A high-fidelity microwave driven two-qubit quantum logic gate in $^{43}Ca^{+}$

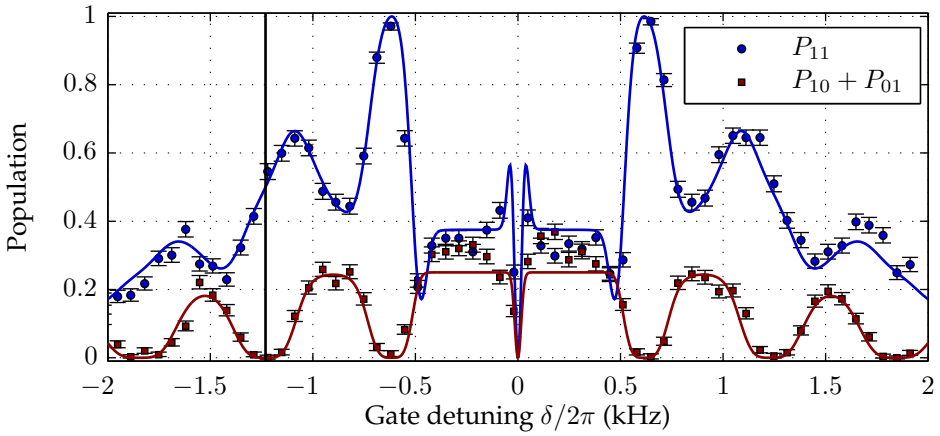

Martin A. Sepiol

A thesis submitted for the degree of
Doctor of Philosophy

Hertford College
Trinity Term, 2016

Department of Physics, University of Oxford

Abstract

A high-fidelity microwave driven two-qubit quantum logic gate in $^{43}Ca^{+}$
Martin A. Sepiol
A thesis submitted for the degree of Doctor of Philosophy
Trinity Term 2016
Hertford College, Oxford

Quantum computers offer great potential for significant speedup in executing certain algorithms compared to their classical counterparts. One of the most promising physical systems in which implementing such a device seems viable are trapped atomic ions. All of the fundamental operations needed for quantum information processing have already been experimentally demonstrated in trapped ion systems. Today, the remaining two obstacles are to improve the fidelities of these operations up to the point where quantum error correction techniques can be successfully applied, as well as to scale up the present systems to a higher number of quantum bits (qubits).

This thesis addresses both issues. On the one hand, it decribes the experimental implementation of a high-fidelity two-qubit quantum logic gate, which is the most technically demanding fundamental operation to realise in practice. On the other hand, the presented work is carried out in a microfabricated surface ion trap – an architecture that holds the promise of scalability.

The gate is applied directly to hyperfine "atomic clock" qubits in $^{43}Ca^{+}$ ions using the near-field microwave magnetic field gradient produced by an integrated trap electrode. To protect the gate against fluctuating energy shifts of the qubit states, as well as to avoid the need to null the microwave field at the position of the ions, a dynamically decoupled Mølmer-Sørensen scheme is employed. After accounting for state preparation and measurement errors, the achieved gate fidelity is 99.7(1)%. In previous work, the same apparatus has been used to demonstrate coherence times of $T_2^* \approx 50\,\mathrm{s}$ and all single-qubit operations with fidelity $> 99.95\%$.

To gain access to the "atomic clock" qubit transition in $^{43}Ca^{+}$, a static magnetic field of 146 G is applied. The resulting energy level Zeeman-structure is spread over many times the linewidth of the atomic transition used for Doppler cooling. This thesis presents a simple and robust method for Doppler cooling and obtaining high fluorescence from this qubit in spite of the complicated level structure. A temperature of 0.3 mK, slightly below the Doppler limit, is reached.

Acknowledgements

The work presented in this thesis would not have been possible without the support and the assistance of many people. I take this opportunity to acknowledge them and extend my sincere gratitude for helping me make this DPhil thesis a possibility, as well as for making my time in Oxford so enjoyable.

First and foremost, I wish to thank my supervisor, Prof. David Lucas, for taking me on as a graduate student, and for his encouragement and enthusiastic involvement in the experiments. He has provided me with much needed guidance during my time as a DPhil student, and always had an open ear for any problems that would arise. I would also like to thank him for the time he took to proof-read this thesis.

The foundations for this work have been laid by Dr. David "D3" Allcock and Dr. Tom Harty. Tom was the post-doc in the "new lab" during the past four years and taught me most of what I know about ion trapping and microwaves. I greatly appreciate his patience and support as we hurdled all the obstacles in the completion of this project.

My grateful thanks are also extended to Prof. Derek Stacey and Dr. Hugo Janacek for their collaboration on the dark resonance cooling experiment – it has been a pleasure to work with such constructive and friendly colleagues. I am further grateful to Prof. Derek Stacey for helping me prepare for the Hertford College's Carreras senior scholarship application, as well as for sharing his teaching expertise with me.

A huge thank you to all other students and post-docs who helped build the experiment up to what it is today. Chris Ballance for his supportive suggestions, comments, and advice; James Tarlton for all his efforts put into designing and assembling electronics; and Mathieu Dos Santos for technical support.

I would also like to thank all remaining labmates and members of the Oxford ion-trap group, past and present, with whom I had the chance to work and socialise: Dr. Diana Prado Lopes Aude Craik, Dr. Norbert Linke, Vera Schäfer, Keshav Thirumalai, Jochen Wolf, Laurent Stephenson, Dr. Joe Goodwin, Dr. Timothy Ballance, Amy Hughes, Clemens Löschnauer, David Nadlinger, Benjamin Merkel and Rustin Nourshargh. Many thanks also to the staff members of the Clarendon mechanical workshop, especially to Simon Moulder and Mathew Newport.

Outside the world of physics, I would like to acknowledge friends and family who supported me during my time in Oxford. Larissa, who has always been there

for me and whose unconditional encouragement and kindness have kept me going. My parents and my sister, for all the moral, practical, and financial support over the years. Felix and Iris, who I have been lucky to have had travelling this same journey alongside me, for many pleasant brunches, dinners, and evenings in my former College, Christ Church.

Finally, I thankfully acknowledge funding by the EPSRC and by Hertford College's Carreras senior scholarship.

Contents

1

Introduction

1.1 Classical vs. quantum information processing

The history of digital information processing is a history of rapid technological development, ranging from vacuum tubes in the 1940s to modern day miniaturised integrated circuits. In spite of the fast progress, the foundations of how a computation is carried out have changed remarkably little and follow a simple recipe.

First, the information to be processed is encoded in terms of binary digits (bits), which in practice take the form of two voltage levels, representing binary 0 and 1, at the terminals of a transistor on a silicon microchip. The calculation itself follows from subjecting these bits to a succession of logical operations called gates, the physical realisation of which is based upon the transistors' capacity to switch between the two voltages. The problem-specific sequence of logical gates is called an algorithm, which when completed yields the desired result.

This way of information processing is based entirely on the familiar laws of classical physics and is therefore termed "classical computation". Notwithstanding its success in meeting society's growing demand for processing power, classical computation has its limits. When facing problems with an exponentially (or super-polynomially) growing number of possible solutions, as is often the case in the real-world, classical algorithms often require an unfeasible amount of

time and resources to examine all potential answers and identify the correct one.

The difficulty of handling this so called "combinatorial explosion" of the solution space by classical computers can be demonstrated in a simple way by looking at the famous travelling salesman problem. There, the objective is to find the shortest round trip route passing through N cities, given the distances between each pair of cities and given that each city should be visited only once. Even for modest N of e.g. 32, and assuming the currently fastest supercomputer speeds of order 10^{16} operations per second [SDSM16], a brute-force algorithm seeking through all $(N-1)!/2$ routes would run for about 13 billion years, which is around the age of the universe. Although more sophisticated methods for tackling intractable tasks of this sort exist, e.g. heuristic and approximation algorithms, taming exponentially or super-polynomially scaling solution spaces remains a near impossible undertaking.

The situation is somewhat different when instead of obeying the rules of classical physics, a computer behaves like a quantum system and thus can make use of non-classical effects like superposition, entanglement and quantum interference. Such a device is called a quantum computer and was first proposed in the seminal work of D. Deutsch in 1985 [Deu85]. Its physical realisation is being actively pursued using technologies based on trapped atoms, superconducting circuits, photons, quantum dots, dopants in solids, nuclear magnetic resonance, and others [LJL+10]. At present it is difficult to predict which platform, if any, will ultimately prove successful, although recent advances in ion trap and superconducting circuit technologies seem to give them an edge over the competition.

At first glance, the "universal quantum computer" envisaged by D. Deutsch employs a similar kind of recipe as its digital counterpart. Information is first encoded in terms of quantum bits (qubits), that is, two-level quantum mechanical systems. These qubits are then driven through an evolution as prescribed by a

problem-specific algorithm. The algorithm consists of a succession of quantum gates that, after successful execution, put the qubits in the desired solution state.

However, the contrast to classical computing comes from the ability to prepare n qubits in a superposition of the 2^n possible input states that n regular bits could adopt. The algorithm's purpose is to entangle and interfere appropriate parts of this superposition in order to generate an output superposition that is based on information pertaining to all input states. A final measurement on the qubits then collapses this output superposition and leaves the qubits in the correct solution state with high probability.

Being able to process superpositions of input states instead of working through each of them individually is called "quantum parallelism" and represents the essence of why quantum computers can efficiently perform some classically infeasible tasks. A prominent example is the search of unstructured databases, which is a category of problems that, amongst others, contains the previously mentioned travelling salesman problem. It was shown by L. A. Grover [Gro96, Gro97] that running a quantum search algorithm on a quantum computer provides a square root speedup over classical brute-force approaches. Grover's algorithm has been demonstrated experimentally with up to four qubits using many physical systems e.g. nuclear magnetic resonance [JMH98, CGK98, VSS+00, EF02, OLK03], trapped ions [BHL+05], photons [KMSW00, WRR+05] and superconducting circuits [DCG+09, WAIQ13].

Another important example of a problem that is believed to be intractable on any classical machine concerns the factoring of large integers into prime numbers. The difficulty of this problem lies at the heart of many classical public-key cryptosystems for data transmission, such as RSA. With the help of Shor's algorithm [Sho97], which is part of a broader class of quantum algorithms based on the quantum Fourier transform [NC11], an exponential speedup over classically known algorithms can be achieved. Shor's algorithm has been

demonstrated experimentally on a small scale using liquid-state nuclear magnetic resonance techniques [VSB+01], cold atoms [GWM+08], photonic systems [LBYP07, LWL+07, PMO09, MLLL+12] and superconducting circuits [LBC+12].

A final example of a task at which quantum computers may excel is the simulation of quantum systems, promising new insights to be gained in the fields of condensed-matter physics, quantum chemistry and many others. For classical computers, such simulations are intractable due to the exponential increase of the required resources with the size of the system to be described. However, this scaling is linear when using qubits, rendering quantum computers well-suited for the job, provided an algorithm allowing for the extraction of the desired information about the system can be found.

It is worth noting that an alternative approach to quantum simulation involves the use of purpose built "quantum simulators", as proposed by R. Feynman in 1982 [Fey82]. These devices are well controlled quantum systems that can be used for the study of experimentally harder to control or inaccessible ones. The chief benefit of quantum simulators is that they are much easier to construct than a universal quantum computer. For this reason, small scale quantum simulation has already been demonstrated on various experimental platforms, as recently summarised in [GAN14].

1.2 How to build a universal quantum computer

D. Deutsch introduced the universal quantum computer as the quantum generalisation of the universal Turing machine, meaning that it could simulate the action of any other computer. In contrast to e.g. quantum simulators or quantum annealing processors [JAG+11], which are devices designed to solve specific problems, a universal quantum computer is general-purpose, in that it can tackle all computable problems. Its physical implementation must satisfy five basic criteria, known as the DiVincenzo criteria [DiV00], that are listed below.

A scalable physical system with well characterised qubits: In principle, any two-level quantum mechanical system can be used as a qubit. The key is to pick one that remains well defined and controllable when accommodated en masse on a suitable platform, without creating unmanageable overheads on classical resources like laser systems, cryogenic refrigerators, electronics, etc. However, most current experiments, including the work presented in this thesis, are carried out using small numbers of qubits. Thus, the question of whether an approach can be declared "scalable" is often a difficult one to answer, as true large scale implementations face many complex engineering issues that in most cases have only recently started to be considered.

The ability to initialise the state of the qubits to a simple fiducial state, such as $|000\ldots\rangle$: This means that at the start of every algorithm, it must be possible to prepare the qubits involved in one of their two states with high fidelity.

Long relevant decoherence times, much longer than the gate operation time: The most crucial requirement for the successful implementation of quantum computing is to ensure that the system is perfectly isolated from its environment. Small amounts of information leakage to, or noise intrusion from, the outside world can cause the fragile quantum states to be disturbed, a process known as decoherence. Yet, the programmer, who is part of the environment, needs to be able to interact with the quantum system in order to carry out the algorithm. Thus, the paradox faced by all technologies trying to realise quantum computing, is that opening the door for interactions that are necessary for performing the computation inevitably allows for some degree of decoherence to slip through and cause errors.

However, just like in classical computation, it is possible to detect and remove small errors at the cost of additional resources, i.e. a larger number of qubits. Such "quantum error correction" techniques can be blended into quantum algo-

rithms and allow for an error-free execution of the calculation, provided the error probability per logic gate is below the so called "fault-tolerant threshold".

Various error correcting schemes have been proposed, some of which have threshold values at the 1% level, such as the "surface code" [FMMC12]. However, the consensus within the scientific community is that quantum gates with error levels of order 10^{-4} are required for the construction of a practicable quantum computer [Kni10, Pre98], as otherwise the amount of extra resources needed to to implement error correction becomes unfeasible.

DiVincenzo's third criterion therefore demands quantum gates to be fast and robust enough for errors from all contributing sources of decoherence to not exceed the fault-tolerant threshold. Fulfilling this requirement is the goal of the research done by many groups and is also the main motivation for the work presented in this thesis.

A universal set of quantum gates: A universal quantum computer operates on a large Hilbert space (the state space of the system) by executing an algorithm that is compiled from a finite set of quantum gates. The set is finite in the sense that if we decide to e.g. increase the number of qubits in the processor, no new types of gates will be required to incorporate them in the calculation. In fact, it has been shown [BBC$^+$95, DBE95, DiV95] that all operations that are possible on a universal quantum computer can be constructed from a complete set of single qubit gates and an entangling operation on two qubits.

A qubit-specific measurement capability: It must be possible to read out the state of a qubit with high fidelity, without affecting the state of the other qubits that take part in the computation.

1.3 Processing information with trapped ion qubits

Individual atoms are compelling candidates for qubits because they are standards, meaning that in contrast to man-made quantum information carriers like quantum dots or superconducting circuits, atoms are identical throughout the universe and don't change their properties over time. When ionised, charged atoms can be strongly confined and suspended in free space with the help of radiofrequency Paul traps under a near-perfect vacuum.

The workhorse of ion trap geometries for quantum computing purposes is the linear Paul trap. It creates a harmonic trapping potential, which in combination with the mutual Coulomb repulsion of the ions induces them to form a linear chain. After employing laser cooling techniques, the system closely resembles an array of (almost) equally spaced atoms at rest in free space, providing excellent isolation from the environment, and hence protection from decoherence.

Within each ion, the qubit is formed by two stable or metastable electronic states, that for the most part can be initialised and read out using optical methods. Quantum gates are implemented with the help of appropriate radiation pulses, that cause transitions between the qubit levels. Two-qubit gates additionally rely on the Coulomb interaction between the ions, which causes the chain to behave much like an array of masses coupled by springs. Thus, the ion array exhibits collective modes of harmonic oscillatory motion with discrete frequencies. The amplitudes of these shared modes, when sufficiently cooled, become quantised and can be used to mediate quantum information between ions. They are equivalent to quantum harmonic oscillator states, and we will sometimes refer to them as the ions' "external" degrees of freedom. Accordingly, the qubit states themselves will be called "internal" degrees of freedom, or simply "spins".

Ensuring that the ions can be laser cooled, and that the qubits encoded within them can be initialised and read out, without having to use many different lasers, restricts the selection of atomic species to ions with a single valence

electron. Consequently, such ions have a simple, alkali-like energy level structure. Commonly used species are the group II alkaline earth metals (Be, Mg, Ca, Sr, and Ba) and some transition metals (Zn, Hg, Cd, and Yb). In general, these ions offer a number of pairs of energy levels that can be used as qubits. The two main types are those formed by two stable electronic states in the ground level, separated by a radiofrequency or microwave transition, and those formed by one stable state in the ground level and one metastable excited state, separated by a forbidden optical transition.

The latter type of qubit is often referred to as an "optical qubit", and in spite of offering some advantages, e.g. a simple readout scheme, it faces two major drawbacks. First, its lifetime is limited to typically $\sim 1\,\mathrm{s}$ by the decay of the excited state. Second, it can only be manipulated by means of a laser that needs to be spectrally narrow in order to maintain phase coherence with the qubit over the course of a computation. The required linewidth is typically at the few Hz level, rendering such a laser technically challenging.

In contrast, the lifetime of ground state qubits is in practice only restricted by ion loss, collisions with residual gas particles in the vacuum chamber, and human error. They can be manipulated either directly, using RF or microwave radiation that can be generated by commercial off-the-shelf equipment, or indirectly, by driving two-photon stimulated Raman transitions with lasers that are much less technically demanding than those used for handling optical qubits.

1.4 Scaling trapped ion architectures

Although it is possible to increase the number of ions accommodated in a single linear Paul trap, this approach leads to a rise in the number of possible vibrational modes of the ion chain. Since each mode has a distinct frequency, long ion arrays have a high spectral density of their motional excitation structure. Consequently, it becomes hard to address and control individual modes with

appropriate radiation pulses without causing off-resonant cross talk between them. This, however, is a necessary requirement for two-qubit gates. Hence, the "spectral crowding" of motional modes restricts algorithms performed in a single linear Paul trap to 10 - 100 qubits [MK13].

A route to scale the trapped ion quantum information processor beyond this limit, is to use the so called "quantum charge-coupled device" (QCCD) approach, as proposed in [WMW+98] and further elaborated in [KMW02]. Here, the underlying platform ought to be microfabricated and can either be a segmented 3D multilayer electrode structure, e.g. [WSGS12], or a 2D planar electrode arrangement [CBB+05]. The key idea is to implement multiple interconnected linear Paul traps in a single device, each capable of holding a manageable number of qubits, and each having a particular purpose, e.g. ion loading, qubit storage, gate execution or readout. Hence, carrying out an algorithm would involve ions being shuttled between different trapping zones by varying the electrode voltages appropriately. The ultimate number of qubits that can be used on a single QCCD is most likely going to be limited by optical and electronic access as well as the complexity of classical hardware needed for qubit control and shuttling.

A potential workaround for this problem, described in [MK13, NFB14], is to use a set of individual modules, each of which can be a linear Paul trap or QCCD, holding as many qubits as technically feasible. Entanglement can be spread between modules in a nondeterministic fashion using photonic links, that are allowed to be noisy and lossy at the cost of computer clock speeds. The main appeal of this approach is that it is freely scalable once a recipe for constructing a fault-tolerant module is developed.

1.5 Thesis focus and structure

Much technological progress has been made in pursuit of a viable QCCD, such as the development and testing of a surface electrode architecture with 150 trap

zones and the successful ion transport through its junctions [AUW+10]. Research is now centred around pushing the fidelities of laser driven quantum gates beyond the fault tolerant threshold, while also devising a way to perform these operations in parallel across multiple trap zones. Since the current approach of surrounding traps with bulk optics is not scalable, surface traps with integrated optical elements are being actively developed [VCA+10, SNJ+11]. However, it would be advantageous to establish alternative methods of manipulating qubits that are less heavily reliant on these prototypical components, or do not use lasers altogether.

This thesis is concerned with such an alternative method, namely that of driving quantum logic operations by means of microwave signals applied to trap electrodes [OLA+08, OWC+11, All11]. More specifically, it describes experimental efforts to increase the fidelity of a two-qubit quantum logic gate, which, as will become apparent in chapter 3, is the most challenging fundamental operation to realise in practice. The underlying architecture allows for full electronic control over all coherent qubit manipulations and thus better prospects for scalability.

The remainder of the thesis is structured into eight chapters.

Chapter 2 introduces how a calcium ion is used as a qubit, discusses its basic atomic properties as well as methods employed for photo-ionization, laser cooling, qubit initialisation and readout.

Chapter 3 explains single and two-qubit quantum logic in a conceptual way and subsequently describes how gates can be implemented using microwave radiation.

Chapter 4 covers the theoretical framework of the Mølmer-Sørensen gate, how it is affected by various sources of error, and how the standard gate recipe can be refined to achieve increased robustness.

Chapter 5 provides an overview of the experimental setup and apparatus.

Chapter 6 describes our experimental investigation of sub-Doppler and

ground state cooling in $^{43}Ca^{+}$, despite its complicated level structure at 146 G.

Chapter 7 presents the results of our two-qubit gate experiments, including details about the pulse sequence used and our tune-up procedure.

Chapter 8 characterises the experimental errors occurring in our two-qubit gate and gives the corresponding error budget.

Chapter 9 concludes this thesis, summarising our results, comparing them to other work in the field and giving an outlook on future work.

2

Calcium ions as qubits

The two main isotopes of calcium used for quantum information purposes are $^{40}Ca^{+}$ and $^{43}Ca^{+}$. In the former, qubits can be encoded in the ground level Zeeman structure [LDH+03] as well as in one of the electric quadrupole transitions [SKGR+03]. In the latter, qubits can be implemented in the ground level hyperfine structure. This work makes use of the latter, $^{43}Ca^{+}$, for two reasons. First, electronically driven quantum logic requires the qubit transition to be in the microwave or radiofrequency domain, calling for a transition within the ground state. Second, hyperfine qubits in $^{43}Ca^{+}$ offer superior coherence properties and a higher fidelity readout scheme compared with Zeeman qubits in $^{40}Ca^{+}$.

The following chapter reviews the basic atomic properties of $^{43}Ca^{+}$ and explains the techniques applied for photo-ionization, laser cooling, qubit initialisation and readout.

2.1 Clock qubits in the ground state of $^{43}Ca^{+}$

All stable isotopes of calcium have a nuclear-spin of zero, the only exception being $^{43}Ca^{+}$, which has $I = 7/2$ and therefore exhibits hyperfine structure. As a consequence, its ground state is split by 3.2 GHz into two manifolds, the energetically higher one consisting of 7 degenerate M_n sublevels, and the energetically lower one consisting of 9 degenerate M_m sublevels.

