSS model not being a true represen-




Figure 6.9: (a) Plots of the mismatch between the microwave and RF field minima in the horizontal (x) and vertical (y) directions as a function of x_{RF}/w_c and the width of the centre bend of the meander electrode, with the lines and surfaces representing linear interpolations between the simulated values. (b) The lines in parameter space along which the horizontal (blue) and vertical (red) distances between the two field nulls equal zero. The lower and upper bends of the meander electrode are $10\mu\text{m}$ thinner and $20\mu\text{m}$ wider respectively than the centre bend (see the main text for an explanation of these choices).

tation of the actual trap.

HFSS uses the finite-element method, which means that it divides the model space up into a finite number of elements, or pixels. For each of these pixels, which are tetrahedral in shape, it evaluates the electromagnetic fields. It then increases the

number of pixels by splitting some of them into two smaller pixels, and then re-evaluates the fields. This process continues until the specified quantities of interest change by a sufficiently small amount between iterations (called “adaptive passes”), indicating that they have been evaluated to the required level of precision.

Figure 6.10 shows how the field null coordinates change with each adaptive pass, where the above iterative process continues until the simulation fails to run with our 64GB of RAM. The achieved precision seems to be $\sim 100\text{nm}$.




Figure 6.10: Plots showing how the microwave and RF field null coordinates change with the number of adaptive passes completed. x is the horizontal position and y is the height (with the trap axis in z). The first two adaptive passes give field distributions of poor resolution and are excluded for clarity. These simulations were carried out with 64GB of RAM, with the computer unable to increase the precision any further than for the final points shown here.

One aspect of the trap that could deviate significantly from the design to the fabricated trap is the electrode dimensions. Our design has $8\mu\text{m}$ -thick electrodes and

$5\mu\text{m}$ -wide gaps between electrodes, but errors in fabrication could result in different values for these parameters (variations can be up to $\sim 1\mu\text{m}$). Figure 6.11 shows how the distance between the microwave and RF field nulls would increase with changes in the electrode dimensions. The distance could be as large as $\sim 1\mu\text{m}$, which, from the field parameters given in table 6.2, would roughly double the residual microwave field at the location of the RF field null.




Figure 6.11: Plots showing how the distance between the microwave and RF field nulls changes, both in x (top) and in y (bottom), with the width of the gaps between the electrodes (left) and the thickness of the electrodes (right). The distances are defined as the coordinates of the microwave null minus those of the RF null.

One other parameter of the trap that we may want to vary is the temperature. The electrical conductivity, and therefore the skin depth, of gold changes with temperature, causing the location of the microwave field minimum to move. Figure 6.12 shows what the mismatch between the locations of the microwave and RF field nulls would be for 20K, 100K and 300K, which are roughly the temperatures for cooling with liquid helium, cooling with liquid nitrogen or for no cooling respectively.




Figure 6.12: Plots showing how the distance between the microwave and RF field null changes, both in x and in y , with the temperature of the trap. The only change made to the design between these simulations is an adjustment of the conductivity of the electrodes corresponding to the change in temperature. The distances are defined as the coordinates of the microwave null minus those of the RF null.

If the trap were to be used for experiments with the same qubit transition as for the experiments presented in this thesis, the frequency of the microwave currents would change from 3.12GHz to 3.20GHz, resulting in a different current distribution and therefore a shift in the location of the microwave field minimum. Simulating with this different frequency gives an offset between the microwave and RF field nulls of -70nm in the x direction and 20nm in the y direction.

Finally, as discussed before, the electrical properties of what is connected to the input of all electrodes other than the meander (i.e. the boundary conditions of the HFSS simulations) is important for accurate modelling of the fields. We have chosen a geometry that should have minimal coupling of microwave currents in the

meander electrode to the DC and single-qubit microwave electrodes. For a complete trap design, the boundary condition of the RF electrode would need to be carefully modelled, using methods presented in, for example, [47].