This degeneracy is lifted when applying an external, static magnetic field that, to a good approximation, interacts solely with the magnetic moment of the ion's electron system. We differentiate between the "weak field regime", where the hyperfine interaction dominates over the Zeeman interaction, and the "strong field regime", where the situation is reversed. The energies of the M_n and M_m sublevels are linear functions of the external field strength. However, their slopes may be different in both regimes, leading to curved trajectories in the so called "intermediate field regime", where the functions join. This behaviour is illustrated in Fig. 2.1 for the $4S_{1/2}$ ground level, and in [Kei07, All11, Har14, Jan15] for all remaining relevant levels in $^{43}Ca^{+}$.

The curvature gives rise to "sweet spots", where the energy of a particular transition becomes independent of the applied magnetic field strength to first order. These "clock transitions" therefore distinguish themselves through resilience against magnetic field noise and consequently through prolonged coherence times, making them well suited for use as frequency standards.

A selection of of qubits based on clock transitions in the ground state of $^{43}Ca^{+}$ is available. Previous work [LKH+07, BHL+16] made use of the $|M_n = 0\rangle$ to $|M_m = 0\rangle$ transition, which has to be operated close to, but not exactly at, its field independent point of 0 G to be spectrally resolvable. The typically-used field strengths of ~ 2 G lead to a field dependency of $4.8\,\mathrm{Hz/mG}$ [Bal14], but only produce weak Zeeman splittings of less than 1 MHz between adjacent hyperfine states. When driving microwave gates on this qubit, the close spacing of states, together with the generally poor polarization control of the microwave field, leads to off-resonant excitation of neighbouring transitions, causing considerable errors. To keep them at a tolerable level, the power of the microwaves has to be lowered, which significantly slows down the gates. As an example, the duration of a single-qubit gate has to be as long as $\sim 100\,\mu$s to suppress the off-resonant excitation error to $\sim 10^{-4}$ [Bal14], rendering two-qubit operations, which are

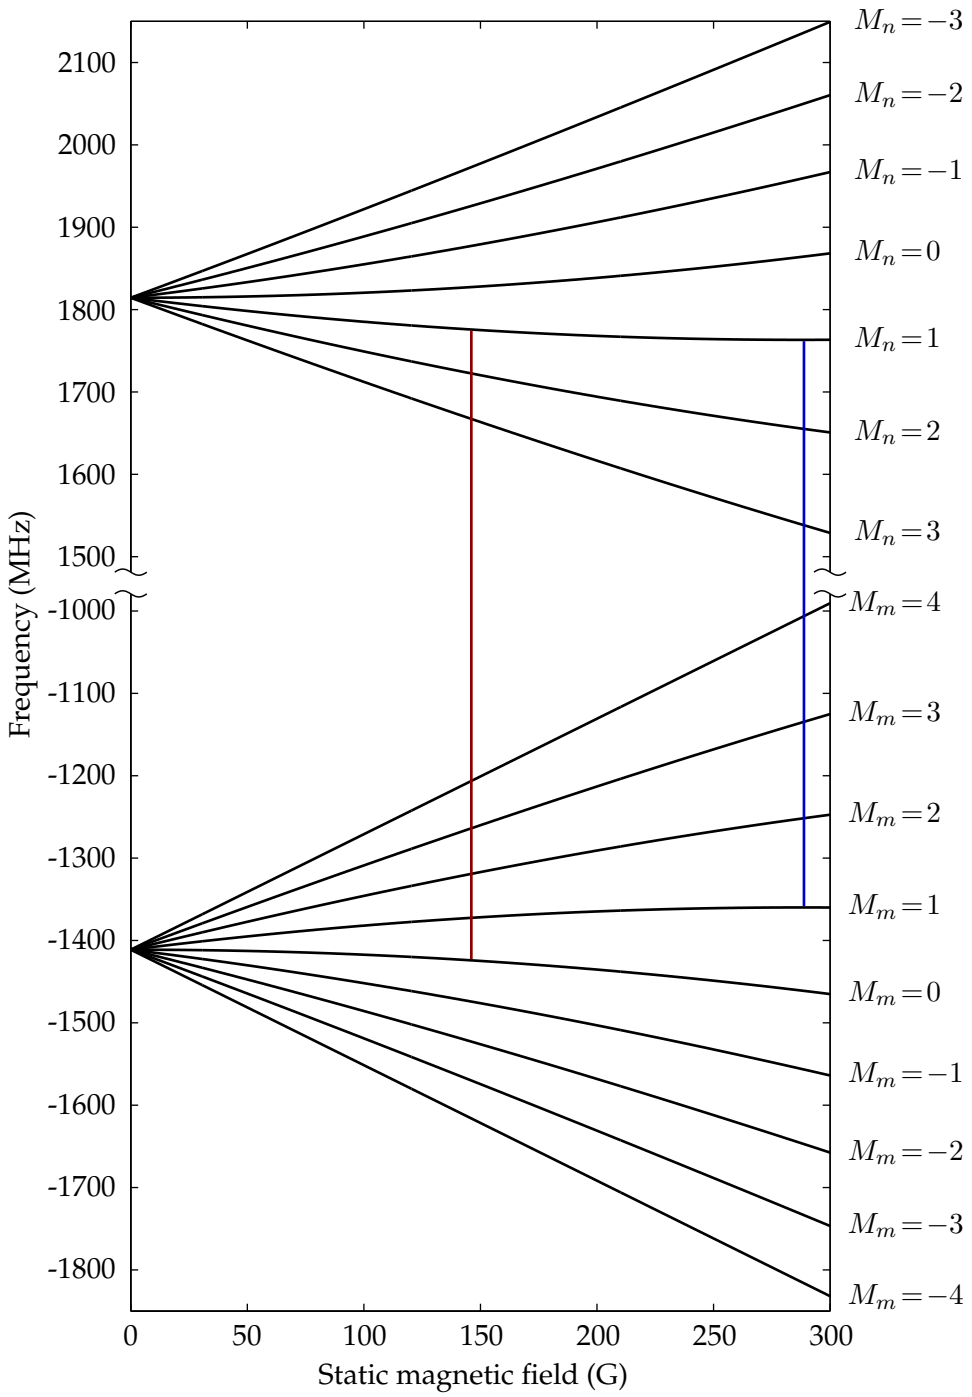

Figure 2.1: Energy splitting of the states in the $4S_{1/2}$ ground level of $^{43}Ca^{+}$ due to the hyperfine and Zeeman interaction as a function of external static magnetic field. The clock transitions at 146 G and 288 G are indicated by the red and blue vertical lines. Within the illustrated range of field strengths, the $|M_n\rangle$ and $|M_m\rangle$ states can still be clearly attributed to their low field $|F = 3, M_F\rangle$ and $|F = 4, M_F\rangle$ counterparts. This is not generally the case for other levels in $^{43}Ca^{+}$. As an example, the states in the $4P_{3/2}$ and $3D_{3/2}$ start crossing at fields between 50 G and 100 G, while the $3D_{5/2}$ level is already well within the strong field regime at 146 G [All11, Har14].

generally orders of magnitude slower, impracticable.

The qubit used in this work is the $|M_n = +1\rangle$ to $|M_m = 0\rangle$ transition, which is magnetic field-independent to first order at 146.0942 G. Such a field gives rise to Zeeman splittings of ~ 50 MHz, as can be seen from Fig. 2.1. This allows for single-qubit gate speeds of a few µs and two-qubit gate speeds at the ms level without causing significant errors due to off-resonant excitation (see section 4.3.4). Furthermore, the ability to operate exactly at the field insensitive point, paves the way for record-breaking coherence times of 50 s [HAB+14], as opposed to 1.2 s measured with the low-field qubit at 1.78 G [LKH+07]. However, these benefits come at the cost of having to cope with an intricate energy level structure that makes the qubit harder to cool, prepare and read out.

Going further up in field strength, the transition between $|M_n = 0\rangle$ and $|M_m = +1\rangle$ becomes field-independent at 146.3314 G. It has similar properties to the qubit transition used in this work, except that it requires an additional microwave π-pulse to be initialised (see section 2.5).

At 287.7827 G the $|M_n = +1\rangle$ to $|M_m = +1\rangle$ transition is a clock transition, offering even higher Zeeman splittings of order 100 MHz, as illustrated in Fig. 2.1. Working at this field would be beneficial, as the associated clock transition is π-polarised and would couple more efficiently to the linearly-polarised microwave radiation produced by our ion trap (see section 5.1). In contrast, the σ-polarised qubits at 146 G are only driven by half the available intensity, making gates preformed on them slower by a factor of $\sqrt{2}$. Furthermore, the 288 G qubit has a 30% stronger transition matrix element (to be introduced in the next section), yielding a further similar gain in gate speed. The disadvantage of this qubit is the lack of a clear cooling scheme, although theory suggests [Jan15] that a similar strategy as in the 146 G case (see section 2.4) should work. However, since to date no experimental efforts have been made to trap $^{43}Ca^+$ at this field, there remains uncertainty about the effectiveness of this approach.

2.2 Magnetic dipole transitions

The states involved in the previously discussed qubits are of the same parity, forbidding electric dipole transitions, and they are both within an S-state, forbidding electric quadrupole transitions. Instead, they are coupled by the magnetic part of the applied radiation $\mathbf{B}(t)$, which interacts with the magnetic moment of the atom $\boldsymbol{\mu}$. Such "magnetic dipole transitions" are described by the operator $-\boldsymbol{\mu} \cdot \mathbf{B}(t)$. The strength with which a transition between the states $|M_n\rangle$ and $|M_m\rangle$ can be driven, is given by the Rabi frequency

$$\Omega \equiv -\frac{1}{\hbar} \langle M_n | \boldsymbol{\mu} \cdot \mathbf{B}(t) | M_m \rangle = \frac{1}{\hbar} \sum_{Q=\pm 1,0} (-1)^{Q+1} \langle M_n | \mu_Q | M_m \rangle \, B_{-Q}(t), \qquad (2.1)$$

which is written in terms of spherical basis components, with $Q \in \{\pm 1, 0\}$, in order to facilitate the evaluation of the matrix elements $\langle M_n | \mu_Q | M_m \rangle$.

The Rabi frequency is the defining factor for the speed with which a gate operation can be implemented on a certain transition. Learning its value requires knowledge of the drive field strength as well of the transition matrix element, the calculation of which is the objective of this section.

Bearing in mind that our qubit exists in the intermediate field regime where neither F, M_F nor M_I, M_J are good quantum numbers, it is not possible to describe its states in terms of the weak field basis $|I, J, F, M_F\rangle$ or the strong field basis $|I, J, M_I, M_J\rangle$. However, by forming appropriate linear combinations of states from either of these two sets, we can construct a new, valid basis for all magnetic fields. The situation is simplified as we are only interested in transitions within the ground level, for which $J = 1/2$, restricting the values of M_J to $\pm 1/2$. Furthermore, the total angular momentum projections M_n and M_m remain good quantum numbers regardless of the field. This means that for each of the two hyperfine manifolds, M_F in the weak field must equal $M_I + M_J$ at strong fields

[Woo80], which implies

$$\begin{aligned} M_n &= M_F = M_I + M_J, \qquad \text{for } F = 3, \\ M_m &= M_F = M_I + M_J, \qquad \text{for } F = 4. \end{aligned} \tag{2.2}$$

Following these considerations, it is convenient to write the new basis as a linear combination of the strong field basis states

$$\begin{aligned} |M_n\rangle &= \alpha_n \left|I, J, M_n + 1/2, -1/2\right\rangle + \beta_n \left|I, J, M_n - 1/2, 1/2\right\rangle, \\ |M_m\rangle &= \alpha_m \left|I, J, M_m + 1/2, -1/2\right\rangle + \beta_m \left|I, J, M_m - 1/2, 1/2\right\rangle. \end{aligned} \tag{2.3}$$

The expansion (Clebsch-Gordan) coefficients $\alpha_{n,m}$ and $\beta_{n,m}$ as well as the energy of each $|M_{n,m}\rangle$ state can be determined by diagonalising the underlying Hamiltonian containing both hyperfine and Zeeman interactions

$$\begin{aligned} H_{\text{hfs}} &= A\mathbf{I}\cdot\mathbf{J} - \boldsymbol{\mu}\cdot\mathbf{B}_{\text{stat}} \\ &= A\sum_{Q=\pm1,0}(-1)^Q I_Q J_{-Q} + \frac{1}{\hbar}g_J\mu_B J_0 B_{\text{stat}} - \frac{1}{\hbar}g_I\mu_N I_0 B_{\text{stat}}, \end{aligned} \tag{2.4}$$

where $\mathbf{B}_{\text{stat}} = B_{\text{stat}}\hat{\mathbf{e}}_\mathbf{0}$ is the static magnetic field that sets the quantisation axis along $\hat{\mathbf{e}}_\mathbf{0} = \hat{\mathbf{z}}$, and where we have decomposed the magnetic moment of the ion into its spherical basis components

$$\boldsymbol{\mu} = \sum_{Q=\pm1,0}\mu_Q = \frac{1}{\hbar}\sum_{Q=\pm1,0}\left(-\mu_B g_J J_Q + \mu_N g_I I_Q\right). \tag{2.5}$$

Here, μ_B is the Bohr magneton, μ_N is the nuclear magneton, $g_J = 2.00225664(9)$ [TPR$^+$03] is the Landé g-factor, and $g_I = 1.315348(9)$ [Har14] is the nuclear g-factor. The eigenvalues of Eq. 2.4, that is, the energies of the $|M_{n,m}\rangle$ states, are given by the Breit-Rabi formula [Woo80]. The corresponding eigenvectors specify the Clebsch-Gordan coefficients $\alpha_{n,m}$ and $\beta_{n,m}$, which can be found tabulated in appendix A for static magnetic fields of 146 G and 288 G.

Having found a suitable basis for describing the qubit states in the intermedi-

ate field, we can use the angular momentum operator relations

$$J_{\pm 1}\left|I,J,M_I,M_J\right\rangle = \mp\frac{1}{\sqrt{2}}\hbar\sqrt{(J\mp M_J)(J\pm M_J+1)}\left|I,J,M_I,M_J\pm 1\right\rangle, \quad (2.6)$$

$$J_0\left|I,J,M_I,M_J\right\rangle = \hbar M_J\left|I,J,M_I,M_J\right\rangle, \quad (2.7)$$

$$I_{\pm 1}\left|I,J,M_I,M_J\right\rangle = \mp\frac{1}{\sqrt{2}}\hbar\sqrt{(I\mp M_I)(I\pm M_I+1)}\left|I,J,M_I\pm 1,M_J\right\rangle, \quad (2.8)$$

$$I_0\left|I,J,M_I,M_J\right\rangle = \hbar M_I\left|I,J,M_I,M_J\right\rangle, \quad (2.9)$$

to determine the transition matrix elements. For σ_+ polarized transitions, the matrix elements are nonzero if $Q = M_n - M_m = +1$, giving

$$\begin{aligned}\langle M_n|\mu_{+1}|M_m\rangle =& \frac{1}{\sqrt{2}}\Big\{\mu_B g_J \alpha_m\beta_n \\ &-\mu_N g_I\left(\alpha_m\alpha_n\sqrt{16-M_n^2}+\beta_m\beta_n\sqrt{15+2M_n-M_n^2}\right)\Big\}.\end{aligned} \quad (2.10)$$

For σ_- polarized transitions, they are nonzero if $Q = M_n - M_m = -1$, giving

$$\begin{aligned}\langle M_n|\mu_{-1}|M_m\rangle =& -\frac{1}{\sqrt{2}}\Big\{\mu_B g_J \alpha_n\beta_m \\ &+\mu_N g_I\left(\beta_m\beta_n\sqrt{16-M_n^2}+\alpha_m\alpha_n\sqrt{15-2M_n-M_n^2}\right)\Big\}.\end{aligned} \quad (2.11)$$

Finally, for π polarized transitions, the transition matrix elements are nonzero if $Q = M_n - M_m = 0$, giving

$$\begin{aligned}\langle M_n|\mu_0|M_m\rangle =& \frac{1}{2}\mu_B g_J\left(\alpha_n\alpha_m-\beta_n\beta_m\right) \\ &+\mu_N g_I\left\{\alpha_n\alpha_m(M_n+1/2)+\beta_n\beta_m(M_n-1/2)\right\}.\end{aligned} \quad (2.12)$$

Their values are tabulated in appendix A for the static magnetic fields of interest.

Finally, we also write down the expressions for

$$\begin{aligned}\langle M_n|\mu_0|M_n\rangle =& \frac{1}{2}\mu_B g_J\left(\alpha_n^2-\beta_n^2\right) \\ &+\mu_N g_I\left\{\alpha_n^2(M_n+1/2)+\beta_n^2(M_n-1/2)\right\},\end{aligned} \quad (2.13)$$

$$\begin{aligned}\langle M_m|\mu_0|M_m\rangle =& \frac{1}{2}\mu_B g_J\left(\alpha_m^2-\beta_m^2\right) \\ &+\mu_N g_I\left\{\alpha_m^2(M_m+1/2)+\beta_m^2(M_m-1/2)\right\}.\end{aligned} \quad (2.14)$$

since they will appear in the form of $\langle M_n|\mu_0|M_n\rangle - \langle M_m|\mu_0|M_m\rangle$ in section 4.1.4. It is important to realise, that for $\alpha_n \approx \alpha_m$ and $\beta_n \approx \beta_m$, as is the case in clock transitions, these matrix elements vanish.

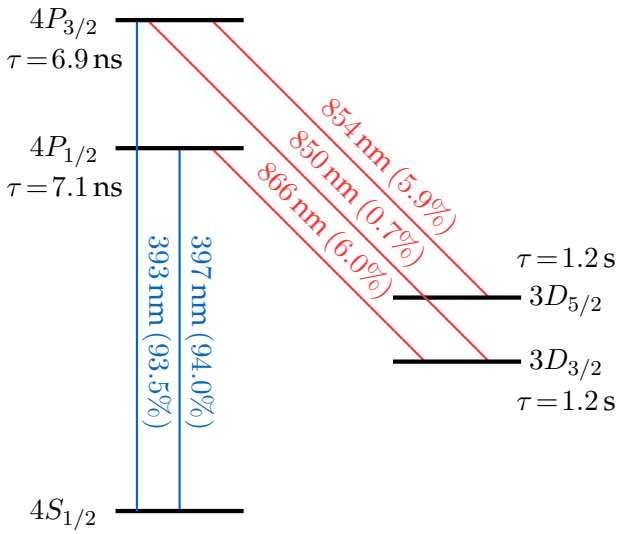

Figure 2.2: Energy levels and dipole transitions needed for quantum information experiments in basic Ca^+. The branching ratios and state lifetimes τ are taken from [Szw09] and rounded to one decimal place.

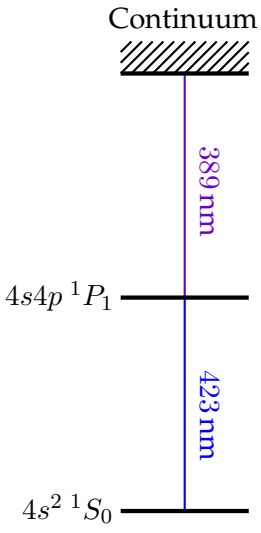

Figure 2.3: Energy levels in neutral calcium relevant for photo-ionization.

2.3 Photo-ionization

In most ion trap experiments, ions are created from a thermal beam of neutral atoms, emanating from an atom source, which in our setup is a resistively heated calcium oven. The beam passes through the trapping region, where the atoms are subjected to laser radiation that knocks one of their electrons off and allows some of them to be captured by the electric fields of the Paul trap. Since natural abundance atom sources only contain 0.135% ^{43}Ca and the maximum level of isotopic enrichment is of order 80% [LRH+04], reliable trapping of pure $^{43}Ca^+$ crystals requires an isotope-selective photo-ionization procedure.

As originally described in [GRB+01, LRH+04] and illustraded in Fig. 2.3, the method used in this experiment utilises two lasers. The first has a wavelength of 423 nm and is resonant with the $4s^2\ ^1S_0$ - $4s4p\ ^1P_1$ transition in neutral calcium. It has a narrower linewidth than the associated 612 MHz [LRH+04] isotope shift between ^{40}Ca and ^{43}Ca, provided a Doppler-free geometry is used. When appropriately tuned, the 423 nm laser can therefore selectively excite the desired

isotope to the $4s4p\ ^1P_1$ level, from where the second laser, with a wavelength of 389 nm, provides enough energy to reach the continuum.

The scheme works in spite of the presence of the hyperfine and Zeeman structure in $^{43}Ca^+$ at 146 G, as explained in section 2.4 of [Har14].

2.4 Laser cooling in the intermediate field

The purpose of laser cooling ions in the context of quantum information experiments is twofold:

First, it reduces the kinetic energy of hot ions after they e.g. were loaded into the trap or suffered a collision with a background gas particle. In general, Doppler cooling fulfils this function, while also extending the ions' lifetime in the Paul trap to typically a few hours, depending on the type of trap and the quality of the surrounding vacuum. When temperatures below the Doppler limit are required, e.g. when carrying out high fidelity two-qubit gates, Doppler cooling is usually followed by resolved sideband techniques that allow for the ground state of the trapping potential to be reached with high probability.

Second, laser cooling is part of the qubit readout method, which relies on removing one of the two qubit states from the Doppler cooling cycle, thus making it possible to discriminate between them by detecting the presence or absence of fluorescence when cooling is initiated. The objective in this case is not to achieve low ion temperatures but, instead, to yield high fluorescence, which generally requires different experimental parameters. A suitable cooling scheme therefore has to be readily convertible between these two configurations.

2.4.1 Doppler cooling

Doppler cooling relies on irradiating hot ions with light that is red-detuned from one of their transitions, causing them to continuously absorb more photons when moving against the radiation, due to the Doppler effect. The associated

momentum transfer effectively slows the atoms down in that particular direction, as the repeated recoil from spontaneous photon emission, which follows every absorption event, is isotropic. Thanks to the restoring forces of the Paul trap, trapped ions can be Doppler cooled using a single beam, provided it has a component along all three principal axes of the confining potential.

In order to slow down ions with a high starting temperature, it is common to employ a dipole transition between the ground state and a short-lived excited state, since this ensures a high scattering rate from the cooling beam. When working with $^{40}Ca^{+}$, we use the $4S_{1/2}$ to $4P_{1/2}$ transition at $397\,\text{nm}$ together with a repumper at $866\,\text{nm}$, which recovers the population occasionally decaying from the $4P_{1/2}$ to the metastable $3D_{3/2}$ state. The corresponding level diagram including lifetimes and branching ratios can be found in Fig. 2.2.

In $^{43}Ca^{+}$, the ground state hyperfine splitting is two orders of magnitude larger than the $\sim 22\,\text{MHz}$ natural linewidths of the transitions linking it to $4P_{1/2}$. Consequently, the 397 nm beam has to be equipped with two frequency components, each connecting one of the two hyperfine ground levels to the $4P_{1/2}(F = 4)$ state. This method is proven to work at low static magnetic fields of up to 5 G [LKH+07, BKRB08], where the frequency spread of the Zeeman structure remains smaller or comparable to the linewidths of the 397 nm transitions, thus keeping the cooling cycle closed.

However, this is not the case at 146 G, where the Zeeman splittings are of order 50 MHz and the 397 nm laser cannot access all levels in the ground state any more. Closing the cooling cycle by adding more frequency components is impracticable, as the system now consists of 16 spectrally resolved Zeeman states in the $4S_{1/2}$ level, further 16 in $4P_{1/2}$ and another 32 in $3D_{3/2}$. We therefore resort to a smarter solution that utilises the same resources as in the low field scenario, but with carefully chosen laser polarizations, frequencies and intensities. In this way, the number of states involved in the cooling cycle can be confined to a

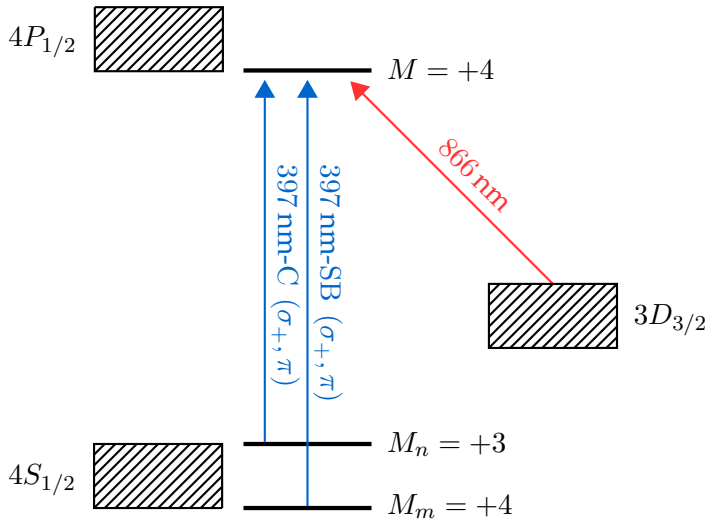

Figure 2.4: Lasers used for Doppler cooling $^{43}Ca^+$ at 146 G. The bulk of the hyperfine structure has been lumped into shaded boxes, with only the most relevant states explicitly shown. The 397 nm beam has two frequency components in order to access both ground state hyperfine manifolds. Their polarization is predominantly σ_+, confining population to the few highest M-states, while the π part keeps it cycling, mainly between $4S_{1/2}\,|M_m = +4\rangle$ and $4P_{1/2}\,|M = +4\rangle$. The 866 nm laser serves to repump population that decayed to the $3D_{3/2}$ level.

manageable amount.

As illustrated in Fig. 2.4, the general idea is to set the 397 nm beam polarization to be predominantly σ_+ with a small amount of π, and tune its two frequency components so that one is close to the $4S_{1/2}\,|M_n = +3\rangle$ to $4P_{1/2}\,|M = +4\rangle$ transition, and the other is close to the $4S_{1/2}\,|M_m = +4\rangle$ to $4P_{1/2}\,|M = +4\rangle$ transition. Here M is the total angular momentum, which describes the $4P_{1/2}\,|M = +4\rangle$ state unambiguously at any static magnetic field strength. The strong σ_+ component's purpose is to optically pump population towards the positive end of the ground state manifold, while the π component prevents it from getting stuck there by primarily driving the $4S_{1/2}\,|M_m = +4\rangle$ to $4P_{1/2}\,|M = +4\rangle$ transition, which acts as the main cooling cycle.

The optimum frequency separation of the two 397 nm components was found theoretically to be 2.930 GHz, and is provided by an EOM. Whenever a distinction between the two beams becomes necessary, we will refer to the carrier as

397 nm-C, while the first higher-frequency sideband, set to have 85% of the intensity of the carrier, will be called 397 nm-SB. The first sideband at lower frequency (and all higher orders) are too far detuned to produce any significant effect on the ion.

In terms of repumping, it is still sufficient to use a single 866 nm beam, as the frequency spread of all states in the $3D_{3/2}$ level is only $\sim$1 GHz and we can retain access to them by saturation broadening the transitions back to $4P_{1/2}$ with a higher laser intesity. It is worth noting that the same approach would not work on the main 397 nm transition, since saturating it deteriorates the cooling performance.

Although this concept seems simple, determining suitable laser intensities and detunings is a non-intuitive task. The 64 states involved in the cooling process give rise to a complicated excitation spectrum, including numerous dark resonances (to be introduced shortly, in section 2.4.2). In order to find an optimal set of parameters, numerical modelling of the complete $4S_{1/2}$-$4P_{1/2}$-$3D_{3/2}$ system using optical Bloch equations was necessary. Meaningful results could only be obtained after taking into account the time-lag associated with repumping population from the $3D_{3/2}$ level relative to the period of the ion's secular motion in the trap. A discussion of this "dynamic effect" and details about modelling $^{43}Ca^{+}$ at 146 G can be found in [Jan15].

Guided by these numerical calculations, we explored the parameter space experimentally, with the results discussed in section 6 and published in [AHS+16]. As explained in the following section, the chief insight was that with appropriate laser settings, the above mentioned concept can be used to exploit the special properties of a particular dark resonance, resulting in low enough temperatures for resolved sideband cooling techniques to take over.

2.4.2 Dark resonance cooling at 146 G

The minimum temperature achievable with Doppler cooling is governed by two main factors. First, the rate at which the ions scatter photons from the cooling beam. This process, although crucial for high-fidelity readout, leads to recoil heating. Second, the change in scattering rate as a function of the Doppler shift, which determines how many more absorption events occur for ions that move against the radiation compared with those moving with the radiation. Consequently, effective cooling requires ions to be operated at a point where their fluorescence profile with respect to the radiation's frequency has a strong gradient, but a low absolute value. The latter condition is incompatible with high fidelity readout, making it necessary to establish a separate set of laser parameters for obtaining maximum fluorescence, while ensuring that it can straightforwardly be switched to the configuration for efficient cooling.

To provide a benchmark against which to compare the performance of the cooling scheme, we initially look at $^{40}Ca^{+}$ under the simplified assumptions that the $3D_{3/2}$ level can be ignored, and that the only relevant heating effect is recoil from spontaneously emitted photons. The fluorescence curve underlying this two-level system is a Lorentzian, with optimal cooling occurring half way up its red detuned slope, where the gradient is steepest. Provided that the 397 nm transition is not saturated, the minimum temperature, called the Doppler limit, in this case is ~ 0.5 mK [Jan15].

Returning to $^{43}Ca^{+}$ at 146 G, the excitation spectrum is more complicated and includes points where both the fluorescence is low and the gradients are steeper than that of the Lorentzian discussed above. These points correspond to two-photon dark resonances [SSB+92], which occur when a Λ-type configuration of three levels is subjected to two radiation fields with equal detunings from their respective resonances. Under such conditions, the ion is optically pumped into a dark state, that is, a state that does not absorb any photons from the applied

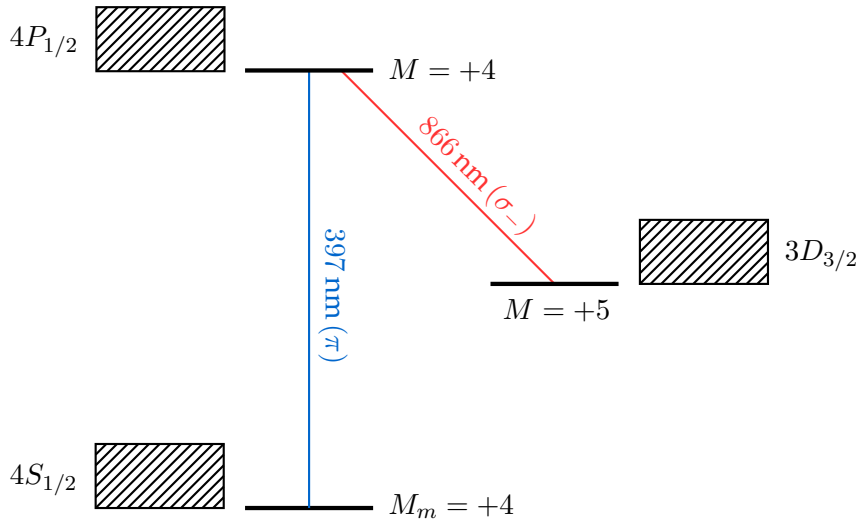

Figure 2.5: Energy level diagram showing the Λ-system underlying the dark resonance used for sub-Doppler cooling in $^{43}Ca^{+}$ at intermediate fields.

fields and does not emit any fluorescence.

In $^{43}Ca^{+}$, the 16 states in the ground level together with the 16 states in $4P_{1/2}$ and the 32 states in $3D_{3/2}$ give rise to numerous Λ-systems, driven by 397 nm and 866 nm light. The associated dark resonances can be made visible upon scanning e.g. the 866 nm laser detuning across the $4P_{1/2}$-$3D_{3/2}$ line.

The specific dark resonance used in this work is spanned by a π transition from $4S_{1/2}$ $|M_m = +4\rangle$ to $4P_{1/2}$ $|M = +4\rangle$, and the σ_- transition from $3D_{3/2}$ $|M = +5\rangle$ to $4P_{1/2}$ $|M = +4\rangle$, as shown in Fig. 2.5. It is suitable for cooling since, unlike most other dark resonances present in the fluorescence profile, its underlying Λ-system is not coupled to other states in the ground-state manifold. This means that, given that the 397 nm beam is $1/3$ π and $2/3$ σ_+ polarised, and the 866 nm beam is $1/2$ σ_+ and $1/2$ σ_- polarised (table 5.3), the only way for population to leave the $4S_{1/2}$ $|M_m = +4\rangle$ or $3D_{3/2}$ $|M = +5\rangle$ states is via $4P_{1/2}$ $|M = +4\rangle$. Consequently, when the resonance condition is met, i.e. when the detunings of the 397 nm and 866 nm lasers are equal, the fluorescence drops sharply to almost zero, with the depth of the feature being limited by the ion's motion and the finite laser linewidths.

Cooling is most effective when the 397 nm laser is tuned a few MHz red of the $4S_{1/2}\,|M_m = +4\rangle$ to $4P_{1/2}\,|M = +4\rangle$ transition and the 866 nm laser is set so the point of operation is close to the bottom of the dark resonance, on the lower frequency slope. Temperatures of 0.3 mK, slightly below the Doppler limit can be reached provided that the two laser beams are counter-propagating, leading to opposite Doppler shifts and thus a bigger change in fluorescence with the ion's motion. Switching to the settings that give maximal fluorescence, requires detuning the 866 nm beam by 360 MHz to push the point of operation out of the dark resonance, while at the same time increasing the intensities in both beams by about one order of magnitude.

A temperature of $T = 0.3\,\mathrm{mK}$ corresponds to the ion being in a thermal state of motion with mean occupation number $\bar{n} = k_\mathrm{B}T/\hbar\omega \sim 2$, assuming the mode frequency ω is $\sim 3\,\mathrm{MHz}$, as is the case in our trap. However, high-fidelity two-qubit gates typically require temperatures around $\bar{n} \sim 0.1$ (see section 4.3). To reach such temperatures, the dark resonance cooling procedure is followed by a resolved sideband cooling step[1], which is the subject of the following section. We will return to the topic of dark resonance cooling in chapter 6, where we present the measurements underlying the above made claims, and where we provide specific values for laser detunings and intensities.

2.4.3 Resolved-sideband cooling

While Doppler cooling can be applied to a cloud of free or trapped atoms equally, resolved-sideband cooling is a technique specific to atoms confined in a trap. A

[1]Since with the help of dark resonance cooling the ion temperature can be brought down to $\sim 0.3\,\mathrm{mK}$, one would hope that increasing the motional mode frequency by going to a stiffer trap might allow for thermal occupations of $\bar{n} \sim 0.1$ to be reached, thus eliminating the necessity of resolved sideband cooling. However, reaching such thermal occupations with dark resonance cooling only requires our current ($\sim 3\,\mathrm{MHz}$) mode frequencies to be raised by a factor of ~ 20. Since in this work we use radial modes, the frequencies of which scale as the trapping RF voltage amplitude [WMW+98], we would have to increase the latter by a similar factor to achieve the desired mode frequency change. As our current trap RF voltage is 72 V, such a change would most likely result in the electrical breakdown of the trap.

description of the underlying theory can be found in e.g. [WI79, LBMW03].

The essential ingredient is, again, the Doppler effect, which causes harmonically confined ions to see incoming radiation with frequency ω_0 in a certain direction, as being modulated with frequencies ω_j of their j oscillatory modes of motion along that direction. In the ions' frame of reference, monochromatic radiation is therefore comprised of a carrier at ω_0 and equally spaced sidebands appearing at $\omega_0 \pm \sum_j m\omega_j$, where m is an integer. Conversely, the excitation spectrum of the atoms in the laboratory frame has a similar structure, but each frequency component is broadened due to the limited lifetimes of the electronic states involved. The strength of the sidebands is determined by the modulation index, which in this case is the product kq of the radiation's wave vector k with the ion's motional mode amplitude q.

Looking back at the previously discussed Doppler cooling procedure, the natural width of the 397 nm transition (22 MHz) was about one order of magnitude larger than the ions' motional mode frequencies of typically a few MHz. Thus, sidebands were not resolved and the resulting cooling process resembled that of free atoms, with the trap only serving as a means to localise them and to allow cooling with a single beam.

The opposite situation arises when employing a transition that has a narrower natural linewidth than the motional mode frequencies of the ions in the trap, such as a dipole forbidden or Raman transition. We can then tune a laser to one of the j^{th} mode's lower frequency (red) sidebands with respect to the carrier, and induce the atoms to absorb predominantly photons with energy $\hbar\,(\omega_0 - m\omega_j)$. Provided that the main decay channel is via the carrier, each absorption event is succeeded by photon reemission with energy $\hbar\omega_0$, in a random direction. The excess amount of energy $\hbar m\omega_j$ that is released with every excitation and emission cycle corresponds to the ions' j^{th} vibrational mode losing m quanta of motion. When applied repeatedly, this process is therefore capable of cooling the mode

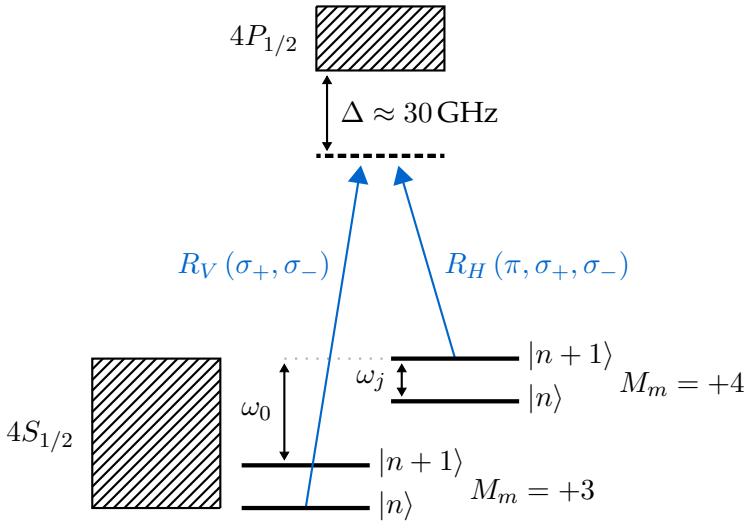

Figure 2.6: Continuous resolved sideband cooling scheme employed in this work, involving the horizontally polarised Raman laser R_H and vertically polarised Raman laser R_V. The energy of the ions' collective motional mode is described by the ladder of harmonic oscillator eigenstates $|n\rangle$ with equidistant spacing $\omega_j \approx 2\pi \times 3\,\mathrm{MHz}$ for typical radial mode frequencies in our trap. The Raman laser frequencies are set to $R_V = R_H + \omega_0 + \omega_j$, where $\omega_0 = 2\pi \times 57.480\,\mathrm{MHz}$, and drive the first red sideband from $4\mathrm{S}_{1/2}\,|M_m = +4, n+1\rangle$ to $4\mathrm{S}_{1/2}\,|M_m = +3, n\rangle$ via a virtual intermediate state that is detuned from $4\mathrm{P}_{1/2}$ by $\Delta \approx -30\,\mathrm{GHz}$. Beams that are not indicated in the diagram are the σ_+-polarized 397 nm and the 866 nm laser, which pump population from $4\mathrm{S}_{1/2}\,|M_m = +3, n\rangle$ back to $4\mathrm{S}_{1/2}\,|M_m = +4, n\rangle$ and thereby close the cooling cycle.

down to its ground state, where the red sidebands vanish.

In order for spontaneous decays to proceed predominantly via carrier transitions and not sidebands of order $m > 1$, we need to ensure weak enough modulation. Suitable conditions prevail if the modulation index lies within the so called "Lamb-Dicke regime", where $\langle (kq)^2 \rangle \ll 1$. To reach it, it is common to employ an initial Doppler cooling step that reduces the motional mode amplitude q in the run-up to the resolved sideband cooling stage, causing the modulation index to drop to the desired level.

In our experiment, we first employ the previously explained dark resonance method to achieve the Lamb-Dicke regime. For the subsequent sideband cooling stage, we use two 397 nm Raman beams, R_H and R_V, where the suffixes denote horizontal and vertical linear polarisation. They drive a stimulated Raman tran-

sition from $|M_m = +4\rangle$ to $|M_m = +3\rangle$, at 57.480 MHz, via a virtual intermediate state that is detuned from $4P_{1/2}$ by $\Delta \approx -30$ GHz, as illustrated in Fig. 2.6. Our beam geometry is shown in Fig. 5.3 of chapter 5, and only gives a projection along the radial direction, which means that the axial modes remain Doppler cooled only. We use 400 μW in each beam, where the $1/e^2$ Gaussian beam diameters are 57 μm for R_H, and 66 μm for R_V, corresponding to 336 and 252 saturation intensities [Szw09], respectively.

Since neither of the two above mentioned states in the ground level spontaneously decays to the other one, we have to artificially introduce a decay channel to open up a continuous cooling cycle. To do so, we apply a weak, σ_+ polarized 397 nm beam and an 866 nm repumper simultaneously with the two Raman beams. The 397 nm beam has two frequency components, each connecting one of the two hyperfine ground state manifolds to $4P_{1/2}$. It contains 0.7 μW of total power over a $1/e^2$ Gaussian beam diameter of 37 μm, corresponding to ~ 1.4 saturation intensities. The 866 nm power is noncritical and we simply use the same power setting as for the dark resonance cooling.

The resulting continuous sideband cooling cycle involves population being transferred from $4S_{1/2}$ $|M_m = +4\rangle$ to $4S_{1/2}$ $|M_m = +3\rangle$ via a first ($m = 1$) red Raman sideband transition, followed by an excitation to the $4P_{1/2}$ $|M = +4\rangle$ state by the σ_+ polarized 397 nm light. From there, the only decay possibility that is not reverted by either the 866 nm or 397 nm beams is the return to $4S_{1/2}$ $|M_m = +4\rangle$, which closes the cycle. Population that is recovered from $3D_{3/2}$ may end up in a different state than $|M = +4\rangle$ in $4P_{1/2}$, thus requiring more time to get optically pumped back to $4S_{1/2}$ $|M_m = +4\rangle$.

2.5 Qubit initialisation

In order to prepare the clock qubit in $^{43}Ca^+$ at 146 G, we first optically pump population to the $4S_{1/2}$ $|M_m = +4\rangle$ state using a pulse of bichromatic, σ_+ polarized

397 nm beam, the same that was already introduced when discussing sideband cooling. However, here we apply it with a power of 5 μW, corresponding to 9.9 saturation intensities, alongside our usual 866 nm repumper for a duration of typically a few hundred μs.

The fidelity with which the $|M_m = +4\rangle$ state can be prepared in this way, is limited by imperfect polarization of the 397 nm beam, causing some population to distribute itself among the highest few M-states in the ground level, usually $M \geq +2$. Most of it can be recovered by applying a similar 397 nm pulse as above, but shortened to 30 μs and with only the frequency component addressing the upper ground state manifold enabled, that is, with the EOM sideband turned off. This clears out population left in $|M_n\rangle$ states and transfers it to the lower ground level, preferentially to states with $M_m > M_n$. We then apply microwave π-pulses resonant with the $|M_m = +3\rangle$ to $|M_n = +3\rangle$ transition, labelled D in Fig. 2.7, and with the $|M_m = +2\rangle$ to $|M_n = +2\rangle$ transition, labelled E in Fig. 2.7, before using another single-frequency 397 nm pulse. This cycle can be repeated as often as necessary to reclaim the majority of the lost population. The π-times of the D and E pulses are 3.1 μs and 2.2 μs, respectively. In previous work [HAB+14, Har14] conducted on the same ion trap, this microwave enhanced optical pumping technique was shown to provide an estimated reduction in the error of preparing $4S_{1/2}$ $|M_m = +4\rangle$ of one order of magnitude, pushing it below the 10^{-4} level.

For some calibration experiments, it is sufficient to prepare the $|M_m = +4\rangle$ state, in which case the initialisation procedure finishes here. However, if experiments are to be carried out on the clock transition, the procedure further involves the transfer of population to one of the qubit states by means of a sequence of microwave π-pulses $\{A, B, C\}$, and optionally Q. The underlying transitions are indicated in Fig. 2.7, while the π-times used in our experiment are 1.5 μs for A, 4 μs for B, and 1.6 μs for C and Q.