Because of the sensitivity of the microwave field distribution to fabrication errors, the decision was made in our group to instead use a ‘U’-shaped microwave electrode for our next trap, so a meander with two bends rather than three, as illustrated in figure 6.13. Such a trap will not null the microwave field, but it should provide good cancellation of the π -polarized radiation, protecting our π -polarized qubit transition from single-qubit rotations during two-qubit gates. This trap will only be useful for π -polarized qubits, but the field cancellation is expected to be less sensitive to fabrication errors. The details of this trap will be provided in a future publication.




Figure 6.13: The trap electrode geometry with a ‘U’-shaped microwave electrode. The blue electrode carries the microwave currents, the red electrodes carry RF voltages, and the black electrodes are connected to the ground plane. The meander and RF electrodes are connected to their sources to the left and right of the diagrams respectively. The design also includes electrodes further away for applying DC fields for trapping.

Chapter 7

Conclusions

7.1 Short-timescale qubit decoherence

The main result of this thesis is the direct probing of qubit decoherence, or memory error, over short timescales. We have measured memory errors below 10^{-5} by interleaving short delays in randomized benchmarking sequences. The decay of qubit state coherence with time in our system does not seem to fit a simple exponential function. The data is fitted well by a model of $1/f$ frequency noise with a low-frequency cutoff.

To put our work into context with other ion trap experiments, we can compare it to the previous use of our system to measure memory errors with Ramsey experiments, which is a typical example of this sort of measurement. In this previous work, the most stringent limit on Ramsey contrast error was $<\sim 5 \times 10^{-2}$ after 2s. Our limit of $<\sim 4 \times 10^{-6}$ after 2ms improves on this result by four orders of magnitude.

We can also compare our results with those of O'Malley *et al.* [75], who performed similar interleaved randomized benchmarking experiments to probe superconducting qubit memory errors. Their most stringent limit is $<\sim 1 \times 10^{-4}$ after 5ns, which we

have improved upon by around an order of magnitude. The reason for us achieving a much more stringent limit may be that our Clifford gate error was much lower, 6×10^{-6} as opposed to $\sim 1 \times 10^{-3}$.

It is commonly assumed in the ion-trapping community that qubit coherence decays exponentially with time, with papers often quoting a measured value for the time constant of this exponential decay, the coherence time T_2^* . We have demonstrated that this assumption does not hold in our system. The evaluated coherence time is strongly dependent upon the range of delay times over which decoherence is measured. Therefore, this quantity has limited significance when attempting to assess the suitability of a system for quantum computing applications.

It is, however, useful to have some single metric with which one can quickly compare the coherence properties of different ion trap quantum computing systems. One might choose some time related to the duration of two-qubit gates, as this is the most relevant timescale for quantum computations (single-qubit gates are generally much shorter than two-qubit gates), and evaluate the average qubit memory error over this time. However, it is not yet clear how long two-qubit gates will be in developed trapped-ion quantum computers. Our near-field microwave two-qubit gates took 3.3ms [36], with laser gates generally taking $30\mu\text{s}$ to $100\mu\text{s}$ [34, 35], and recent experiments achieved fast high-fidelity gates in just $1.6\mu\text{s}$ [83]. Therefore, there is not yet a clear choice of timescale over which to characterize qubit decoherence.

One might instead characterize the qubit decoherence in any given system by the average memory error over the duration of a two-qubit gate performed in that system. While this will likely be helpful when we are nearer to developing a prototype quantum computer, at present most research is focussed on particular challenges in the field, rather than tackling them all in a single system, with very few setups capable of performing high-fidelity two-qubit gates. Therefore, such a characterization

would not be possible with most systems that are currently under development.

We argue that a good metric with which to characterize qubit decoherence is the time at which the average memory error passes some threshold. The fault-tolerant thresholds for the best quantum error correction codes have not changed much in recent years, so it seems likely any memory error threshold based on the current best fault-tolerant threshold of $\approx 1\%$ [32] will remain relevant in the future.

We need the memory error to be much lower than the fault-tolerant threshold to keep the number of physical qubits per logical qubit reasonably low; we choose a value of 10^{-3} . For our system, the time at which this threshold is passed is $\approx 400\text{ms}$. This would be long enough to apply ~ 1000 two-qubit laser gates of duration $100\mu\text{s}$, or ~ 100 of our 3.3ms microwave gates. These large numbers suggest that the overhead of correcting for memory errors (which involves regular application of correction operations) would be relatively small in a computational context.

For the superconducting qubit results, the 10^{-3} memory error time is $\sim 50\text{ns}$ [75]. This is about the same as the two-qubit gate times for superconducting qubits [84], which suggests that a great deal of progress is still required on reducing gate times and/or increasing memory times before superconducting qubits are viable for quantum computing applications.

In addition to the memory preservation results presented in this thesis, our ion trap system has previously been used to achieve single- and two-qubit gates, as well as state preparation and measurement with one or two ions, all with errors below the fault-tolerant threshold. These results are summarized in table 7.1, and suggest that our system is a good candidate for quantum computing applications.

Metric	Value	Reference
SPAM error (one qubit)	7×10^{-4}	[33]
SPAM error (per qubit for two qubits)	3×10^{-3}	[36]
Single-qubit gate error	1×10^{-6}	[33]
Two-qubit gate error	3×10^{-3}	[36]
10^{-3} memory error time	400ms	This work

Table 7.1: Summary of the results achieved in our experimental setup with an intermediate-field atomic clock qubit in trapped $^{43}\text{Ca}^+$ ions.

7.2 Next-generation trap design

We have presented a design study into a new surface ion trap electrode geometry for achieving high-fidelity near-field microwave two-qubit gates. This geometry features an ‘S’-shaped meander electrode [46] to allow for passive nulling of the microwave magnetic field while still providing a large field gradient. The geometry allows for ion chains to be trapped, transported and split with feasible DC and RF voltages.

In previous work on this sort of electrode geometry, the measured distance between the RF and microwave field null was $3.8\mu\text{m}$ [48]. With such a large distance between these nulls, trapped ions would undergo significant single-qubit rotations during each two-qubit gate and/or experience RF micromotion, both of which would reduce the gate fidelity. The deviation from the simulated value of $0.9\mu\text{m}$ was largely due to the microwaves applied to the meander electrode coupling to other electrodes and then reflecting back into the trapping region, changing the microwave field distribution in an unpredictable way. Our design has improved isolation between the meander and the other electrodes, which should reduce the sensitivity of the field distribution to the boundary conditions of the electrodes. In addition, we have developed a design process that allows for the simulated RF and microwave field nulls to be completely overlapped.

The limiting factor in the performance of our trap design may be with fabrication imperfections. A change in the electrode widths or thicknesses of $1\mu\text{m}$ would substantially change the microwave field distribution, moving the microwave field null away from the RF field null. Our group plans to use a ‘U’-shaped electrode in future work, which would only null the microwave field in one direction, but the field distribution should be less sensitive to fabrication errors.

We note that, while we restricted our design study to a single layer of electrodes on the surface trap, there has been some promising research undertaken in recent years on fabricating surface traps with multiple layers of electrodes [85, 86]. This could be useful in developing new traps for near-field microwave gates, for example by having microwave electrodes running underneath the top layer and then emerging to form shapes on the surface. Given that microwave systems are likely to be easier to scale to large numbers of qubits than laser systems, and that high-fidelity single- and two-qubit gates have been demonstrated with near-field microwaves in our setup [33, 36], the development of new, possibly multi-layer, surface traps featuring microwave electrodes seems to be a promising area of future research.

Limerick

Quantum computing is fast

But how long does data last?

Despite what you'd think

We found noise white and pink

So the promise of ions is vast!