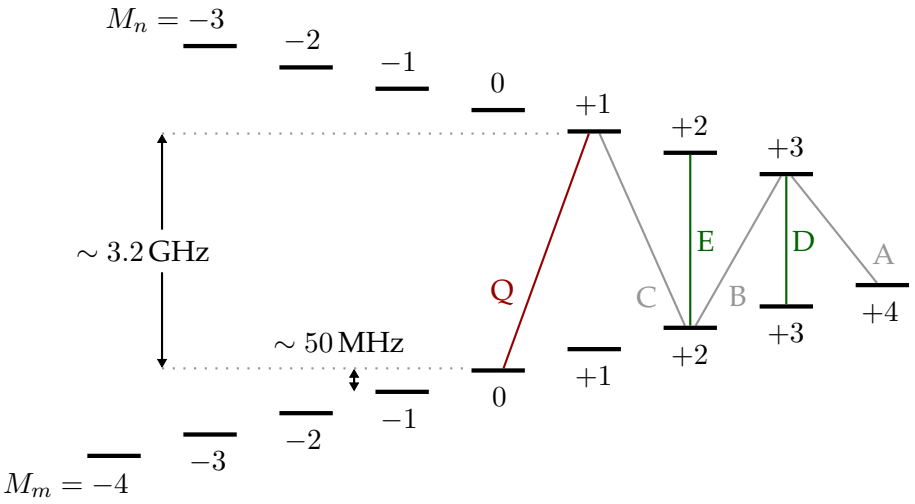

Figure 2.7: Hyperfine and Zeeman structure (not to scale) of the $4S_{1/2}$ ground level of $^{43}Ca^{+}$ at 146 G. The clock transition Q, the transfer pulses A, B and C used for state preparation, as well as the D and E pulses for enhanced optical pumping, are indicated.

2.6 Qubit readout

The state of qubits encoded in ions with low lying D states can be measured with the "electron shelving" technique, proposed by H. Dehmelt [Deh75]. It involves driving a cycling transition between one of the qubit states and an auxiliary state with a resonant laser pulse, while at the same time detecting a fraction of the scattered photons, typically $\sim 0.1\%$, with e.g. a CCD camera or photomultiplier tube. The second qubit state is mapped onto a "shelf" state, which is not part of the cycling transition, prior to applying the laser pulse. Thus the presence or absence of fluorescence can be used to distinguish between the internal states of the qubits.

In $^{43}Ca^{+}$ at 146 G, we use the Doppler-cooling transition as our cycling transition and the $3D_{5/2}$ level as the shelf. For all experiments described in the remainder of this thesis, we choose the population in the $|M_n = +1\rangle$ qubit state to be shelved, and that in $|M_m = 0\rangle$ to remain in the ground level and take part in the cycling transition.

The shelving procedure begins with the transfer of population from $|M_n = +1\rangle$

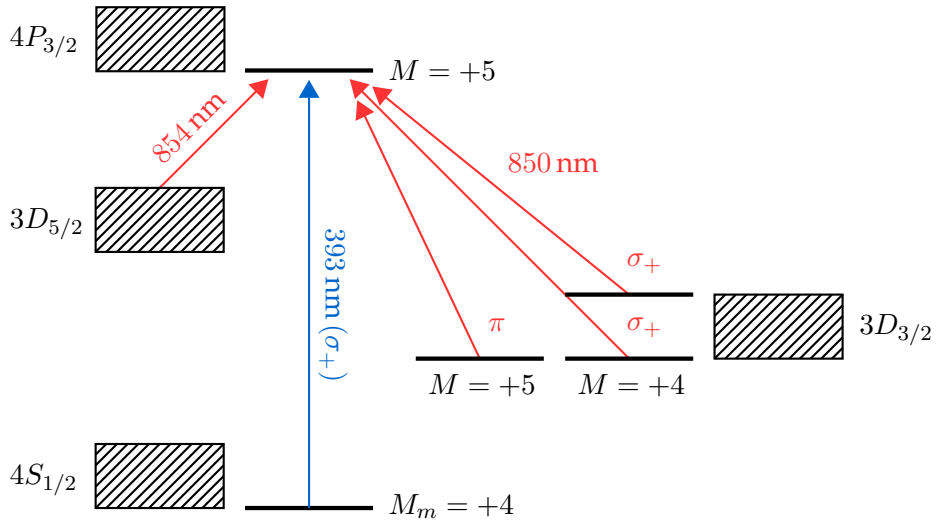

Figure 2.8: Lasers required for shelving population from $4S_{1/2}\,|M_m = +4\rangle$ to $3D_{5/2}$ in $^{43}Ca^+$ at intermediate fields. The 850 nm repumper consists of three frequency components combined into one σ_+ and one π polarized beam, which are pulsed on sequentially together with the 393 nm light. Deshelving after the readout process involves an 854 nm pulse.

to $|M_m = +4\rangle$ and $|M_m = 0\rangle$ to $|M_n = +1\rangle$ by applying the microwave π-pulse series $\{C, B, A, Q\}$, as illustrated in Fig. 2.7. We then apply a σ_+ polarized 393 nm laser pulse, resonant with the $4S_{1/2}\,|M_m = +4\rangle$ to $4P_{3/2}\,|M = +5\rangle$ transition. Population residing in $4S_{1/2}\,|M_n = +1\rangle$ is unaffected by the pulse as it is ~ 3 GHz detuned, while that in $4S_{1/2}\,|M_m = +4\rangle$ gets excited to $4P_{3/2}\,|M = +5\rangle$. From there, it can either decay to $3D_{5/2}$, $3D_{3/2}$ or back to $4S_{1/2}\,|M_m = +4\rangle$, with the associated branching ratios given in Fig. 2.2.

Decays to the $3D_{3/2}$ level can only proceed via σ_+-transitions to the two states with $M = 4$, or via a π-transition to the state with $M = 5$. Population lost in this way can be reclaimed using an 850 nm laser containing all three frequency components. As explained in [MSW+08, Har14], repumping is most effectively incorporated into the shelving process, when the laser pulses involved are applied in a sequence of the form

$$\left\{\left(393\,\text{nm}, 850\,\text{nm}\ \sigma_+, 850\,\text{nm}\ \pi\right)^i, 393\,\text{nm}\right\}, \tag{2.15}$$

where i signifies the number of repeats, typically $i = 6$ for this work, and

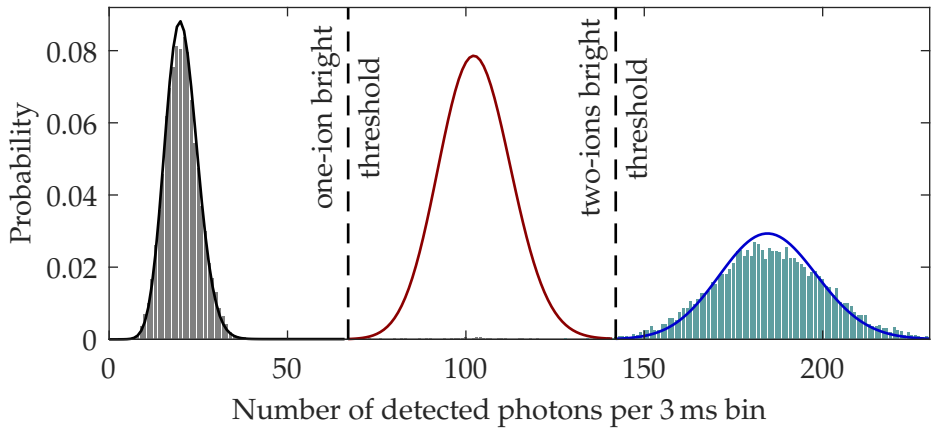

Figure 2.9: Photon-count histograms measured using two $^{43}Ca^{+}$ ions with a fluorescence detection bin time of 3 ms. The histograms have been converted into probability distributions, providing the probability of detecting a certain number of photons when both ions are dark (grey histogram) or bright (light blue histogram). The solid black and blue curves represent Poisson distributions with the same means as the corresponding histograms. The red solid curve is the predicted Poisson distribution for one ion being bright, and the other dark. These curves do not take into account spontaneous decay from the metastable $3D_{5/2}$ shelf, which would extend the black and red distributions towards higher photon numbers [MSW+08]. By setting thresholds between the histograms, we can distinguish different qubit states in our detection process.

$\{850\,\text{nm}\ \sigma_+\}$ contains two frequency components, separated by $\sim 230\,\text{MHz}$. The infidelity of this technique is fundamentally limited to $\sim 1 \times 10^{-4}$ due to the tradeoff between off-resonant shelving of population in $4S_{1/2}\ |M_n = +1\rangle$ to $4P_{3/2}$, and spontaneous decay of the metastable shelf state.

In this work, we set the 393 nm pulse time to 10 μs, and its power to 35 nW, which corresponds to 0.0658 saturation intensities when taking into account its $1/e^2$ Gaussian beam diameter of 37 μm. Both 850 nm pulses are 6 μs long and their power is set to 20 μW, which is equivalent to 119 saturation intensities when factoring in their beam diameter of 66 μm.

Successful shelving is followed by a 3 ms of Doppler-cooling and simultaneous fluorescence detection using a photomultiplier (PMT). The collected PMT counts over many repeats of the experiment give rise to photon-count histograms, as illustrated for the two-ion case in Fig. 2.9. We discriminate between qubit states

by setting thresholds between histograms. More specifically, when detecting photon numbers below the "one-ion bright" threshold, both ions are registered as dark, i.e. shelved. When detecting photon numbers above the "two-ions bright" threshold, both ions are registered as bright, i.e. fluorescing. Photon numbers between thresholds are therefore ascribed to one ion being bright, and the other dark.

When carefully optimised, the above readout method was shown [HAB+14, Har14] to give an error of only $\sim 5 \times 10^{-4}$ for a single ion in the same trap as was used in this work. However, the experiments presented in this thesis have a higher readout error, for reasons that will be discussed in section 7.5.

3

Microwave driven quantum logic

This chapter is aimed at providing a conceptual introduction to single- and two-qubit quantum logic, in preparation for the more detailed theoretical treatment given in the subsequent chapter. We focus on the gate scheme used in this work and only briefly touch upon related methods for the sake of completeness.

3.1 A universal set of gates in a trapped ion system

In classical computation, a logic gate is an operation that maps the state of its input bits onto a state of output bits by following a truth table. The quantum mechanical equivalent, i.e. the quantum logic gate, implements a truth table on the basis states of its qubits, by means of a unitary evolution.

A set of quantum gates is classified as universal, if all unitary transformations on arbitrarily many qubits can be expressed as compositions of gates taken from this set. Consequently, any operation that is possible on a universal quantum computer, can be reduced to sequences of these gates. As already mentioned in section 1.2, a universal set requires a complete basis of single-qubit gates and an entangling operation on two qubits, both of which will be introduced in the following.

3.1.1 Single-qubit gates

Single qubit gates are operations that manipulate the state of a single qubit, allowing for the generation of superposition states. They can be implemented by means of driving Rabi oscillations between the two qubit levels with resonant radiation pulses.

When visualised on the Bloch sphere, where the north and south pole are aligned along the z-axis and represent logical $|0\rangle$ and $|1\rangle$ respectively, single qubit gates can be illustrated as rotations of the Bloch vector. The relevant parameters are the phase of the radiation, which sets the axis of rotation, and the strength and duration of the pulse, i.e. pulse area, which sets the amount by which the vector is turned. Any point on the Bloch sphere can be reached, i.e. any state of a single qubit can be produced, with a combination of rotations around two orthogonal axes.

In our setup, it is most convenient to induce rotations around two orthogonal axes lying in the equatorial plane, which we will refer to as σ_x and σ_y gates. To choose between them, we shift the driving radiation's phase by 90°. Our ion trap only permits simultaneous rotations on both ions, as it was designed primarily for two-qubit gate experiments and hence, does not include the necessary means for driving quantum logic operations on individual qubits.

We carry out single qubit manipulations with typical pulse lengths between 1 and 2 µs. In previous work, using the same ion trap as in this thesis, the single qubit gate error was measured to be $1.0(3) \times 10^{-6}$ [HAB+14] using the technique of randomised benchmarking, which is briefly explained in section 4.2.4. This result is two orders of magnitude below the fault tolerant threshold and constitutes the current state of the art.

3.1.2 Two-qubit gates

A two-qubit quantum logic gate is an operation that causes two physically separated qubits to undergo a coherent interaction that is conditioned on their state. Its key purpose is to enable the generation of entanglement between the qubits involved, complementing the single-qubit operations in establishing a universal set of gates.

When looking for a mechanism that would be adequate for implementing two-qubit gates in a trapped ion system, we can consider the direct coupling between the magnetic dipoles of each ion. However, since neighbouring ions are typically separated by 1 μm and 10 μm (8 μm in our trap), this interaction is in the mHz range [KAN+14], and therefore too weak for practical use.

An alternative method, which was originally proposed by Cirac and Zoller [CZ95] and nowadays forms the backbone of most gate schemes used in practice, takes advantage of the much stronger Coulomb interaction. It is this interaction that induces the ions to form a stationary crystal with some residual oscillatory motion that can be broken down into collective normal modes. As already mentioned in section 1.3, each of these modes behaves like a quantum harmonic oscillator when sufficiently cooled, creating a system that is formally equivalent to the Jaynes-Cummings model [JC63] of quantum optics. Thus, in a similar way as the two-level atom can be coherently coupled to an optical cavity mode by applying suitable radiation pulses, so can the internal levels of each ion be coupled to a vibrational mode of the ion chain. The major insight of Cirac and Zoller, was that this shared motional mode can be employed as a bus through which quantum information can be transferred between ions, thus establishing the required connection between qubits for quantum logic operations.

Most trapped ion two-qubit gate schemes have since been developed on this foundation [TWK+98, SM99, JPK00, MW01, LDM+03, GRZC03, BSPR12]. However, the ones that proved to be most successful experimentally, rely on a

refined method that uses the concept of mediating the interaction between qubits via their collective motion in conjunction with the generation of a spin-dependent geometric phase, as will be explained in the following section. For completeness we also mention [CZ00, vS03] as examples of two-qubit gate proposals that utilise the long range Coulomb interaction between ions, but don't require the coupling to a common motional mode.

Realising two-qubit entangling operations is the most technically demanding hurdle on the way towards building a quantum information processor, mainly because quantum information channelled through the bus-mode is prone to noise from the environment. The currently lowest error in a trapped ion system is held by a laser-driven gate implementation and amounts to $8(4) \times 10^{-4}$ [Gae16].

3.1.3 Two qubit gates using spin-dependent forces

Although simple in theory, the original two-qubit gate proposed by J. I. Cirac and P. Zoller is technically challenging to realise, as it requires the ions to be addressed individually with laser beams, and fails when the bus-mode is not cooled to the ground state. By the time it was experimentally realised in 2003 [SKHR$^+$03], a new class of gates that circumvents these problems, and that is now favoured, had already been developed.

The novelty of this class is that it utilises a spin-dependent force to drive the ion system around closed loops in position/momentum space of the bus mode, i.e. to transiently excite a vibrational mode. When these loops close, the motion has returned to its original state, but the ions have acquired a geometric phase, the Berry-phase [Ber84], which is proportional to the area enclosed by the phase space trajectory. Furthermore, since the force is conditional on the ions' spin, so is the acquired phase factor. When tuned correctly, it can be used to create entanglement.

This is the basic principle of operation behind the "Mølmer-Sørensen gate"

[SM99], the "light-shift gate" [LDM+03] and the "driven geometric phase gate" [BSPR12], amongst others. What differentiates these schemes, is the method of how the spin-dependent force is generated, and in what qubit basis it is acting. What they have in common, is that they do not fail when the ions' motion is in a thermal state, which significantly relaxes experimental requirements compared with the Cirac-Zoller scheme. Moreover, none of these schemes need individual ion addressing, since both ions are subjected to the radiation field simultaneously.

3.2 Microwave driven two-qubit gates

Until recently, the above mentioned two-qubit gate schemes have been driven almost exclusively using lasers. However, laser-based control is technically very demanding due to the stringent requirements on beam quality, pointing stability, as well as frequency, phase and intensity noise. In spite of its success, the associated technical overhead raises the question of feasibility when considering implementations on a larger scale. Although proposals for architectures relying solely on lasers exist, e.g. [LKOW07], it is nevertheless desirable to develop a method that is easier to scale up but has equivalent capabilities.

Microwaves are a promising candidate, as they are a mature, well-established technology that would allow us to fall back on a vast repertoire of commercial off-the-shelf components with excellent specifications. In contrast to optical setups, microwave networks are much easier to make interferometrically stable and can partially be incorporated into the ion trap itself. Furthermore, since the qubit transition can be driven directly with microwaves, without having to employ dipole allowed transitions like in a Raman process, errors due to spontaneous scattering of photons [OIB+07] are eliminated.

3.2.1 Field gradients and Lamb-Dicke parameter

A prerequisite for implementing any two-qubit gate scheme is coherent control over a selected vibrational mode of the ion chain. As was already addressed in section 2.4.3 when discussing resolved-sideband cooling, we can move up and down the motional energy ladder when driving sideband transitions. In the context of two-qubit gates, the aim is to coherently excite and deexcite the bus-mode by means of sideband Rabi oscillations, induced with the help of appropriate radiation pulses. The corresponding Rabi frequency is the figure of merit for the strength of the coupling between the radiation field and the bus-mode.

In physical terms, this coupling takes the form of a momentum transfer from the radiation field onto the ion chain's motion, conveyed through a force on the atoms' magnetic dipoles. The strength of the force is given by the gradient of the field over the extent of the motional mode's wavefunction.

For plane-wave radiation, we can quantify the coupling strength by means of the Lamb-Dicke parameter

$$\eta \equiv 2\pi q^0/\lambda, \tag{3.1}$$

which relates the extent of the motional mode's ground state wavefunction q^0 to the wavelength of the external field λ.

Working with the right size of the Lamb-Dicke parameter is crucial for a successful gate implementation. If too small, the applied field gradient is weak over the distance q^0 and consequently the motional sideband Rabi frequencies are low, leading to slow gates. If too large, nonlinearities of the field gradient become apparent in the region to which the ions' motion is confined. Since the size of this region varies depending on the state the bus mode is in, so does the force the ions experience, causing errors during the gate operation. However, this error only comes into play if $\eta \ll 1$ is not satisfied, i.e. outside of the Lamb-Dicke regime (section 2.4.3).

In order to estimate a suitable value of η for optimal gate operation, we assume a typical trap secular frequency of a few MHz, so that q^0 is of order 10 nm. The wavelength of the plane-wave radiation field should be as small as possible, but at least an order of magnitude larger than q^0 to provide a linear gradient over a wide enough area to ensure an operation within the Lamb-Dicke regime. Thus, λ should ideally be in the optical domain, where $\eta \approx 0.1$, which is why traditionally two-qubit gates have been driven using lasers.

Attempting to use plane-wave microwave radiation with wavelengths in the cm range, results in extremely low Lamb-Dicke parameters of order 10^{-7}, giving rise to unusably small sideband Rabi frequencies. To date, two ways to overcome this problem have established themselves, both of which are discussed in the remainder of this chapter and the first of them being the method of choice for the experiments described in this thesis.

3.2.2 Beating small field gradients - the near-field approach

One way to increase field gradients in the microwave domain, is to work in the near-field of microwave current-carrying structures. In this region the spatial variations of the field are set by the dimensions of the conductors, which can be made significantly smaller than the free-space wavelength.

Although already mentioned in 1998 [WMW+98], this idea was not seriously pursued for another 10 years, until microfabrication technologies allowed for the manufacturing of planar ion traps that opened the possibility of trapping ion crystals in the near-field region above the microwave electrodes. The first proposal for a two-qubit gate based on this concept appeared in 2008 [OLA+08] and the experimental demonstration employing the Mølmer-Sørensen scheme followed three years later [OWC+11, WOC+13a] with an error of $24(3)\%$.

We can model a typical system used for near-field microwave gates as a linear string of ions, oriented parallel to a straight current carrying wire with

a separation d of a few tens of μm. To facilitate the comparison of coupling strengths between this system and others, we define an effective Lamb-Dicke parameter as

$$\tilde{\eta} \equiv \Omega_{\mathrm{SB}}/\Omega, \tag{3.2}$$

which is typically $\approx q^0/d$ and relates how strongly the field can drive sideband transitions relative to carrier transitions, reflecting the essence of the Lamb-Dicke parameter in the plane-wave case. The carrier and sideband Rabi frequencies, Ω and Ω_{SB}, will be made explicit later, in section 4.1.4, justifying the approximation of $\tilde{\eta}$ to q^0/d.

Effective Lamb-Dicke parameters for near-field microwave radiation are usually of order 10^{-4}, three orders of magnitude larger than for free-space radiation, but at the same time three orders of magnitude smaller than those achieved with lasers. Sideband Rabi frequencies are therefore high enough to make two-qubit gates possible, but only using strong driving fields, since otherwise the gate duration would be unfeasibly long. As will be discussed in the following chapter, strong driving fields imply off-resonant excitation errors and large AC-Zeeman shifts.

Typical gate speeds are in the ms range, allowing for various noise processes that set in on these time scales to cause errors. A significant such noise process occurs when stray electric fields from the trap's surface couple to the ions' charges, inducing randomly fluctuating forces that lead to uncontrolled increase of the motional mode's thermal occupation $\bar{n}$. The effect is dubbed "anomalous heating" and its fundamental cause is, to date, not fully understood. According to past measurements by numerous groups in the ion-trap community, recently compared in [BKRB15], the perturbative effect of the electric field noise approximately follows the empirical power law of d^{-4}. This is particularly problematic for near-field driven gates, which rely on small values of d to produce sufficiently strong gradients.

On the upside, low Lamb-Dicke parameters are beneficial as they guarantee a gate operation within the Lamb-Dicke regime. In fact, we are more likely to run into problems with anharmonicities of the trapping potential at higher thermal occupations of the motional mode, than to see errors from nonlinearities of the field gradient.

Finally, a complication of the near-field scheme relates to the fact that the field produced by a microwave current-carrying conductor is not only localised to the qubit that it intends to drive. Instead, it spreads across the whole trap, causing unwanted operations on different qubits in the processor. When deliberately trying to address individual qubits with microwaves, such crosstalk errors can be suppressed with the techniques outlined in [WOC+13b], provided the ions are arranged a linear array. An alternative approach, that was proposed and implemented by our research group in Oxford [ACLH+14, ACLS+16], employs specific electrode layouts that include multiple trapping zones. The idea is to equip each zone with a set of microwave control electrodes, the purpose of which is to produce additional microwave near-fields that actively null residual fields at the centre of the respective zone.

3.2.3 Beating small field gradients - the MAGIC approach

Magnetic gradient induced coupling (MAGIC) was proposed in 2001 [MW01] and represents an alternative way to increase the Lamb-Dicke parameter for free-running microwave radiation. It is premised on an additional static magnetic field gradient applied along the ion chain, creating an effective (complex) Lamb-Dicke parameter

$$\left|\tilde{\eta}\right| = \sqrt{\eta^2 + \kappa^2} \approx \kappa, \tag{3.3}$$

where κ is proportional to the strength of this gradient and can reach sizeable values in the 10^{-3} regime. Furthermore, the static gradient leads to a position-dependent Zeeman shift, giving every ion a unique resonance frequency and

thereby allowing for individual qubit addressing.

Initial measurements of the spin-motion coupling induced by this method were carried out in 2009 [JBT+09], followed three years later by a two-qubit gate demonstration [KPS+12] with an error of $36(5)\%$ for neighbouring ions and $43(4)\%$ for next-nearest neighbouring ions. The main limitation was magnetic field noise as the MAGIC scheme cannot be implemented on magnetic field-insensitive qubits. In order to work around this problem, current experimental efforts make use of dressed states [TBJ+11] that span a decoherence-free subspace. Recently, a two-qubit Mølmer-Sørensen gate making use of this approach achieved an error of $1.5(1.2)\%$ [WRW+16].

The chief benefits of using the MAGIC scheme are the possibility to straightforwardly spectrally address individual qubits [PSVW14], and the ability to drive quantum logic gates using global radiation fields. Furthermore, the ion to trap surface separation is not very stringent and can be made much larger than in the near-field approach, thus reducing the gate sensitivity to anomalous heating.

On the flipside, having distinguishable qubit resonances also means that each evolves differently in time. For higher numbers of ions, accurate bookkeeping of frequencies and phases may therefore lead to high experimental complexity. Moreover, the MAGIC scheme is more vulnerable to magnetic field noise than the near-field approach, as the dressed states so far demonstrated fall significantly short of providing the same degree of robustness that is given by naturally-occurring clock states. This manifests itself in coherence times of order $1\,\mathrm{s}$ [TBJ+11] compared with e.g. $50\,\mathrm{s}$ for clock qubits [HAB+14]. Additionally, employing dressed states goes together with further experimental overhead, as multiple additional microwave fields are required for their generation.

4

Two-qubit gate theory

In this chapter we build the theoretical foundation underlying two-qubit gates when implemented using near-field oscillating magnetic fields. We focus specifically on the Mølmer-Sørensen scheme, as it is the scheme of choice for the experiments conducted in this work. Apart from discussing the origin of various sources of infidelity and their impact on the gate, we also consider alterations to the standard Mølmer-Sørensen recipe, aimed at increasing the gate's robustness against some of these errors.

4.1 Microwave mediated spin-motion coupling

The ability to couple the ions' internal states with their collective motional states forms the basis of a two-qubit gate. In the following section, we show how this coupling can be established with the help of an external, oscillating magnetic field, by deriving the associated interaction Hamiltonian. We keep the derivation general, and do not yet restrict ourselves to specific qubit states. The results obtained in this chapter are therefore applicable to any pair of states separated by a magnetic dipole transition.

For convenience, we simplify the notation introduced in chapter 2. Instead of referring to the qubit states in the upper and lower ground state manifold as $|M_n\rangle$ and $|M_m\rangle$, we henceforth use $|1\rangle$ and $|0\rangle$.

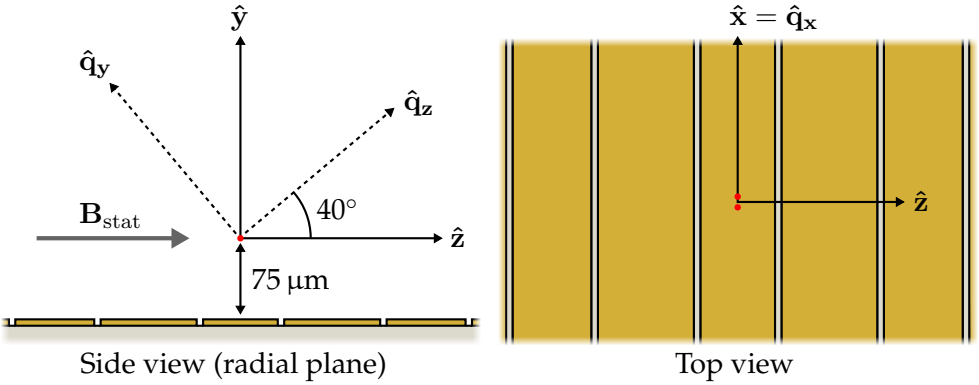

Figure 4.1: Sketch of the orientation of the $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ coordinate system as well as the principal pseudo-potential axes $\{\hat{\mathbf{q}}_\mathbf{x}, \hat{\mathbf{q}}_\mathbf{y}, \hat{\mathbf{q}}_\mathbf{z}\}$ with respect to the trap surface. The direction of the external static magnetic field $\mathbf{B}_{\text{stat}}$ is along $\hat{\mathbf{z}}$, while the ions, indicated by red spheres, line up along $\hat{\mathbf{x}}$, 75 µm above the trap surface.

4.1.1 Quantised modes of motion

We consider a linear, surface electrode Paul trap with weaker axial confinement than radial. Let the electrode geometry be related to an orthogonal coordinate system $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ so that the trap's surface is parallel to the xz-plane, $\hat{\mathbf{y}}$ points away from the surface in the upward direction, and $\hat{\mathbf{x}}$ lies within the trap's symmetry plane. The generated trapping potential can be modelled as a harmonic pseudo-potential with principal axes $\{\hat{\mathbf{q}}_\mathbf{x}, \hat{\mathbf{q}}_\mathbf{y}, \hat{\mathbf{q}}_\mathbf{z}\}$, where the weak axis $\hat{\mathbf{q}}_\mathbf{x}$ matches $\hat{\mathbf{x}}$ and the two radial directions $\hat{\mathbf{q}}_\mathbf{y}$ and $\hat{\mathbf{q}}_\mathbf{z}$ are tilted from $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$ by $\sim 40°$ to facilitate Doppler cooling, as explained in section 5.1. The orientation of both coordinate systems with respect to the trap's surface is sketched in Fig. 4.1.

When confined and sufficiently cold, N identical ions of mass m will arrange themselves collinearly along the weak axis, forming a stationary chain with some residual oscillatory motion. This collective motion can be modelled as a set of independent harmonic oscillators, provided the anharmonicities of the trapping potential and of the ions' mutual Coulomb interaction are sufficiently small [Jam98]. In the vast majority of ion trapping experiments, including the one presented in this thesis, this picture holds true. Consequently, the oscillations of the ion chain can be broken down into N normal modes with characteristic

frequencies, along each of the three principal trap axes. The position $\mathbf{r_n}$ of the n^{th} ion in the vibrating string is

$$\mathbf{r_n} \approx \mathbf{r_n^{(0)}} + \sum_{l=\{x,y,z\}} \sum_{j=1}^{N} \hat{\mathbf{q}}_l \, b_n^{(l,j)} q_{l,j}^0 (a_{l,j} + a_{l,j}^\dagger), \tag{4.1}$$

where $\mathbf{r_n^{(0)}}$ is the ion's equilibrium position, and where the second term in the sum represents a displacement, already decomposed into normal modes, indexed with j. The operators $a_{l,j}$ and $a_{l,j}^\dagger$ are the normal modes' ladder operators, $\omega_{l,j}$ is their frequency, $q_{l,j}^0 = \sqrt{\hbar/2m\omega_{l,j}}$ is the spatial extent of their respective ground state wave function, and $b_n^{(l,j)}$ is the normal modes' eigenvector, determining the phase and relative amplitude of the n^{th} ion's oscillation.

For a single ion, the normal mode decomposition reduces to describing its motion in terms of simple translations along the three principal trap axes, known as centre of mass modes (COM). When multiple ions are present in the chain, the motion can be more complex and the resulting normal mode structure much richer. In the case of two-ion crystals for example, apart from centre of mass modes with eigenvectors $b^{(l,1)} = \frac{1}{\sqrt{2}}(1,1)$, there exists one breathing mode in the axial direction and two rocking modes in the radial directions, all three bearing the eigenvectors $b^{(l,2)} = \frac{1}{\sqrt{2}}(-1,1)$.

4.1.2 The system Hamiltonian

We wish to find the Hamiltonian describing the collective system of N ions confined in a harmonic trap and subject to an external oscillating magnetic field. The three building blocks are, the internal energy of the ions H_0, the motion of the ion chain H_{motion}, and the interaction with the external field $H_{\text{I}}(t)$

$$H(t) = H_0 + H_{\text{motion}} + H_{\text{I}}(t). \tag{4.2}$$

Since each qubit can be treated as an individual two-level system, the internal

energy contribution of the chain is given by

$$H_0 = \sum_{n=1}^{N} H_{0,n} = \sum_{n=1}^{N} \frac{1}{2}\hbar\omega_0\sigma_{z,n}, \tag{4.3}$$

where ω_0 is the qubit transition frequency. The motional part, as explained in the previous section, is composed of j normal modes, or equivalently, of j independent quantum harmonic oscillators along each of the three principal trap axes $\hat{\mathbf{q}}_l$. The associated Hamiltonian reads

$$H_{\text{motion}} = \sum_{l=\{x,y,z\}} \sum_{j=1}^{N} \hbar\omega_{l,j}\Big(\frac{1}{2} + a_{l,j}^{\dagger}a_{l,j}\Big). \tag{4.4}$$

For the interaction part, we introduce the drive field $\mathbf{B}(\mathbf{r}, t)$, which oscillates at a frequency ω with a phase ϕ. Its coupling to the ions' magnetic dipole moments $\boldsymbol{\mu}$ is given by

$$H_{\text{I}}(t) = \sum_{n=1}^{N} H_{I,n}(t) = -\sum_{n=1}^{N} \boldsymbol{\mu}\cdot\mathbf{B}(\mathbf{r_n})\cos(\omega t - \phi). \tag{4.5}$$

4.1.3 The microwave magnetic field

Let us now take a closer look at the drive field, as seen by the n^{th} ion. Provided that the ion's displacement from its equilibrium position $\mathbf{r_n^{(0)}}$ is sufficiently small, we can linearise the field as follows

$$\mathbf{B}(\mathbf{r_n}) \approx \Big\{\mathbf{B}\big(\mathbf{r_n^{(0)}}\big) + \frac{d\mathbf{B}}{d\mathbf{q}}\big(\mathbf{r_n^{(0)}}\big)\,\Big(\mathbf{r_n} - \mathbf{r_n^{(0)}}\Big)\Big\}, \tag{4.6}$$

where $d\mathbf{B}(\mathbf{r_n^{(0)}})/d\mathbf{q}$ is the Jacobian evaluated at the n^{th} ion's equilibrium position. It takes the form

$$\frac{d\mathbf{B}}{d\mathbf{q}} = \begin{bmatrix} \partial B_{+1}/\partial q_x & \partial B_{+1}/\partial q_y & \partial B_{+1}/\partial q_z \\ \partial B_{-1}/\partial q_x & \partial B_{-1}/\partial q_y & \partial B_{-1}/\partial q_z \\ \partial B_z/\partial q_x & \partial B_z/\partial q_y & \partial B_z/\partial q_z \end{bmatrix}, \tag{4.7}$$

when working in the spherical basis defined by the unit vectors $\{\hat{\mathbf{e}}_0 = \hat{\mathbf{e}}_z, \hat{\mathbf{e}}_{-1}, \hat{\mathbf{e}}_{+1}\}$. This is a convenient choice that will later allow us to make direct use of the previously evaluated transition matrix elements from section 2.2.

In typical setups used for near-field quantum logic, the drive field is generated by a current flowing through a straight section of trap electrode, which lies parallel to the weak trap axis. Thus, $\mathbf{B}(\mathbf{r},t)$ is translationally symmetric along $\hat{\mathbf{q}}_\mathbf{x} = \hat{\mathbf{x}}$ and all terms proportional to $\partial B_Q/\partial q_1$ can be discarded. Furthermore, we can drop the explicit dependence on $\mathbf{r}_\mathbf{n}^{(\mathbf{0})}$ and instead introduce the abbreviations $B_Q(\mathbf{r}_\mathbf{n}^{(\mathbf{0})}) \equiv B_Q^{(0)}$ and $\partial B_Q(\mathbf{r}_\mathbf{n}^{(\mathbf{0})})/\partial q_l \equiv \partial B_Q^{(0)}/\partial q_l$, where the superscript (0) denotes the given quantity's value on the x-axis.

Using these simplifications together with Eq. 4.1 result in a linearised magnetic field, as seen by the n^{th} ion

$$\begin{aligned}\mathbf{B}(\mathbf{r_n}) &= \sum_{Q=\pm1,0} (-1)^Q B_Q \hat{\mathbf{e}}_{-Q} \\ &= \sum_{Q=\pm1,0} (-1)^Q \Big\{ B_Q^{(0)} + \sum_{l=\{y,z\}} \sum_{j=1}^{N} \frac{\partial B_Q^{(0)}}{\partial q_l}\, b_n^{(l,j)} q_{l,j}^0 \big(a_{l,j} + a_{l,j}^\dagger\big) \Big\} \hat{\mathbf{e}}_{-Q},\end{aligned} \tag{4.8}$$

which confirms the claim made earlier in section 3.2.1, that the magnetic field gradients are responsible for exciting the ion chain's motion.

4.1.4 Coupling spin and motion with microwave field gradients

Within the qubit subspace $\{|0\rangle, |1\rangle\}$, we rewrite the interaction Hamiltonian for the n^{th} ion, from Eq. 4.5, in terms of its matrix elements

$$H_{\text{I},n}(t) = \begin{bmatrix} H_{00}(t) & H_{01}(t) \\ H_{10}(t) & H_{11}(t) \end{bmatrix} = H_{01}(t)\sigma_{x,n} + \frac{1}{2}\big(H_{11}(t) - H_{00}(t)\big)\sigma_{z,n}, \tag{4.9}$$

where $\sigma_{x,n}$ and $\sigma_{z,n}$ are the Pauli operators acting on the internal states of the n^{th} ion, and where $H_{ij} = \langle i|H_{\text{I},n}(t)|j\rangle$. The constant additive energy offset $\frac{1}{2}(H_{11}(t) + H_{00}(t))\mathbb{1}$ was dropped and $H_{01}(t)$ was set equal to $H_{10}(t)$ because $H_{\text{I},n}(t)$ is Hermitian and has real matrix elements.

The $\sigma_{x,n}$ term in Eq. 4.9 constitutes a coupling between the $|1\rangle$ and $|0\rangle$ states, induced by the component of the magnetic field that matches the polarization of the qubit transition. It is given by

$$H_{10}(t) = (-1)^{Q+1} B_{-Q} \langle 1|\mu_Q|0\rangle \cos(\omega t - \phi), \tag{4.10}$$

where analytic expressions for the transition matrix elements $\langle 1|\mu_Q|0\rangle$ can be found in section 2.2, and their numerical values in appendix A.

The $\sigma_{z,n}$ term represents a change in qubit frequency and reads

$$\frac{1}{2}\big(H_{11}(t) - H_{00}(t)\big) = -\frac{1}{2}B_z\left\{\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle\right\}\cos(\omega t - \phi). \tag{4.11}$$

As already explained in section 2.2, this term vanishes for clock transitions.

With Eq. 4.10 and Eq. 4.11 inserted into Eq. 4.9, and using the definitions

$$\bar{\mu}_Q \equiv -(-1)^{Q+1}\langle 1|\mu_Q|0\rangle\,, \tag{4.12}$$

$$\bar{\mu}_z \equiv \frac{1}{2}\left\{\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle\right\}, \tag{4.13}$$

the interaction Hamiltonian for the n^{th} ion can be written as

$$H_{\mathrm{I},n}(t) = -\big(\bar{\mu}_Q B_{-Q}\sigma_{x,n} + \bar{\mu}_z B_z \sigma_{z,n}\big)\cos(\omega t - \phi). \tag{4.14}$$

When inserting the linearised form of the drive field from Eq. 4.8, we can further expand the above expression

$$\begin{aligned} H_{\mathrm{I},n}(t) = -\cos(\omega t - \phi)\Bigg\{&\bar{\mu}_Q B^{(0)}_{-Q}\sigma_{x,n} + \bar{\mu}_z B^{(0)}_z \sigma_{z,n} \\ &+ \sum_{l=\{y,z\}}\sum_{j=1}^{N}\left(\bar{\mu}_Q\frac{\partial B^{(0)}_{-Q}}{\partial q_l}\sigma_{x,n} + \bar{\mu}_z\frac{\partial B^{(0)}_z}{\partial q_l}\sigma_{z,n}\right) b_n^{(l,j)} q^0_{l,j}\Big(a_{l,j} + a^\dagger_{l,j}\Big)\Bigg\}. \end{aligned} \tag{4.15}$$

In order to separate out the effect that the magnetic field has on the internal and external states of the ions, we move into the interaction picture with respect to $H_0 + H_{\text{motion}}$ and make a rotating wave approximation, neglecting all terms rotating at $2\omega_0$. The Hamiltonian describing the coupling between the n^{th} ion and an oscillating magnetic field in the interaction picture (identified with the tilde accent) becomes

$$\begin{aligned} \tilde{H}_{\mathrm{I},n}(t) = &\frac{1}{2}\hbar\Omega_Q e^{-i(\delta t-\phi)}\sigma_{+,n} + \frac{1}{2}\hbar\Omega_z e^{-i(\omega t-\phi)}\sigma_{z,n} \\ &+ \sum_{l=\{y,z\}}\sum_{j=1}^{N}\frac{1}{2}\hbar\Omega_{Q,l,j,n}e^{-i(\delta t-\phi)}\sigma_{+,n}\Big(a_{l,j}e^{-it\omega_{l,j}} + a^\dagger_{l,j}e^{it\omega_{l,j}}\Big) \\ &+ \sum_{l=\{y,z\}}\sum_{j=1}^{N}\frac{1}{2}\hbar\Omega_{z,l,j,n}e^{-i(\omega t-\phi)}\sigma_{z,n}\Big(a_{l,j}e^{-it\omega_{l,j}} + a^\dagger_{l,j}e^{it\omega_{l,j}}\Big) \\ &+ \text{h.c.}, \end{aligned} \tag{4.16}$$

where we have defined the drive field detuning from the qubit transition as $\delta \equiv \omega - \omega_0$, and where we introduced the Rabi frequencies

$$\Omega_Q \equiv -\frac{1}{\hbar}\bar{\mu}_Q B^{(0)}_{-Q}, \tag{4.17}$$

$$\Omega_z \equiv -\frac{1}{\hbar}\bar{\mu}_z B^{(0)}_z, \tag{4.18}$$

$$\Omega_{Q,l,j,n} \equiv -\frac{1}{\hbar}\bar{\mu}_Q \frac{\partial B^{(0)}_{-Q}}{\partial q_l} b^{(l,j)}_n q^0_{l,j}, \tag{4.19}$$

$$\Omega_{z,l,j,n} \equiv -\frac{1}{\hbar}\bar{\mu}_z \frac{\partial B^{(0)}_z}{\partial q_l} b^{(l,j)}_n q^0_{l,j}. \tag{4.20}$$

This Hamiltonian contains four well separated resonances that can be targeted by appropriately tuning the microwave frequency ω.

The first term in Eq. 4.16, together with its Hermitian conjugate, has a resonance at $\delta = 0$ and describes carrier excitation, that is, Rabi oscillations with frequency Ω_Q between $|1\rangle \otimes |n_{l,j}\rangle$ and $|0\rangle \otimes |n_{l,j}\rangle$. As previously explained in section 3.1.1, this resonance is responsible for driving single qubit rotations around an axis on the Bloch sphere that is set by the radiation's phase ϕ. However, we will see in section 4.2.1, that implementing two-qubit gates requires radiation that is detuned from the qubit transition by the motional mode frequency $\omega_{l,j}$. In this context, the first term in Eq. 4.16 describes off-resonant excitation of the qubit transition, which causes an unwanted population transfer between the qubit states, and shifts their frequency. The former effect will be treated in section 4.3.4, while the latter is called an AC-Zeeman shift and will be explained in more detail in section 4.3.5.

The second term in Eq. 4.16 represents a periodic change of the qubit's energy, following the oscillation of the $B^{(0)}_z$ component of the magnetic field with frequency ω. Since in our experiments we work on a clock transition, for which $\bar{\mu}_z = 0$ (see section 2.2) and therefore $\Omega_{z,l,j,n} = 0$, this term is negligible.

The third term and its Hermitian conjugate, when tuned to either of their two resonances at $\delta = \pm\omega_{l,j}$, are equivalent to the James-Cummings Hamiltonian of quantum optics, and will be written out explicitly in section 4.2.1. For $\delta = -\omega_{l,j}$,

we obtain Rabi flopping between $|1\rangle \otimes |n_{l,j}\rangle$ and $|0\rangle \otimes |n_{l,j}+1\rangle$ with frequency $\sqrt{n+1}\,\Omega_{Q,l,j,n}$, which is a red sideband transition that swaps the internal state of the ions while removing one quantum of motion from the j^{th} normal mode along the principal trap axis $\hat{\mathbf{q}}_l$. Its counterpart is the the blue sideband transition at $\delta = +\omega_{l,j}$, describing Rabi oscillations between $|1\rangle \otimes |n_{l,j}\rangle$ and $|0\rangle \otimes |n_{l,j}-1\rangle$, also with frequency $\sqrt{n}\,\Omega_{Q,l,j,n}$.

The remaining resonance at $\omega = \omega_{l,j}$ is generated by the fourth term in Eq. 4.16 and its Hermitian conjugate. It too is capable of exciting the ion chain's vibrational mode, but without a simultaneous change in the qubit's internal state.

In principle, both pairs of sidebands give rise to a similar interaction and therefore either can be used to generate entanglement, and hence, to implement a two-qubit gate. When working with a $B_z(t)$ field tuned close to $\omega_{l,j}$, we obtain a $\sigma_z \otimes \sigma_z$ or geometric phase gate. When employing a $B_{-Q}(t)$ field, we have to use two frequency tones close to $\omega_0 \pm \omega_{l,j}$ to get a $\sigma_\phi \otimes \sigma_\phi$ or Mølmer-Sørensen gate [OLA+08]. Apart from the different drive field characteristics, the only mathematical difference between these schemes lies in the basis in which they operate.

In this work we use the Mølmer-Sørensen scheme, as $\sigma_z \otimes \sigma_z$ gates have two practical disadvantages when driven by near-field microwaves. First, they cannot be implemented on clock qubits, as otherwise $\bar{\mu}_z = 0$ and thus $\Omega_{z,l,j,n} = 0$. This is problematic, since the weak field gradients produced by near-field radiation only permit slow gates, making the long coherence times offered by clock transitions crucial. Second, because the driving field currents have to be near resonant with the motional mode frequency, it is possible for accompanying oscillating electric potentials to couple directly to the ions' motion, causing heating, especially on COM modes [OLA+08].