Bibliography

- [1] P. Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22(5):563–591, 1980.
- [2] Y. I. Manin. *Vychislimoe i nevychislimoe (Computable and Noncomputable)*. Sov. Radio, 1980.
- [3] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7):467–488, 1982.
- [4] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D.A. Lidar, and M. Troyer. Defining and detecting quantum speedup. *Science*, 345(6195):420–425, 2014.
- [5] S. M. Paul, D. S. Mytelka, C. T. Dunwiddie, C. C. Persinger, B. H. Munos, S. R. Lindborg, and A. L. Schacht. How to improve RD productivity: The pharmaceutical industry’s grand challenge. *Nature Reviews Drug Discovery*, 9:203–214, 2010.
- [6] J. W. Scannell, A. Blanckley, H. Boldon, and B. Warrington. Diagnosing the decline in pharmaceutical R&D efficiency. *Nature Reviews Drug Discovery*, 11:191–200, 2012.
- [7] J. A. DiMasi, H. G. Grabowski, and R. W. Hansen. Innovation in the pharmaceutical industry: New estimates of R&D costs. *Journal of Health Economics*, 47:20–33, 2016.

- [8] The Long View - R&D productivity. Technical report, Bernstein Research, 2010.
- [9] F. Pammolli, L. Magazzini, and M. Riccaboni. The productivity crisis in pharmaceutical R&D. *Nature Reviews Drug Discovery*, 10:428–438, 2011.
- [10] K. Smietana, M. Siatkowski, and M. Møller. Trends in clinical success rates. *Nature Reviews Drug Discovery*, 15:379–380, 2016.
- [11] L. Brocchieri and S. Karlin. Protein length in eukaryotic and prokaryotic proteomes. *Nucleic Acids Research*, 33(10):3390–3400, 2005.
- [12] G. Wu. *Amino Acids: Biochemistry and Nutrition*. CRC Press, 2013.
- [13] Planck Collaboration. Astrophysics Special feature Planck 2015 results. *Astronomy & Astrophysics*, 13, 2016.
- [14] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549:242–246, 2017.
- [15] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proc. R. Soc. Lond. A*, 400:97–117, 1985.
- [16] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. R. Soc. Lond. A*, 439:553–558, 1992.
- [17] L. K. Grover. A fast quantum mechanical algorithm for database search. *Proceedings of the twenty-eighth annual ACM symposium on theory of computing - STOC '96*, pages 212–219, 1996.
- [18] P. Shor. Proceedings of the 35th Annual Symposium on Foundations of Computer Science. page 124, Los Alamitos, 1994. IEEE Computer Society.
- [19] A. Steane. Multiple Particle Interference and Quantum Error Correction. *Proc. R. Soc. Lond. A*, (452):2551–2577, 1996.

- [20] D. P. DiVincenzo. The Physical Implementation of Quantum Computation. 2000.
- [21] W. Paul and M. Raether. Das elektrische Massenfilter. *Zeitschrift fur Physik*, 140(3):262–273, 1955.
- [22] D. J. Wineland, J. J. Bollinger, and W. M. Itano. Spin squeezing and reduced quantum noise in spectroscopy. *Physical Review A*, 46(11), 1992.
- [23] F. Mintert and C. Wunderlich. Ion-Trap Quantum Logic Using Long-Wavelength Radiation. *Physical Review Letters*, 87(25):257904, 2001.
- [24] M. Johanning, A. Braun, N. Timoney, V. Elman, W. Neuhauser, and C. Wunderlich. Individual addressing of trapped ions and coupling of motional and spin states using rf radiation. *Physical Review Letters*, 102(7), 2009.
- [25] A. Khromova, C. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and C. Wunderlich. Designer spin pseudomolecule implemented with trapped ions in a magnetic gradient. *Physical Review Letters*, 108(22), 2012.
- [26] N. Timoney, I. Baumgart, M. Johanning, A. F. Varón, M. B. Plenio, A. Retzker, and C. Wunderlich. Quantum gates and memory using microwave-dressed states. *Nature*, 476:185–188, 2011.
- [27] 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), 2016.
- [28] 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, 1998.
- [29] D. P. L. Aude Craik, N. M. Linke, M. A. Sepiol, T. P. Harty, J. F. Goodwin, C. J. Ballance, D. N. Stacey, A. M. Steane, D. M. Lucas, and D. T. C. Allcock.