As a final remark, now that we found expressions for the Rabi frequencies, we are in the position to back up our earlier claim, made in section 3.2.2, about

the form of effective Lamb-Dicke parameter for near-field radiation. We defined $\tilde{\eta}$ in Eq. 3.2 as the ratio of sideband to carrier Rabi frequency

$$\tilde{\eta} \equiv \frac{\Omega_{Q,l,j,n}}{\Omega_Q} = \left(B^{(0)}_{-Q}\right)^{-1} \frac{\partial B^{(0)}_{-Q}}{\partial q_l} b_n^{(l,j)} q^0_{l,j} \approx q^0/d, \tag{4.21}$$

where we have approximated the magnetic field and its gradient by the standard formulas for a magnetic field generated by a straight current carrying wire. Thus, $B^{(0)}_{-Q} \approx -\mu_0 I/(2\pi d)$ and $\partial B^{(0)}_{-Q}/\partial q_l \approx \mu_0 I/(2\pi d^2)$, with d being the separation of the linear string of ions from the conductor.

4.2 The basic Mølmer-Sørensen gate

Building on the results from the previous section, we now derive the Mølmer-Sørensen Hamiltonian in the absence of errors, integrate it analytically, and study its effects on a pair of qubits.

4.2.1 The gate Hamiltonian

The Mølmer-Sørensen gate relies on a spin-dependent force, which is generated by simultaneously applying a pair of radiation fields tuned close to the first red ($\omega_0 - \omega_{l,j}$) and blue ($= \omega_0 + \omega_{l,j}$) sideband transitions. In order to construct the associated Hamiltonian, we make the following simplifications.

First, we ignore all off-resonant terms in Eq. 4.16, leaving only the contribution proportional to $\Omega_{Q,l,j,n}$. Second, we assume that the interaction is restricted to a single mode of motion j along the principal trap axis $\hat{\mathbf{q}}_l$, and with frequency $\omega_{l,j}$. Third, the Rabi frequencies of the red and blue sideband drives, Ω_{R} and Ω_{B}, and their detunings, δ_{R} and δ_{B}, are set to be equal and constant in time. The resulting expressions for the bichromatic field components are

$$\tilde{H}_{\mathrm{B},n}(t) = \frac{1}{2}\hbar\Omega_{\mathrm{B}}\sigma_{+,n}\left(a_{l,j}e^{-i(\delta_{\mathrm{B}}t+2\omega_{l,j}t-\phi_{\mathrm{B}})} + a^\dagger_{l,j}e^{-i(\delta_{\mathrm{B}}t-\phi_{\mathrm{B}})}\right) + \mathrm{h.c.}, \tag{4.22}$$

$$\tilde{H}_{\mathrm{R},n}(t) = \frac{1}{2}\hbar\Omega_{\mathrm{R}}\sigma_{+,n}\left(a_{l,j}e^{i(\delta_{\mathrm{R}}t+\phi_{\mathrm{R}})} + a^\dagger_{l,j}e^{i(\delta_{\mathrm{R}}t+2\omega_{l,j}t+\phi_{\mathrm{R}})}\right) + \mathrm{h.c.}, \tag{4.23}$$

where $\Omega_{\mathrm{B}} = \Omega_{\mathrm{R}} = \Omega_{Q,l,j,n}$ and $\delta_{\mathrm{B}} = \delta_{\mathrm{R}} = \delta$, in accordance with the simplifications stated above. We note that the latter relation is a redefinition of δ.

The detuning of the blue sideband drive is $\delta_{\mathrm{B}} \equiv \omega_{\mathrm{B}} - \omega_{l,j} - (\omega_0 + \omega_{\mathrm{Z}})$ and that of the red sideband drive is $\delta_{\mathrm{R}} \equiv (\omega_0 + \omega_{\mathrm{Z}}) - \omega_{\mathrm{R}} - \omega_{l,j}$. The qubit frequency shift ω_{Z} is due to the AC-Zeeman effect, which will be explained in more detail in section 4.3.5. Thus, the bichromatic field components are symmetrically offset from their respective resonances in opposite directions, as shown in Fig. 4.2.

After discarding the fast rotating terms at $2\omega_{l,j}$, the Mølmer-Sørensen Hamiltonian for the n^{th} ion becomes

$$\begin{aligned}\tilde{H}_{\mathrm{MS},n}(t) &= \tilde{H}_{\mathrm{B},n}(t) + \tilde{H}_{\mathrm{R},n}(t) \\ &= \frac{1}{2}\hbar\Omega_{Q,l,j,n}\Big(\sigma_{+,n}e^{i\bar{\phi}} + \sigma_{-,n}e^{-i\bar{\phi}}\Big)\Big(a_{l,j}e^{i\delta t}e^{-i\frac{1}{2}\Delta\phi} + a_{l,j}^{\dagger}e^{-i\delta t}e^{i\frac{1}{2}\Delta\phi}\Big) \\ &= \frac{1}{2}\hbar\Omega_{Q,l,j,n}\sigma_{\bar{\phi},n}\Big(a_{l,j}e^{i\delta t}e^{-i\frac{1}{2}\Delta\phi} + a_{l,j}^{\dagger}e^{-i\delta t}e^{i\frac{1}{2}\Delta\phi}\Big),\end{aligned} \tag{4.24}$$

where we have defined

$$\Delta\phi \equiv \phi_{\mathrm{B}} - \phi_{\mathrm{R}}, \tag{4.25}$$

$$\bar{\phi} \equiv \frac{1}{2}(\phi_{\mathrm{B}} + \phi_{\mathrm{R}}), \tag{4.26}$$

$$\sigma_{\bar{\phi},n} \equiv \sigma_{x,n}\cos(\bar{\phi}) - \sigma_{y,n}\sin(\bar{\phi}), \tag{4.27}$$

so that $\Delta\phi$ is the difference phase and $\bar{\phi}$ is the mean phase of the bichromatic field. The latter is also referred to as the "spin phase".

For up to $N = 2$ ions, which is the maximum number of ions used in our experiment, the basic Mølmer-Sørensen gate Hamiltonian can be written as

$$\tilde{H}_{\mathrm{MS}}(t) = \sum_{n=1}^{N}\tilde{H}_{\mathrm{MS},n}(t) = \frac{1}{2}\hbar\Omega^{(N)}S_{\bar{\phi}}\Big(a_{l,j}e^{i\delta t}e^{-i\frac{1}{2}\Delta\phi} + a_{l,j}^{\dagger}e^{-i\delta t}e^{i\frac{1}{2}\Delta\phi}\Big). \tag{4.28}$$

The above expression makes use of the following definitions

$$\Omega^{(N)} \equiv -\frac{1}{\hbar}\bar{\mu}_Q\frac{\partial B_{-Q}^{(0)}}{\partial q_l}\big|b\big|q_{l,j}^0, \tag{4.29}$$

$$S_{\bar{\phi}} \equiv \sum_{n=1}^{N}\mathrm{sgn}(b_n^{(l,j)})\sigma_{\bar{\phi},n}, \tag{4.30}$$

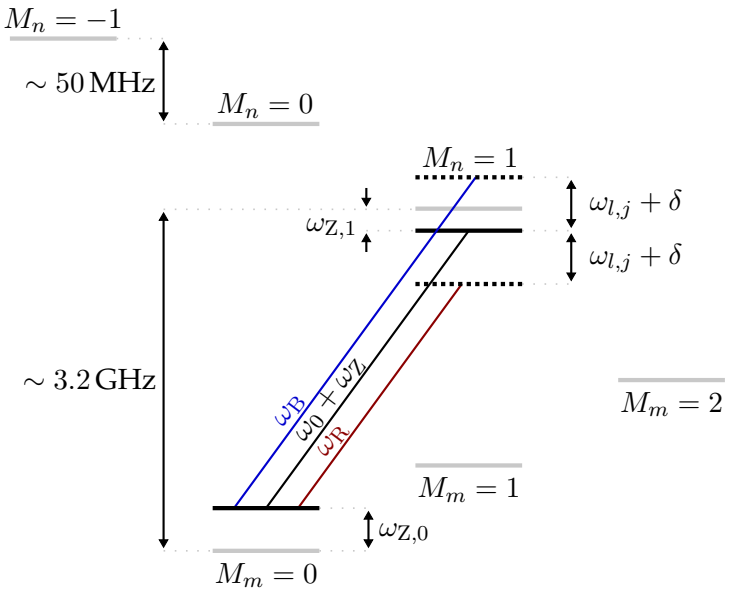

Figure 4.2: Energy level diagram (not to scale) of part of the $4S_{1/2}$ ground level in $^{43}Ca^{+}$ at 146 G, showing the qubit states $|0\rangle$ ($M_m = 0$) and $|1\rangle$ ($M_n = 1$) as well as the four other states connected to them via spectator transitions. The bichromatic microwave field off-resonantly drives the qubit and spectator transitions, causing the lower and upper qubit state energies to (AC-Zeeman) shift by $\omega_{\mathrm{Z},0}$ and $\omega_{\mathrm{Z},1}$, respectively. When taking the combined AC-Zeeman shift $\omega_{\mathrm{Z}} = \omega_{\mathrm{Z},0} + \omega_{\mathrm{Z},1}$ into account, the frequency of the blue sideband microwaves is $\omega_{\mathrm{B}} = (\omega_0 + \omega_{\mathrm{Z}}) + (\omega_{l,j} + \delta)$ and of the red sideband microwaves is $\omega_{\mathrm{R}} = (\omega_0 + \omega_{\mathrm{Z}}) - (\omega_{l,j} + \delta)$, where ω_{Z} can be negative.

where $S_{\tilde{\phi}}$ is the collective N-ion spin operator, and where $\Omega^{(N)}$ is the redefined N-ion Rabi frequency. We note that $\Omega^{(N)}$ is now independent of n, and that we have dropped the indices Q, l, j for convenience. In defining the above quantities, we have used the relation $b_n^{(l,j)} = \mathrm{sgn}(b_n^{(l,j)})\big|b\big|$, with $\big|b\big| = 1$ for a single ion, and with $\big|b\big| = 1/\sqrt{2}$ for two ions.

4.2.2 Analytic solution

Under the influence of the bichromatic field applied for a time Δt, the two-ion system's state vector evolves according to the propagator $\tilde{U}_{\mathrm{MS}}(t_0 + \Delta t, t_0)$, which is defined as

$$\left|\tilde{\psi}(t_0 + \Delta t)\right\rangle \equiv \tilde{U}_{\mathrm{MS}}(t_0 + \Delta t, t_0) \left|\tilde{\psi}(t_0)\right\rangle . \tag{4.31}$$

To find it, we integrate the associated Schrödinger equation in the interaction picture

$$\frac{d}{dt}\tilde{U}_{\mathrm{MS}}(t_0+\Delta t,t_0) = -\frac{i}{\hbar}\tilde{H}_{\mathrm{MS}}(t)\tilde{U}_{\mathrm{MS}}(t_0+\Delta t,t_0), \tag{4.32}$$

which can be done analytically using the Magnus expansion, as explained in appendix B. The resulting Magnus series terminates after the first two terms, and the only non-vanishing commutator reads

$$[H_{\mathrm{MS}}(t),H_{\mathrm{MS}}(t')] = \frac{1}{2}i\Big(\hbar\Omega^{(N)}S_{\bar{\phi}}\Big)^2\sin\Big(\delta\left(t-t'\right)\Big). \tag{4.33}$$

The exact Mølmer-Sørensen propagator is therefore

$$\tilde{U}_{\mathrm{MS}}(t_0+\Delta t,t_0) = \mathcal{D}\big[S_{\bar{\phi}}\alpha(t_0+\Delta t,t_0)\big]e^{-iS_{\bar{\phi}}^2\Phi(t_0+\Delta t,t_0)}, \tag{4.34}$$

where we defined the motional displacement $\alpha(t_0+\Delta t,t_0)$, the geometric phase $\Phi(t_0+\Delta t,t_0)$, and the displacement operator $\mathcal{D}[S_{\bar{\phi}}\alpha(t_0+\Delta t,t_0)]$ in the following way

$$\mathcal{D}\big[S_{\bar{\phi}}\alpha(t_0+\Delta t,t_0)\big] \equiv e^{S_{\bar{\phi}}\{\alpha(t_0+\Delta t,t_0)\ a_{l,j}^{\dagger}-\alpha^*(t_0+\Delta t,t_0)\ a_{l,j}\}}, \tag{4.35}$$

$$\alpha(t_0+\Delta t,t_0) \equiv -i\frac{\Omega^{(N)}}{\delta}e^{i\frac{1}{2}\Delta\phi}e^{-i\frac{1}{2}\delta(2t_0+\Delta t)}\sin\left(\frac{1}{2}\delta\Delta t\right), \tag{4.36}$$

$$\Phi(t_0+\Delta t,t_0) \equiv \left(\frac{\Omega^{(N)}}{2\delta}\right)^2\big\{\delta\Delta t-\sin(\delta\Delta t)\big\}. \tag{4.37}$$

The propagator consists of two parts. The first is the displacement $\mathcal{D}$ that is generated by the first order term in the Magnus expansion and that coherently translates the initial motional state of the ions around circular paths in phase space, as shown in Fig. 4.3. The second is the Berry geometric phase, which originates from the second-order term of the Magnus series and is therefore a consequence of the interaction Hamiltonian not commuting with itself at different times. This phase accumulates as the loops are being traversed and its final value is equivalent to the area enclosed by the trajectory.

The crucial point is that the displacement depends on the collective spin operator $S_{\bar{\phi}}$, or more precisely, on its eigenvalue λ. Consequently, spin states

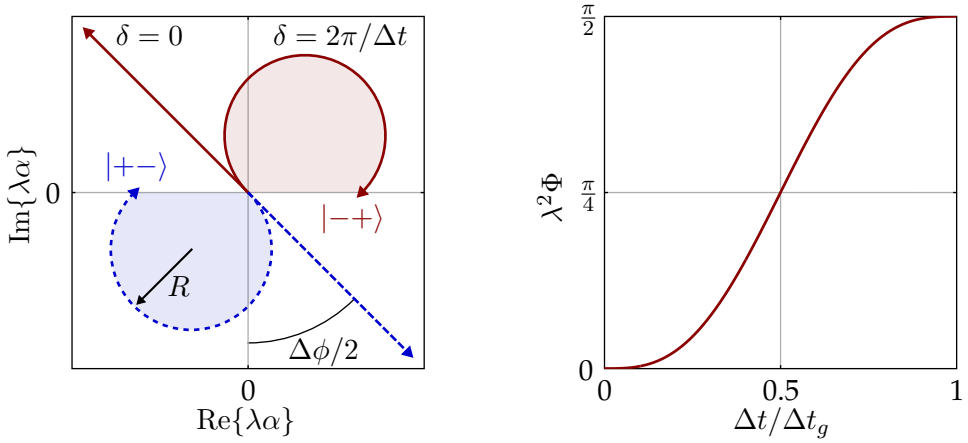

Figure 4.3: Left: Action of the spin-dependent force in phase space, spanned by the quadratures of the complex displacement α, in a rotating frame. The eigenvalues λ of the collective spin operator $S_{\bar{\phi}}$ are zero for the spin states $|{+}{+}\rangle$ and $|{-}{-}\rangle$, leaving them undisplaced. The eigenvalues associated with the $|{+}{-}\rangle$ and $|{-}{+}\rangle$ states are nonzero, and have inverse signs. Hence, these states are displaced in opposite directions, at an angle given by $\Delta\phi/2$. If the force is resonant, i.e. $\delta = 0$, this displacement is linear. If $\delta \neq 0$, the phase space trajectories are curved with a radius $R = (|\lambda|\Omega^{(2)})/(2\delta)$. The geometric phase picked up in one round trip is proportional to the enclosed area, represented by the shaded region within the circles. Right: Geometric phase $\lambda^2\Phi$ accumulation over the course of the gate time Δt_g for a single loop gate.

with differing eigenvalues λ give rise to differently sized loops in phase space, and thus acquire a different geometric phase. Accumulating the right amount of this spin-dependent phase is key for generating an entangling operation between the internal states of two ions.

Furthermore, while the phase space trajectories are being traversed, the internal states get entangled with the motion, which leads to errors when measuring the spin. We avoid these errors by ensuring that the paths form closed loops and return to the origin at the end of the gate operation, i.e. $\alpha(\Delta t, 0) = 0$, leaving the internal states separable from the motion.

4.2.3 Gate implementation

To demonstrate how the propagator given in Eq. 4.34 can be used to implement a two-qubit gate, we consider a force that only couples to the two-ion rocking mode, for which $S_{\bar{\phi}} = \sigma_{\bar{\phi},1} - \sigma_{\bar{\phi},2}$. We note that this is the mode that we work

with in practice, for reasons that will be explained in section 4.3.6.

The three parameters that have to be set appropriately are the Rabi frequency $\Omega^{(2)}$, the gate detuning δ, and the interaction time Δt. When complying with

$$\Delta t = \frac{2\pi K}{\delta} \equiv \Delta t_g, \tag{4.38}$$

$$\delta = 2\Omega^{(2)}\sqrt{K} \equiv \delta_g, \tag{4.39}$$

it is ensured that the K loops in phase space close, and that the area encompassed by them gives rise to a round-trip geometric phase of $\lambda^2\Phi(\Delta t_g, 0) = \pi/2$. For convenience, we will refer to the correct interaction time as Δt_g, and to the correct gate detuning as δ_g.

As a side note, we will see in section 4.3, that using a higher number of loops in phase space reduces the gate's sensitivity to various sources of error. It is clear from Eq. 4.38 and Eq. 4.39, that in order to increase the number of loops from K_1 to K_2, while keeping the gate time constant, the Rabi frequency has to be increased by a factor $\sqrt{K_2/K_1}$. Consequently, the microwave power of both the red and the blue sideband drive has to be raised by a factor of K_2/K_1. This puts a practical limit on the number of loops we can achieve in our experiment, as the ion trap can only dissipate a finite amount of microwave power. In this work, we use $K = 4$ loops, requiring input powers of approximately 2 W in each sideband.

Under the conditions given by Eq. 4.38 and Eq. 4.39, the gate performs the operation

$$\begin{aligned} |{+}{+}\rangle &\mapsto \, |{+}{+}\rangle \\ |{+}{-}\rangle &\mapsto -i\, |{+}{-}\rangle \\ |{-}{+}\rangle &\mapsto -i\, |{-}{+}\rangle \\ |{-}{-}\rangle &\mapsto \, |{-}{-}\rangle\,, \end{aligned} \tag{4.40}$$

which is equivalent to a two-qubit phase gate. However, this phase gate is not acting in the computational basis $\{|11\rangle, |10\rangle, |01\rangle, |00\rangle\}$, but in the eigenbasis of

$S_{\bar{\phi}}$, spanned by

$$|\pm\rangle_1 \otimes |\pm\rangle_2 = \left[\frac{1}{\sqrt{2}}\big(\,|1\rangle_1 \pm e^{-i\bar{\phi}}\,|0\rangle_1\,\big)\right] \otimes \left[\frac{1}{\sqrt{2}}\big(\,|1\rangle_2 \pm e^{-i\bar{\phi}}\,|0\rangle_2\,\big)\right], \tag{4.41}$$

where $|i\rangle_1 \otimes |j\rangle_2 = |ij\rangle$, and where the subscripts signify whether the state belongs to the first or second ion in the chain. The eigenvalues of $S_{\bar{\phi}}$ are $\lambda = \pm 2$ for $\{|+-\rangle, |-+\rangle\}$, and $\lambda = 0$ for $\{|++\rangle, |--\rangle\}$. Thus, the conditional operation results from the motional state associated with the spin states $\{|+-\rangle, |-+\rangle\}$ being driven around closed loops in phase space and generating a phase $-i$, while the motional state associated with $\{|++\rangle, |--\rangle\}$ remains unaffected by the gate radiation, as shown in Fig. 4.3.

In the computational basis, the above truth table turns into

$$|11\rangle \;\mapsto\; \frac{1}{\sqrt{2}}e^{-i\pi/4}\left[\,|11\rangle + e^{-2i\bar{\phi}}i\,|00\rangle\,\right] \tag{4.42}$$

$$|10\rangle \;\mapsto\; \frac{1}{\sqrt{2}}e^{-i\pi/4}\left[\,|10\rangle + i\,|01\rangle\,\right] \tag{4.43}$$

$$|01\rangle \;\mapsto\; \frac{1}{\sqrt{2}}e^{i\pi/4}\left[\,|10\rangle - i\,|01\rangle\,\right] \tag{4.44}$$

$$|00\rangle \;\mapsto\; \frac{1}{\sqrt{2}}e^{i\pi/4}\left[e^{2i\bar{\phi}}\,|11\rangle - i\,|00\rangle\,\right], \tag{4.45}$$

which describes a mapping into the Bell state basis.

4.2.4 Measuring fidelity

We can assess the quality of a quantum gate by using the "fidelity" as a measure of similarity between the intended and the actual output state of the gate. The fidelity $\mathcal{F}$ is defined as

$$\mathcal{F} \equiv \langle\psi|\rho|\psi\rangle\,, \tag{4.46}$$

where the desired output state $|\psi\rangle$ is a pure state, and where ρ is the density operator of the actual state produced by the gate, which may be a mixed state. Conversely, we define the gate error as $\epsilon \equiv 1 - \mathcal{F}$.

In general, the full characterisation of the performance of a quantum gate requires process tomography [RKS+06], that is, measuring the average gate

fidelity over a complete basis of input states. This entails the implementation of the full truth table given by Eq. 4.42-4.45, which is impossible in our experiment. The reasons are, first, that our ion trap is not designed to carry out addressed single-qubit rotations on individual ions, and second, that we are currently relying on a PMT for state detection, which prevents us from distinguishing the $|10\rangle$ state from the $|01\rangle$ state. Instead, we resort to a common, reduced version of process tomography, where we evaluate the gate performance based on a single input state and assume that the measured fidelity is, by symmetry, similar for all remaining inputs.

We pick $|11\rangle$ as our initial state, and write down the fidelity of the associated ideal output state, given by Eq. 4.42, relative to an arbitrary state with a density matrix ρ

$$\begin{aligned}\mathcal{F} &= \frac{1}{2}\left(\rho_{11,11}+\rho_{00,00}\right)+\frac{1}{2}\left(e^{-2i\bar{\phi}}i\rho_{11,00}+c.c\right)\\ &= \frac{1}{2}\left(\rho_{11,11}+\rho_{00,00}\right)+\sin\left(2\bar{\phi}-\arg\left\{\rho_{11,00}\right\}\right)\left|\rho_{11,00}\right|.\end{aligned} \tag{4.47}$$

To arrive at the above expression, we made use of the Hermiticity of the density matrix $\rho_{11,00} = \rho^*_{00,11}$, and introduced the abbreviation $\rho_{ii,jj} = \langle ii|\rho|jj\rangle$. The phase of the intended Bell state $2\bar{\phi}$ and the phase of the actual gate output $\arg\{\rho_{11,00}\}$ should ideally match, so that

$$\mathcal{F} = \frac{1}{2}\left(\rho_{11,11}+\rho_{00,00}\right)+\left|\rho_{11,00}\right|. \tag{4.48}$$

In practice, we can stick with Eq 4.48 even if our gate produces a Bell state with the wrong phase, since such an error can be straightforwardly corrected using global single-qubit rotations, provided the phase error remains constant in time.

We measure $\mathcal{F}$ using two separate experiments. First, we determine the left-hand term of the sum in Eq. 4.48 by projectively measuring the state of the qubits after the gate is completed. To this end, we use our standard readout method described in section 2.6. Adding the measured population in $|00\rangle$, labelled P_{00}, to

that measured in $|11\rangle$, labelled P_{11}, yields

$$P_{11} + P_{00} = \mathrm{Tr}\left\{|11\rangle\langle 11|\,\rho\right\} + \mathrm{Tr}\left\{|00\rangle\langle 00|\,\rho\right\} = \rho_{11,11} + \rho_{00,00}. \tag{4.49}$$

In the ideal case, we obtain $P_{00} = P_{11} = 1/2$, meaning that the contribution to the overall fidelity from the $1/2\,(\rho_{11,11} + \rho_{00,00})$ term in Eq. 4.48 is at most $1/2$.

Second, we determine the right-hand term of the sum in Eq. 4.48, by interposing a global $\pi/2$ rotation[1] $R(\pi/2, \phi_a)$ with phase ϕ_a between the gate and the readout sequence. This additional rotation is also known as the "analysis" pulse. Measuring the populations P_{11}, P_{00}, and $P_{10} + P_{01}$, where the latter corresponds to finding one of the ions in $|0\rangle$ and the other in $|1\rangle$, allows for the construction of the so-called "parity"

$$\begin{aligned}\mathcal{P} &\equiv P_{11} + P_{00} - (P_{10} + P_{01}) \\ &= \mathrm{Tr}\left\{\left[\,|11\rangle\langle 11| + |00\rangle\langle 00| - |10\rangle\langle 10| - |01\rangle\langle 01|\,\right] R(\pi/2,\phi_a)\rho R^{\dagger}(\pi/2,\phi_a)\right\} \\ &= 2\mathrm{Re}\left\{\rho_{10,01}\right\} - 2\left|\rho_{11,00}\right|\cos(2\phi_a + \arg\left\{\rho_{11,00}\right\}).\end{aligned} \tag{4.51}$$

As we scan the phase of the analysis pulse, the parity oscillates at a rate $2\phi_a$ between ± 1, with a phase offset that is given by $\arg\{\rho_{11,00}\}$ (see Fig. 7.13 in section 7.4). The magnitude of $\rho_{11,00}$, which characterises the coherence between $|00\rangle$ and $|11\rangle$, can be deduced from the contrast of this oscillation

$$\left|\rho_{11,00}\right| = \frac{1}{4}\left(\max\left\{\mathcal{P}\right\} - \min\left\{\mathcal{P}\right\}\right). \tag{4.52}$$

For a perfect gate, $\left|\rho_{11,00}\right| = 1/2$ holds. Thus, the contribution to the overall fidelity from the right-hand term of the sum in Eq. 4.48 is, again, at most $1/2$.

[1] The two-ion rotation matrix is given by

$$\begin{aligned}R(\theta,\phi) &= R_1(\theta,\phi) \times R_2(\theta,\phi) \\ &= \begin{bmatrix}\cos(\frac{\theta}{2}) & -ie^{i\phi}\sin(\frac{\theta}{2}) \\ -ie^{-i\phi}\sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2})\end{bmatrix}_1 \times \begin{bmatrix}\cos(\frac{\theta}{2}) & -ie^{i\phi}\sin(\frac{\theta}{2}) \\ -ie^{-i\phi}\sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2})\end{bmatrix}_2,\end{aligned} \tag{4.50}$$

where $R_n(\theta,\phi)$ is a one-ion rotation on the n^{th} ion. For a $\pi/2$-pulse driven by resonant radiation with phase ϕ, the pulse duration is set so that $\theta = \Omega t = \pi/2$, with Rabi frequency Ω.

Finally, we note that an alternative method for determining gate fidelities is randomized benchmarking, proposed in [EAZ05] and first implemented using trapped ions in [KLR+08]. It involves the execution of long sequences of randomly chosen gates from a specified set. The decay in average fidelity with which a given sequence produces the correct final state is then measured as a function of the sequence length. It is possible to isolate the fidelity of a particular kind of gate by measuring the excess error from interleaving additional gates of the type under investigation within the benchmarking sequence [GMT+12].

Although the above method has certain advantages over process tomography, our current two-qubit gate is too slow to be incorporated into sufficiently long sequences for randomized benchmarking to be feasible. Moreover, it would also require methods of individually addressing the ions, as well as re-cooling them in-between gate operations without disturbing the qubit coherence [BDS+03].

4.3 Error analysis

In this section we turn to the investigation of the most relevant sources of error occurring in the Mølmer-Sørensen gate, the way it was implemented in this work. These can be classified into two main categories, coherent and incoherent errors.

The former do not stop the two-ion system's dynamics from being unitary, but nevertheless cause the final state of the gate to deviate from the intended output. Such errors are usually the result of poorly calibrated experimental parameters, or of parameters that drift on timescales longer than the gate duration. Examples include offsets in the gate time, detuning, Rabi frequency or AC-Zeeman shift compensation. We determine the impact of coherent errors on the gate's fidelity by, if need be, modifying the basic gate Hamiltonian to include the error source under investigation, and subsequently integrating the corresponding Schrödinger equation.

In contrast, incoherent errors lead to non-unitary and irreversible dynamics.

They are caused by stochastic processes, i.e. noisy experimental parameters that fluctuate on timescales comparable to the operation time of the gate. Examples include motional heating and motional dephasing. We model incoherent errors using a master equation approach, which gives rise to the following Lindblad equation [Lin76]

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[\tilde{H}_{\mathrm{MS}}, \rho] + \sum_k \left\{ L_k \rho L_k^\dagger - \frac{1}{2}\left(\rho L_k^\dagger L_k + L_k^\dagger L_k \rho \right) \right\}. \tag{4.53}$$

The Lindblad operators L_k are specific to the type of error under investigation, while $\tilde{H}_{\mathrm{MS}}$ is given by Eq. 4.28.

4.3.1 Acquiring the wrong geometric phase - Rabi frequency fluctuations

We have seen in section 4.2.3 that in order for the Mølmer-Sørensen gate to be implemented correctly, the three parameters δ, $\Omega^{(2)}$, and Δt must comply with the relations laid down in Eq. 4.38 and Eq. 4.39. If the latter of the two is not satisfied, then the accumulated geometric phase deviates from its ideal value $\lambda^2 \Phi(\Delta t_g, 0) = \pi/2$. The two possible causes for this are, first, a miscalibration or fluctuations of the Rabi frequency, leading to pure geometric phase errors if no other sources of infidelity are present. Second, a miscalibration or fluctuations of the gate detuning, which additionally upset Eq. 4.38 and thereby result in both a geometric phase error and an error from the loops in phase space not closing. The latter of the two cases will be discussed in more detail in the next section. Here, we focus on the basic Mølmer-Sørensen gate in the presence of a pure geometric phase error, for which the fidelity is given by

$$\mathcal{F} = \frac{1}{2} + \frac{1}{2}\sin\left[2\pi K \left(\frac{\Omega^{(2)}}{\delta} \right)^2 \right]. \tag{4.54}$$

The microwave sideband power, which determines the Rabi frequency $\Omega^{(2)}$, is the first parameter that we choose in the gate set-up procedure (see section 7.3). It

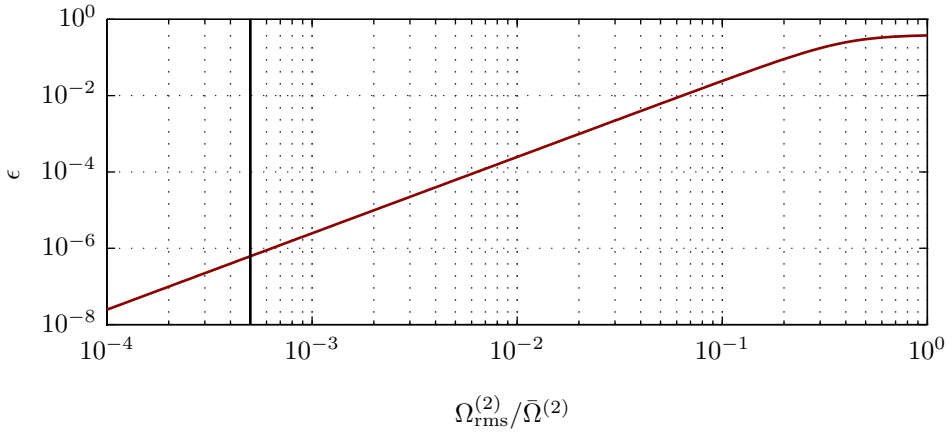

Figure 4.4: Simulated Mølmer-Sørensen gate error caused by normally-distributed shot-to-shot fluctuations of the Rabi frequency with an r.m.s. deviation $\Omega^{(2)}_{\mathrm{rms}}$ around its average $\bar{\Omega}^{(2)}$. The gate time, detuning and average Rabi frequency are assumed to be constant, and are set as prescribed by Eq. 4.38 and Eq. 4.39. The resulting infidelity is a consequence of accumulating the wrong amount of geometric phase, and does not depend on the ions' temperature or the number of loops traversed in phase space. Indicated by the black line is the mean fractional r.m.s. deviation of the Rabi frequency fluctuations estimated in this work, amounting to $\sim 5 \times 10^{-4}$, and therefore giving rise to a negligible gate error at the 10^{-6} level.

is the basic parameter with reference to which Δt and δ are subsequently adjusted. Thus, we only have to worry about $\Omega^{(2)}$ changing over time, rather than it having any constant calibration offsets. Variations in the Rabi frequency are usually a consequence of unstable sideband power, which is likely to be caused by thermal changes of the properties of components used in the microwave supply chain, or thermal changes within the ion trap itself. In this work, we eliminate long term microwave power drifts using a feedback loop, described in section 5.4.3, leaving only small residual fluctuations of $\Omega^{(2)}$ around its stable mean $\bar{\Omega}^{(2)}$, from shot to shot of the gate experiment.

In order to determine the effect of such fluctuations on the fidelity of the basic Mølmer-Sørensen gate, we evaluate Eq. 4.54 for 10^6 randomly sampled values of $\Omega^{(2)}$ from a normal distribution with an r.m.s. deviation of $\Omega^{(2)}_{\mathrm{rms}}$ around $\bar{\Omega}^{(2)}$. We assume that $\bar{\Omega}^{(2)}$ conforms to Eq. 4.39, and that the two sideband powers do not change independently (this case will be treated separately, in section 4.3.3). The

simulation results are plotted in Fig. 4.4 and show the gate error as a function of the fractional r.m.s. deviation of the Rabi frequency.

In our experiment, we are able to keep the Rabi frequency stable at the sub-Hertz level, as can be seen from the power stability measurements presented in section 8.7. More specifically, we estimate $\Omega^{(2)}_{\mathrm{rms}}/\bar{\Omega}^{(2)} \approx 5 \times 10^{-4}$, which according to the above simulation corresponds to a negligible gate error of approximately 1×10^{-6}. It is worth noting that this error remains unaffected when modifying the basic Mølmer-Sørensen scheme to incorporate Walsh modulation or dynamical decoupling, which will both be introduced later in this chapter.

4.3.2 Unclosed loops in phase space - timing miscalibration and detuning fluctuations

If $\Delta t \neq (2\pi K)/\delta$, because of an offset in δ or Δt, the phase space trajectories are left open, and the internal states of the ions remain entangled with their motional state at the end of the gate operation. To determine how this effect translates into a gate error, we first write down the fidelity of the basic Mølmer-Sørensen scheme for general δ, $\Omega^{(2)}$, and Δt

$$\mathcal{F} = \frac{1}{8}\left\{3 + e^{-16|\alpha(\Delta t,0)|^2\left(\bar{n}+\frac{1}{2}\right)} + 4e^{-4|\alpha(\Delta t,0)|^2\left(\bar{n}+\frac{1}{2}\right)}\sin\big(4\Phi(\Delta t,0)\big)\right\}. \quad (4.55)$$

In deriving the above expression, we have assumed that the ions' spins start in $|11\rangle$, and that the motional mode is in an initial thermal state of mean occupancy $\bar{n}$. We note that in the special case of $\alpha(\Delta t, 0) = 0$, i.e. when the loops close, Eq. 4.55 reduces to Eq. 4.54, which describes pure geometric phase errors only.

Fig. 4.5 shows the gate error as a function of a constant timing miscalibration, which in practice results from inaccuracies of setting the appropriate Δt for a given $\Omega^{(2)}$. In this work, we find that the systematic, fractional offset of Δt from its ideal value is 0.021(4), as assessed in section 8.6. This leads to an error of 0.2(1)% for a single-loop, basic Mølmer-Sørensen gate, and an initial motional mode temperature of $\bar{n} = 0.1$.

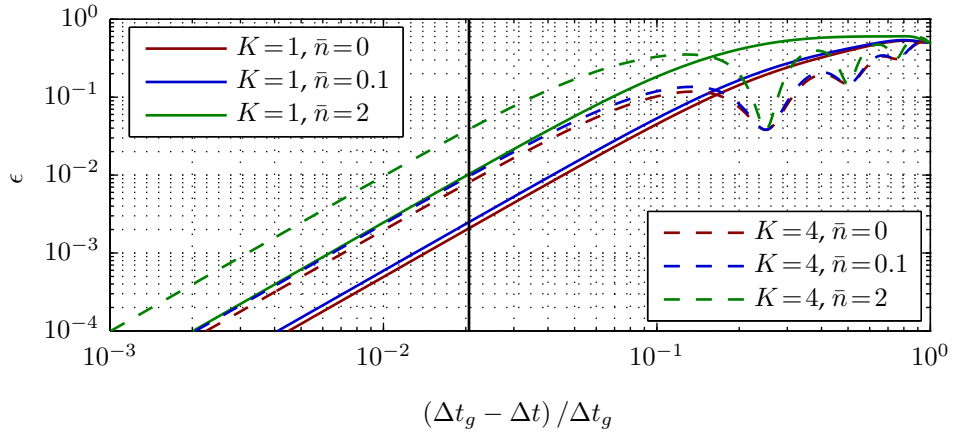

Figure 4.5: Basic Mølmer-Sørensen gate error caused by a constant gate time offset from the optimal value Δt_g. The plot is generated by evaluating Eq. 4.55, with δ fixed at its respective optimum value δ_g. The gate infidelity stems from both an incorrect geometric phase, and the imperfect closure of the loops in motional phase space. We estimate the gate time in our experiment to have a systematic, fractional miscalibration error of 0.021(4), indicated by the black line. This corresponds to a substantial gate error of 0.2(1)% for a single loop in phase space, and 1.0(4)% for $K = 4$ loops in phase space, when assuming an initial ion temperature of $\bar{n} = 0.1$.

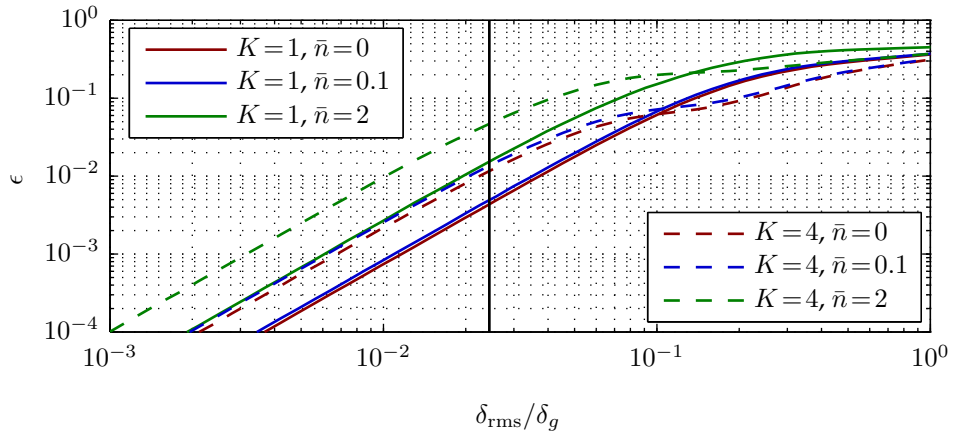

Figure 4.6: Basic Mølmer-Sørensen gate error as a function of the r.m.s. deviation δ_{rms} of the normally-distributed shot-to-shot fluctuations in the gate detuning around its ideal value δ_g. For the simulation at hand, the gate time and Rabi frequency were kept constant at $\Delta t = (\pi\sqrt{K})/\Omega^{(2)}$. As in Fig. 4.5, the depicted gate error stems from an incorrect geometric phase, as well as unclosed loops in motional phase space. The black line indicates the fluctuations measured in this experiment, that is, $\delta_{\mathrm{rms}}/\delta_g \approx 30\,\mathrm{Hz}$ r.m.s./$1.2\,\mathrm{kHz}$, which corresponds to a single-loop gate error at the $\sim 0.5\%$ level.

In contrast, the deviations of the gate detuning are subject to fluctuations, arising from frequency instabilities of the motional mode on which the gate is performed (see section 8.2). With regard to timescales, frequency jitter that is faster than the gate time results in motional dephasing, which will be treated separately in section 4.3.7, while slow drifts are generally taken out by manual adjustments. Here, we model the remaining variations, that is, those that are fast enough to cause δ to change from shot to shot of the gate experiment, but not during the gate pulse itself. Fig. 4.6 illustrates the associated gate error, simulated by repeatedly evaluating Eq. 4.55 for 10^6 values of δ, drawn randomly from a normal distribution with an r.m.s. deviation δ_{rms} around δ_g. From the measurements presented in section 8.2, we estimate the fractional detuning fluctuations to be $\sim$0.025 r.m.s., corresponding to an error at the $\sim 0.5\%$ level for a single loop.

In summary, both the timing and the detuning error discussed above are unacceptably high, prompting us to consider Walsh modulation (section 4.4.1) as a method of suppressing them.

Two final remarks about Fig. 4.5 and Fig. 4.6. First, for a given Rabi frequency, it is generally more favourable to accumulate the geometric phase by means of multiple smaller loops in phase space rather than a single bigger one. Second, the required level of precision in gate time and detuning exponentially increases with the initial thermal occupation $\bar{n}$ of the ion chain's motion. Consequently, it is advantageous to keep the vibrational mode as cold as possible. Moreover, the gate's sensitivity to many other sources of infidelity also decreases with temperature. In this work, we therefore employ continuous (laser-driven) resolved sideband cooling techniques to reach motional mode temperatures around $\bar{n} = 0.1$. For completeness, we will in future sections also mention the gate errors for $\bar{n} = 2$, as achieved by dark-resonance cooling only.

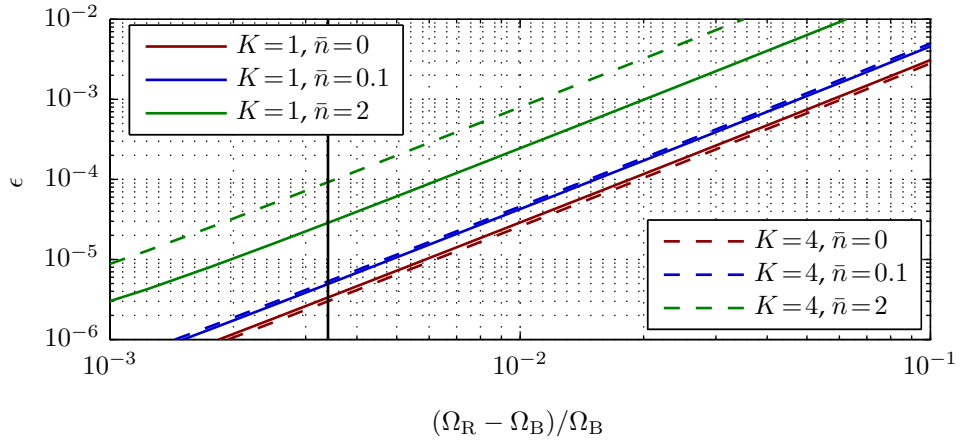

Figure 4.7: Simulated error of a basic Mølmer-Sørensen gate, stemming from unequal red and blue sideband Rabi frequencies, symmetrically offset about their ideal value $\Omega^{(2)}$. All remaining parameters are assumed to be optimally tuned to $\Omega^{(2)}$, following Eq. 4.38 and Eq. 4.39. The black line indicates the typical point of operation in this work.

4.3.3 Rabi frequency imbalance

To model the gate error resulting from differing Rabi frequencies of the two components of the bichromatic field, we assume for simplicity, that they are symmetrically offset about $\Omega^{(2)}$. In other words, we scale the Rabi frequency of the blue sideband from Eq. 4.22 by $(1-\beta)$, so that $\Omega_B = (1-\beta)\Omega^{(2)}$, and that of the red sideband from Eq. 4.23 by $(1+\beta)$, so that $\Omega_R = (1+\beta)\Omega^{(2)}$. Numerical integration of the associated Schrödinger equation allows us to plot the gate error as a function of the fractional discrepancy between the two sideband strengths. Fig. 4.7 illustrates the result of this integration, when considering two spin states and a single motional mode, truncated to 40 harmonic oscillator Fock states.

We truncate the Fock state space of the ions' harmonic motion for the sake of simulation speed. The cutoff is chosen so that the truncation error is negligible within the region of interest of the simulated quantity, which in this case is the gate error. As expected, simulating gates that start out with high thermal occupations of their motional mode, require more Fock states to be included in the numerical evaluation compared with ground state cooled gates.