- High-fidelity spatial and polarization addressing of $^{43}\text{Ca}^+$ qubits using near-field microwave control. *Physical Review A*, 95(022337), 2017.
- [30] 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), 2008.
 - [31] 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*, 476:181–184, 2011.
 - [32] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3), 2012.
 - [33] 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), 2014.
 - [34] 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), 2016.
 - [35] 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 $\text{Be } 9^+$ Ion Qubits. *Physical Review Letters*, 117(6), 2016.
 - [36] 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):140501, 2016.
 - [37] Q. Turchette, C. Wood, B. King, C. Myatt, D. Leibfried, W. Itano, C. Monroe, and D. Wineland. Deterministic Entanglement of Two Trapped Ions. *Physical Review Letters*, 81(17):3631–3634, 1998.

- [38] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Experimental Comparison of Two Quantum Computing Architectures. *PNAS*, 114(13):3305–3310, 2017.
- [39] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, 2002.
- [40] J. M. Amini, H. Uys, J. H. Wesenberg, S. Seidelin, J. Britton, J. J. Bollinger, D. Leibfried, C. Ospelkaus, A. P. VandeVender, and D. J. Wineland. Toward scalable ion traps for quantum information processing. *New Journal of Physics*, 12, 2010.
- [41] P. Kaufmann, Ti. F. Gloger, D. Kaufmann, M. Johanning, and C. Wunderlich. High-Fidelity Preservation of Quantum Information During Trapped-Ion Transport. *Physical Review Letters*, 120(1):10501, 2018.
- [42] 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), 2012.
- [43] 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), 2012.
- [44] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.
- [45] I. V. Inlek, C. Crocker, M. Lichtman, K. Sosnova, and C. Monroe. Multispecies Trapped-Ion Node for Quantum Networking. *Physical Review Letters*, 118(25), 2017.
- [46] M. Carsjens, M. Kohnen, T. Dubielzig, and C. Ospelkaus. Surface-electrode Paul trap with optimized near-field microwave control. *Applied Physics B*, 114:243–250, 2014.

- [47] M. Wahnschaffe. *Engineered microwave control for trapped ions*. PhD thesis, Universität Hannover, 2016.
- [48] M. Wahnschaffe, H. Hahn, G. Zarantonello, T. Dubielzig, S. Grondkowski, A. Bautista-Salvador, M. Kohnen, and C. Ospelkaus. Single-ion microwave near-field quantum sensor. *Applied Physics Letters*, 110(3):034103, 2017.
- [49] T. P. Harty. *High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field 43Ca^+* . PhD thesis, University of Oxford, 2013.
- [50] M. A Sepiol. *A high-fidelity microwave driven two-qubit quantum logic gate in 43Ca^+* . PhD thesis, University of Oxford, 2016.
- [51] S. Bravyi and A. Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71(2), 2005.
- [52] E. Magesan, J. M. Gambetta, and J. Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical Review Letters*, 106(18), 2011.
- [53] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, Marcus P. Da Silva, George A. Keefe, Mary B. Rothwell, Thomas A. Ohki, Mark B. Ketchen, and M. Steffen. Efficient measurement of quantum gate error by interleaved randomized benchmarking. *Physical Review Letters*, 109(8), 2012.
- [54] S. Wimperis. Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments. *Journal of Magnetic Resonance, Series A*, 109(2):221–231, 1994.
- [55] L. S. Cutler and C. L. Searle. Some Aspects of the Theory and Measurement of Frequency Fluctuations in Frequency Standards. *Proceedings of the IEEE*, 54(2):136–154, 1966.
- [56] B. Keitch. *A quantum memory qubit in Calcium- 43* . PhD thesis, 2007.

- [57] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. PhD thesis, University of Oxford, 2011.
- [58] H. A. Janacek. *Optical Bloch equations for simulating trapped-ion qubits*. PhD thesis, University of Oxford, 2015.
- [59] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hogervorst. Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments. *Applied Physics B*, 73(8):861–863, 2001.
- [60] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Makayama, J.-P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photoionization for calcium ion trapping. *Physical Review A*, 69:012711, 2004.
- [61] I. Siemers, M. Schubert, R. Blatt, W. Neuhauser, and P. E. Toschek. The trapped state of a trapped ionline shifts and shape. *EPL*, 18(2):181–186, 1992.
- [62] 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(023043), 2016.
- [63] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B: Lasers and Optics*, 66(2):20, 1998.
- [64] H. Dehmelt. Proposed laser fluorescence spectroscopy of mono-ion oscillator II. *Bulletin of the American Physical Society*, 20(1):60, 1975.
- [65] 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), 2008.
- [66] C. Monroe, C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, and D. J. Wineland. Scalable entanglement of trapped ions. *AIP Conference Proceedings*, 551:173–186, 2001.

- [67] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, 2010.
- [68] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Surface-Electrode Architecture for Ion-Trap Quantum Information Processing. *Quantum Information and Computation*, 5(6):419–439, 2005.
- [69] D. T. C. Allcock, T. P. Harty, C. J. Ballance, B. C. Keitch, N. M. Linke, D. N. Stacey, and D. M. Lucas. A microfabricated ion trap with integrated microwave circuitry. *Applied Physics Letters*, 102(044103), 2013.
- [70] 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), 2006.
- [71] I. A. Boldin, A. Kraft, and C. Wunderlich. Measuring Anomalous Heating in a Planar Ion Trap with Variable Ion-Surface Separation. *Physical Review Letters*, 120(2):23201, 2018.
- [72] D. T. C. Allcock, J. A. Sherman, D. N. Stacey, A. H. Burrell, M. J. Curtis, G. Imreh, N. M. Linke, D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas. Implementation of a symmetric surface-electrode ion trap with field compensation using a modulated Raman effect. *New Journal of Physics*, 12(053026), 2010.
- [73] J. R. Bainter. Active filter has a stable notch, and response can be regulated. *Electronics*, pages 115–117, 1975.
- [74] P. E. Fleischer and J. Tow. Design Formulas for Biquad Active Filters using Three Operational Amplifiers. *IEEE Proc*, pages 662–663, 1973.
- [75] P. J. J. O’Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I. C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Qubit

- Metrology of Ultralow Phase Noise Using Randomized Benchmarking. *Physical Review Applied*, 3(4), 2015.
- [76] J. Labaziewicz, Y. Ge, D. R. Leibbrandt, S. X. Wang, R. Shewmon, and I. L. Chuang. Temperature dependence of electric field noise above gold surfaces. *Physical Review Letters*, 101(18), 2008.
- [77] J. Chiaverini and J. M. Sage. Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range. *Physical Review A*, 89(1), 2014.
- [78] E. A. Guttierrez-D., M. J. Deen, and C. L. Claeys. *Low temperature electronics: physics, devices, circuits, and applications*. Academic Press, 2001.
- [79] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10):5025–5033, 1998.
- [80] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Physical Review A*, 78(3), 2008.
- [81] J. P. Home and A. M. Steane. Electrode Configurations for Fast Separation of Trapped Ions. *Quantum Inf. Comput.*, 6:289, 2006.
- [82] M. Palmero, S. Martínez-Garaot, U. G. Poschinger, A. Ruschhaupt, and J. G. Muga. Fast separation of two trapped ions. *New Journal of Physics*, 17(9):93031, 2015.
- [83] 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*, 555:75–78, 2018.
- [84] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner,

- A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508:500–503, 2014.
- [85] N. D. Guise, S. D. Fallek, K. E. Stevens, K. R. Brown, C. Volin, A. W. Harter, J. M. Amini, R. E. Higashi, S. T. Lu, H. M. Chanhvongsak, T. A. Nguyen, M. S. Marcus, T. R. Ohnstein, and D.W. Youngner. Ball-grid array architecture for microfabricated ion traps. *Journal of Applied Physics*, 117(17), 2015.
- [86] P. Maunz. High Optical Access Trap 2.0. Technical report, Sandia National Laboratory, Albuquerque, 2016.

Appendix A

Ground level structure of $^{43}\text{Ca}^+$

M_m	M_n	Frequency (GHz)	df/dB (MHz G ⁻¹)	$\langle M_n \mu_q M_m \rangle$ (μ_B)
-4	-3	3.589 033 180	2.519 894	+1.342 420
-3	-3	3.543 000 725	2.237 093	+0.602 802
-3	-2	3.495 825 733	1.940 390	+1.195 577
-2	-3	3.495 742 031	1.939 817	-0.204 417
-2	-2	3.448 567 039	1.643 113	-0.810 866
-2	-1	3.400 062 158	1.330 090	-1.040 756
-1	-2	3.399 978 455	1.329 517	-0.363 393
-1	-1	3.351 473 573	1.016 493	-0.932 839
-1	0	3.301 519 669	0.684 932	-0.876 791
0	-1	3.301 435 966	0.684 359	-0.528 271
0	0	3.251 482 062	0.352 797	-0.993 060
0	+1	3.199 941 077	0	-0.702 127
+1	0	3.199 857 374	-0.000 573	+0.702 254
+1	+1	3.148 316 389	-0.353 370	+0.993 034
+1	+2	3.095 026 842	-0.730 728	-0.514 410
+2	+1	3.094 943 139	-0.731 301	+0.887 366
+2	+2	3.041 653 592	-1.108 659	-0.919 344
+2	+3	2.986 424 605	-1.514 736	+0.308 504
+3	+2	2.986 340 903	-1.515 309	-1.085 679
+3	+3	2.931 111 915	-1.921 385	+0.728 641
+4	+3	2.873 631 427	-2.362 039	+1.299 654


Table A.1: The resonance frequencies, frequency dependence on the static magnetic field B , and coupling strengths for all magnetic dipole transitions in the ground $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at 146.0942G, with symbols as defined in equation 2.19. The second order frequency dependence on magnetic field for the $|F=4, M=0\rangle - |F=3, M=+1\rangle$ transition is 2.416kHz G⁻².


M_m	M_n	Frequency (GHz)	df/dB (MHz G ⁻¹)	$\langle M_n \mu_q M_m \rangle$ (μ_B)
-4	-3	3.949 771 857	2.569 912	+1.355 692
-3	-3	3.867 368 475	2.337 131	+0.552 243
-3	-2	3.781 465 945	2.086 061	+1.233 104
-2	-3	3.781 301 064	2.085 488	-0.169 444
-2	-2	3.695 398 534	1.834 419	-0.756 705
-2	-1	3.605 294 276	1.560 347	-1.099 756
-1	-2	3.605 129 395	1.559 774	+0.305 995
-1	-1	3.515 025 138	1.285 701	+0.889 435
-1	0	3.420 035 468	0.982 371	+0.453 176
0	-1	3.419 870 588	0.982 371	+0.453 176
0	0	3.324 880 918	0.679 614	+0.971 138
0	+1	3.224 113 796	0.340 094	+0.789 034
+1	0	3.223 948 915	0.339 521	+0.616 028
+1	+1	3.123 181 794	0	+1.001 026
+1	+2	3.015 435 346	-0.388 298	+0.601 736
+2	+1	3.015 270 465	-0.388 871	-0.799 970
+2	+2	2.907 524 018	-0.777 169	+0.961 755
+2	+3	2.791 111 000	-1.233 169	-0.379 132
+3	+2	2.790 946 119	-1.233 742	-1.012 843
+3	+3	2.674 533 102	-1.689 743	+0.798 542
+4	+3	2.546 787 218	-2.246 218	-1.267 363

Table A.2: The resonance frequencies, frequency dependence on the static magnetic field B , and coupling strengths for all magnetic dipole transitions in the ground $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at 287.7827G, with symbols as defined in equation 2.19. The second order frequency dependence on magnetic field for the $|F = 4, M = +1\rangle - |F = 3, M = +1\rangle$ transition is 2.4514kHz G⁻².

Appendix B


Magnetic field control circuits





Narrow Notch Filter (17 kHz)

Bainter design



```

1) - NOTCH1_IN
2) 2 * NOTCH1_IN <- check here for saturation
3) NOTCH1_IN

```

J. Bainter: "Active filter has stable notch, and response can be regulated"
Electronics, October/2, 1975

Low-Pass Notch Filter

Bainter design

double-pole at 16 KHz
double-zero at 51 KHz

$$G = K3 * (s^2 + w_z^2) / (s^2 + s * w_0 / Qp + w_p^2)$$

$$w_z = 2\pi f * 51 \text{ KHz}$$

$$w_p = 2\pi f * 16 \text{ KHz}$$


$$Qp = 0.99$$

notch at 51 KHz

lowpass with pole at 16 KHz

lower freq: unity gain

higher freq: -30 dB



opamp output at low frequencies:
1) -NOTCH2_IN
2) 10 * NOTCH2_IN
3) 10 * NOTCH2_IN
4) NOTCH2_IN
max. input: 0.1 V rms


J. Bainter: "Active filter has stable notch, and response can be regulated"
Electronics, October/2, 1975

11/12/2017 18:00 Date:

4/12 Sheet: current stabiliser v4 LS8 TITLE:

Compensator

Biquad Active All-Pass Filter
(Tow-Thomas Design)





maximum opamp output:
1) NOTCH2 (lowpass)
2) 10T NOTCH2 (compensator) <- check COMP_OUT for saturation
3) NOTCH2 (bandpass)

P. E. Fleischer and J. Tow:
Design Formulas for Biquad Active Filters
Using Operational Amplifiers
Proc. IEEE, 562, May 1973


$$G = (m s^2 + c s + d) / (s^2 + a s + b)$$

n = 62.79
c = 7138 rad/s
d = 7138 (rad/s)^2
a = 886.5
b = 21865 rad/s
zeros at ~5 Hz, poles at ~25 Hz






Feedforward DAC



Voltage Reference


provides 2.5V reference for DACs



11/12/2017 18:00 Date:


9/12 Sheet: current stabiliser v4 LS8 TITLE:

SPI Isolators





Voltage Regulators


+12V, -12V for OpAmp Power Supply
+5V for DACs/ADC, SPI Isolators and Reference LTC6655 Power Supply




If output ringing occurs:
put 10uF tantalum capacitor (high ESR)
in parallel to the 10u ceramic input capacitor



low ESR ceramic (X7R, or X5R) input/output capacitors




lowpass,
pole at 1.6 kHz




Decoupling capacitors


11/12/2017 18:00		Date:
12/12	Sheet:	current stabiliser v4 LS8 TITLE:

IMPORTANT!!! The configuration here seems to result in a strong 3MHz oscillation at shunt output. In Lab 2, we have removed C5 and replaced C4 with 1nF (and the equivalent change for the feedback loop)




Feedback Shunt






Connections to the main circuit board



The BeagleBone only supports +3.3V on all I/O pins.
Do not connect to 5V logic signals!

Software is for MOSI_PIC and MISO_PIC pins swapped around - should do this manually for v2 boards and swap traces when designing v3 (JET)

Feedback / ADC SPI connections
 <CLK_FB, MOSI_FB, MISO_ADC,
 CS_FB_COARSE, CS_FB_FINE, CS_ADC>



Feedforward SPI connections (CLK_FF, MOSI_FF, CS_FF)