Unbalanced Rabi frequencies are the consequence of unbalanced microwave power, which in our experiment can be equalised and sustained to within 0.03 dB. In section 8.7, we measure that our best gate data was taken with Ω_{R} deviating from Ω_{B} by ~ 1 Hz. More specifically, the fractional discrepancy was $3.4(6) \times 10^{-3}$, which according to the above simulation gives a negligible gate error of $5(2) \times 10^{-6}$ for an initial motional mode temperature of $\bar{n} = 0.1$. When using the dynamically-decoupled Mølmer-Sørensen scheme, which will be explained in section $4.4.2$, instead of the basic gate scheme, this error is pushed down even further, to $2.8(3) \times 10^{-7}$.

4.3.4 Off-resonant excitation - Effect on populations

An unwanted byproduct when generating the oscillating microwave magnetic field gradient, is the accompanying field amplitude, which can off-resonantly drive transitions other than those intended for the gate. This results in undesired (though coherent) population transfer between the qubit states, as well as between the qubit and spectator states, as illustrated in Fig. 4.8.

In section $5.1.2$ we characterise the magnetic field poduced by our trap, from which we infer that the field amplitude accompanying the field gradient gives rise to a Rabi frequency of 0.37 MHz on the qubit and of 14 MHz on π-spectator transitions. The Rabi frequencies on the σ_- spectator transitions are similar to that on the qubit, up to small differences in the associated transition matrix elements.

These off-resonant Rabi oscillations have a considerable amplitude ($\sim 10\%$), causing the populations to periodically leave the qubit subspace. To ensure that by the end of the gate operation, all population has returned to its appropriate state, we work with pulse shapes that provide smooth switching of the gate microwaves over periods of 3 μs. An additional benefit of pulse shaping is the prevention of spectral broadening, which usually appears as a consequence of

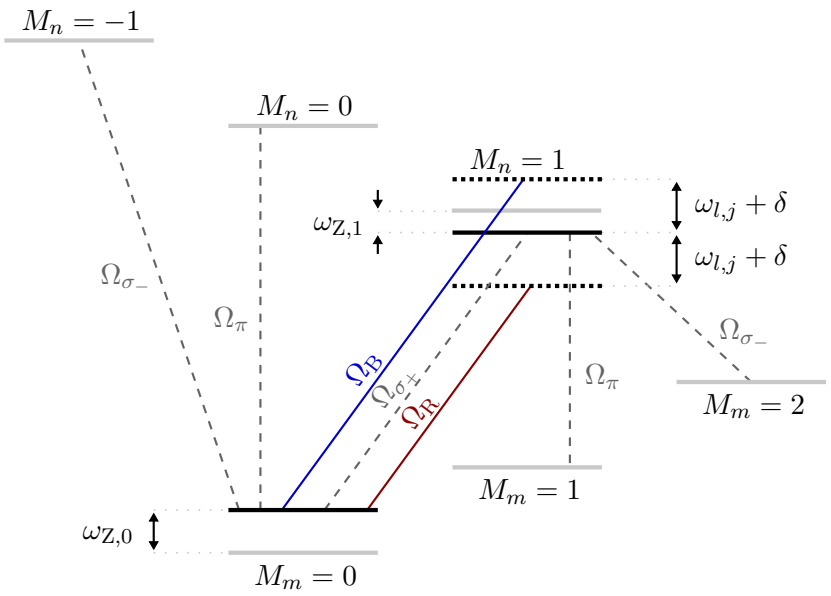

Figure 4.8: Energy level diagram (not to scale) of part of the $4S_{1/2}$ ground level in $^{43}Ca^+$ at 146 G, showing the qubit states and the four other states connected to them via spectator transitions. The Rabi frequencies caused by the field amplitudes accompanying the microwave field gradients, as well as the AC-Zeeman shifts $\omega_{Z,0}$ and $\omega_{Z,1}$, are indicated. In this experiment, $\Omega_{\sigma_+} \approx \Omega_{\sigma_-} = 0.37\,\text{MHz}$, and $\Omega_\pi = 14\,\text{MHz}$.

the abrupt start and stop of the field emission in a rectangular pulse.

However, off-resonant excitation may also be the consequence of spectral impurities of the driving field. As an example, frequency spurs close to spectator transitions connecting to the qubit states can lead to population escaping the qubit subspace. Here, the solution lies in improving the spectral purity of the microwaves by e.g. filtering (see section 5.4.2.1).

As explained in section 8.4, our attempt to independently measure the amount of off-resonant excitation present in our experiment was limited by the statistical sensitivity of the underlying measurement, yielding an upper bound of $\lesssim 10^{-3}$.

4.3.5 Off-resonant excitation - AC-Zeeman shifts

AC-Zeeman shifts have a similar origin as off-resonant excitation, but instead of affecting the populations, they offset in the qubit's energy levels. The $|0\rangle$ state is shifted by $\omega_{Z,0}$, while the $|1\rangle$ state is shifted by $\omega_{Z,0}$, as illustrated in Fig. 4.8. We

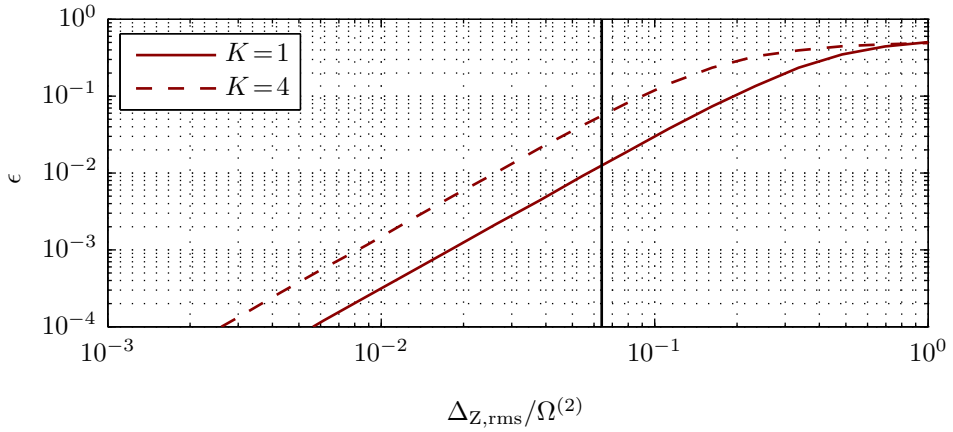

Figure 4.9: Simulated basic Mølmer-Sørensen gate error due to normally-distributed shot-to-shot fluctuations of an uncompensated AC-Zeeman shift, with all other aspects of the gate configured optimally. The black line corresponds to a fractional r.m.s. deviation of the fluctuations $\Delta_{\mathrm{Z,rms}}/\Omega^{(2)} = 19.7\,\mathrm{Hz}$ r.m.s./$308\,\mathrm{Hz}$, as measured in our experiment. Both graphs correspond to an initial thermal occupation of the motional mode of $\bar{n} = 0$, which at the black line gives a gate error $\sim 3 \times 10^{-4}$ below the $\bar{n} = 0.1$ case for $K = 1$, and $\sim 6 \times 10^{-4}$ below the $\bar{n} = 0.1$ case for $K = 4$.

label the net AC-Zeeman shift on the qubit transition $\omega_{\mathrm{Z},0} + \omega_{\mathrm{Z},1}$ as ω_Z. As long as this shift is constant, it can be measured and taken into account by adjusting the two microwave field frequency tones to $\omega_{\mathrm{B}} = (\omega_0 + \omega_{\mathrm{Z}}) + (\omega_{l,j} + \delta)$ and $\omega_R = (\omega_0 + \omega_{\mathrm{Z}}) - (\omega_{l,j} + \delta)$. Any residual, uncompensated AC-Zeeman shifts lead to frequency deviations Δ_{Z} of the drive field, which can be modelled by adding the term

$$\tilde{H}_{\mathrm{Z}} = \frac{1}{2}\hbar\Delta_{\mathrm{Z}} \sum_{n=1}^{N} \sigma_{z,n}, \tag{4.56}$$

to $\tilde{H}_{\mathrm{MS}}(t)$ from Eq. 4.28.

In practice, Δ_{Z} results from fluctuations of ω_Z, since they are generally the only part of the AC-Zeeman shift that cannot be effectively taken out. For the purpose of modelling the impact of such fluctuations on the gate's fidelity, we assume them to be slowly varying compared with the gate time, and normally-distributed. We then integrate the Schrödinger equation associated with $\left[\tilde{H}_{\mathrm{MS}}(t) + \tilde{H}_{\mathrm{Z}}\right]$ for 10^4 values for Δ_{Z}, randomly sampled from a normal distribution with an r.m.s. deviation $\Delta_{\mathrm{Z,rms}}$. The integration is done numerically for two spin states and one

motional mode, truncated to 30 harmonic oscillator Fock states, in the absence of any other sources of infidelity. Fig. 4.9 shows the resulting gate error as a function of $\Delta_{\mathrm{Z,rms}}$. It bears similarity to Fig. 4.6, which shows the impact of a gate detuning miscalibration, since both effects boil down to a frequency error of the drive field. The difference is that uncompensated AC-Zeeman shifts offset both ω_{R} and ω_{B} with the same sign, whereas an error in gate detuning contributes with opposite sign.

Although we implemented a feedback loop to control the stability of the microwave power (see 5.4.3), the AC-Zeeman shift for our final gate experiments still showed shot-to-shot fluctuations of $\Delta_{\mathrm{Z,rms}} = 2\pi \times 19.7(8)\,\mathrm{Hz}$ r.m.s. (see section 8.5) about $\omega_{\mathrm{Z}} = 2\pi \times 20.78\,\mathrm{kHz}$. The corresponding simulated error for a basic MS gate with $K = 4$ loops is $5.6(5) \times 10^{-2}$, as indicated by the black line in Fig. 4.9. This error is unacceptably high, which is why we employ the technique of dynamical decoupling, introduced in section 4.4.2, to reduce it.

4.3.6 Motional Heating

For most of the gate duration, the ions' internal states are entangled with their motional state, allowing perturbations that affect the motion to cause gate errors. A significant such perturbation, already introduced in section 3.2.2, is caused by electric field noise from the trap's surface, which leads to uncontrolled heating of the ion chain's motional mode during the gate operation.

This process is therefore an incoherent source of error, and we model it, following [Har14], by setting $L_1 = \sqrt{\dot{\bar{n}}}\, a_{l,j}$ and $L_2 = L_1^\dagger$ in Eq. 4.53. The resulting master equation reads

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[\tilde{H}_{\mathrm{MS}}, \rho] + \frac{1}{2}\dot{\bar{n}}\left(2a^\dagger \rho a + 2a\rho a^\dagger - aa^\dagger\rho - a^\dagger a\rho - \rho a^\dagger a - \rho a a^\dagger\right), \quad (4.57)$$

where $\dot{\bar{n}}$ is the linear heating rate of the gate mode in quanta/s, and where we have dropped the indices l, j associated with the ladder operators, for convenience. The results from numerically integrating this master equation for two

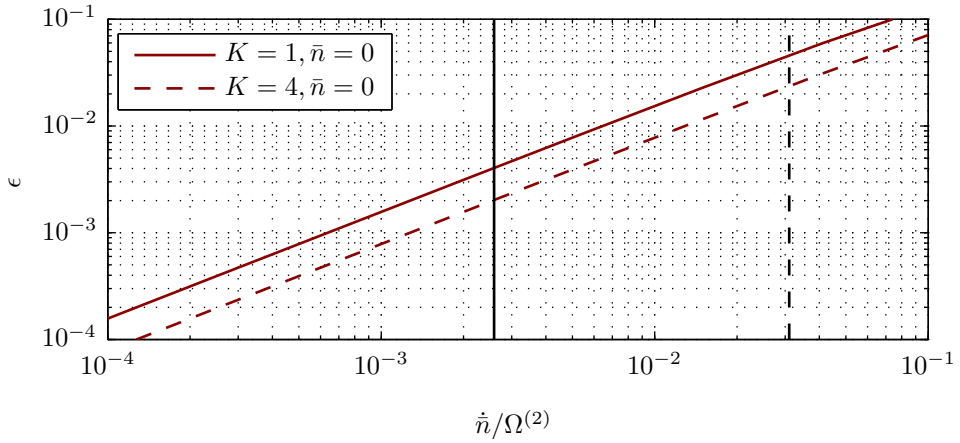

Figure 4.10: Simulated gate error due to motional heating, assuming no other sources of infidelity are present. The solid black line indicates the 5 quanta/s upper limit of the heating rate of the two-ion radial rocking mode used in this work. The dashed black line corresponds to the 60 quanta/s heating rate of the two-ion radial COM mode. We normalise the heating rates with respect to $\Omega^{(2)} = 2\pi \times 308\,\mathrm{Hz}$. Initial temperatures other than $\bar{n} = 0$ have only minimal effect on the gate fidelity. More specifically, when starting at $\bar{n} \approx 2$ the extra error at the solid black line is $\sim 1 \times 10^{-3}$ for both $K = 1$ and $K = 4$. When starting at $\bar{n} \approx 0.1$ the additional error is at the 10^{-8} level, and hence, insignificant.

spin states and one motional mode, truncated at 100 harmonic oscillator Fock states, are plotted in Fig. 4.10.

In this work, we perform our gate on the two-ion radial rocking mode ($b^{(l,2)} = \frac{1}{\sqrt{2}}(-1, 1)$), since its $\lesssim 5$ quanta/s heating rate is significantly lower than the 60(15) quanta/s measured for two-ion radial COM mode ($b^{(l,1)} = \frac{1}{\sqrt{2}}(1, 1)$). Details about the measurements of the above stated heating rates can be found in section 8.1. The simulated gate error associated with a 5 quanta/s heating rate amounts to $\lesssim 0.2\%$ for $K = 4$ loops, while it is a factor of 2 worse for a single loop, and about an order of magnitude worse for the two-ion radial COM mode. These errors are not affected when adding Walsh modulation or dynamical decoupling to the basic Mølmer-Sørensen scheme (see sections 4.4.1 and 4.4.2).

4.3.7 Motional decoherence

Processes that cause an ion chain's motional superposition state to dephase while it is entangled with its spins, as is the case during the gate operation, represent an incoherent source of error. An example of such a process are fluctuations of the vibrational mode's frequency over timescales comparable to the gate duration. We model the error using Eq. 4.53 with $L = \sqrt{2/\tau_{\mathrm{C}}}\, a^{\dagger}_{l,j} a_{l,j}$, as explained in e.g. [TMK+00, Har14]. The master equation becomes

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[\tilde{H}_{\mathrm{MS}}, \rho] + \frac{1}{\tau_{\mathrm{C}}}\left(2a^{\dagger}a\rho a^{\dagger}a - \left(a^{\dagger}a\right)^2\rho - \rho\left(a^{\dagger}a\right)^2\right), \tag{4.58}$$

and describes a decay in coherence between two motional states $|n\rangle$ and $|n'\rangle$ at a rate $(n-n')^2/\tau_{\mathrm{C}}$. The "motional coherence time" τ_{C} is therefore the coherence time for a $[|n\rangle + |n'\rangle]$ motional superposition state.

After integrating Eq. 4.58 numerically, in a similar way as in the previous section, we find the gate error as a function of τ_{C}, which is plotted in Fig. 4.11. The motional coherence time in this work has not been measured directly, but is estimated to have an upper bound of 400 ms, as explained in section 8.3. According to the simulation, this corresponds to an error of $\lesssim 1.1 \times 10^{-3}$ for $K = 4$ loops starting in $\bar{n} = 0$, and $\lesssim 1.3 \times 10^{-3}$ for $\bar{n} = 0.1$. We note that these errors do not change when switching from the basic Mølmer-Sørensen scheme to a Walsh- modulated or dynamically-decoupled gate.

4.4 Error suppression

In this section, we discuss methods for increasing the two-qubit gate's robustness to certain sources of error. Techniques that have already been mentioned include employing a clock transition to minimise the impact of external fluctuations in the static magnetic field, and reducing effects of thermal motion of the ion chain by laser cooling it close to the ground state. Here, we introduce two crucial modifications to the standard Mølmer-Sørensen sequence, each suppressing a

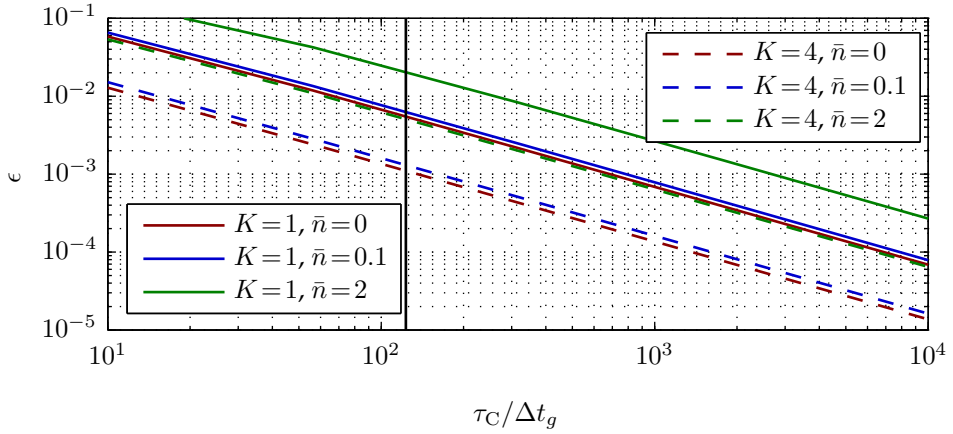

Figure 4.11: Simulated basic Mølmer-Sørensen gate error as a function of motional coherence time in the absence of other sources of infidelity. The black line indicates the estimated upper bound for $\tau_C/\Delta t_g \approx 400\,\text{ms}/3.25\,\text{ms}$ in our experiment, corresponding to a gate error of 1.3×10^{-3} for $K = 4$ and $\bar{n} = 0.1$.

major source of infidelity in our gate, namely, fluctuations of the gate detuning and of uncompensated AC-Zeeman shifts.

4.4.1 Composite multi-loop gate sequences

As we have seen in section 4.3.2, timing and detuning errors prevent us from ensuring that the motional states trace out closed loops in phase space and cause significant gate errors of order ∼0.5%. Although the reduction of timing errors is usually a matter of more careful calibration, bringing detuning fluctuations under control is hard due to the difficulty of stabilising the ion trap's radial modes. This is further complicated by the weak coupling between the ions and the microwave magnetic field. The implications are, first, a long duration of the gate, and second, a small ($\sim 1\,\text{kHz}$) gate detuning, which means that the radial mode frequency needs to be stable at the $\sim 10\,\text{Hz}$ level.

As a remedy, we use composite multi-loop pulse sequences that minimise residual displacements in phase space [HCD+12]. We have already seen in section 4.3.2 that, for a given Rabi frequency, acquiring the geometric phase by means of multiple small loops instead of a single big loop can be beneficial. An

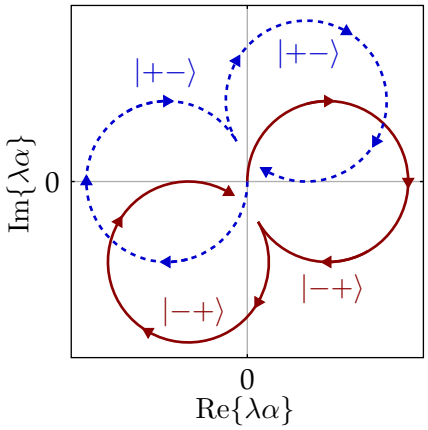

Figure 4.12: Phase space representation of a two-loop gate in the presence of a timing or detuning error, when modulated according to a $W_1(t/\Delta t)$ Walsh function, using π-phase shifts of $\bar{\phi}$.

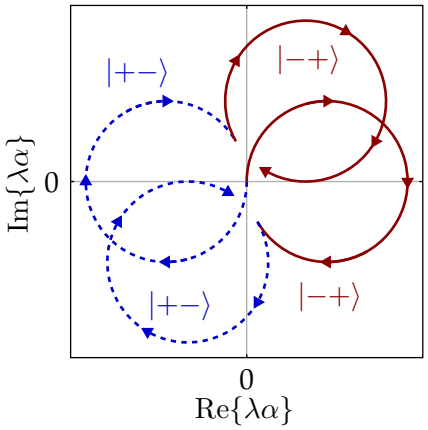

Figure 4.13: Phase space representation of a similar gate to that shown in Fig. 4.12, when modulated according to a $W_1(t/\Delta t)$ Walsh function, using global π-pulses instead of phase shifting $\bar{\phi}$.

extension of this idea are pulse sequences with interleaved sign inversions of the spin-dependent force, or alternatively, spin flips of both qubits. They are inserted between particular loops, as prescribed by a Walsh function $W_k(x)$ [HKVB11].

This function alternates between ± 1 depending on the index k, as can be seen in Fig. 4.14. It is defined on on an interval $0 \leq x \leq 1$, which in our case represents the duration Δt of the K-loop gate. Every zero-crossing of the Walsh function indicates either a sign change of the spin-dependent force, or a global spin flip. In practice, the former is realised by a π-phase shift of $\bar{\phi}$, i.e. a π-phase shift of both ϕ_{R} and ϕ_{B}, and the latter by performing a global π-pulse around an axis set by $\bar{\phi} \pm \pi/2$.

Phase shifts produce mappings $S_{\bar{\phi}} \mapsto S_{\bar{\phi}+\pi} = (-1)S_{\bar{\phi}}$. Since the timing of these shifts is given by the Walsh function, the above mappings are equivalent to $S_{\bar{\phi}} \mapsto W_k(t/\Delta t)S_{\bar{\phi}}$, and therefore correspond to sign inversions of the Hamiltonian. To illustrate how these inversions can be used for error suppression, consider Fig. 4.12, showing the phase space trajectories of a two-loop gate modulated by the Walsh function $W_1(t/\Delta t)$ in the presence of an error in Δt or δ. The sign inversion of the Hamiltonian happens after the first loop and reverses the

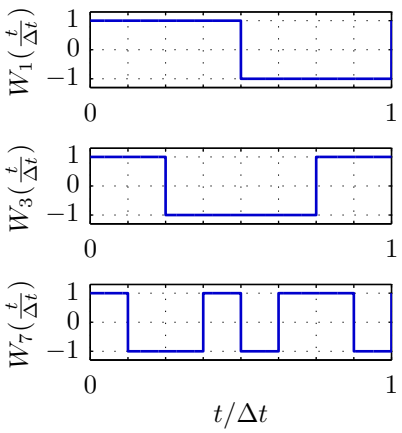

Figure 4.14: Three Walsh functions $W_k(t/\Delta t)$ plotted for modulation indices $k = \{1, 3, 7\}$ on the interval $0 \leq t \leq \Delta t$. In the absence of timing or detuning errors, the zero-crossings fall between completed loops in phase space.

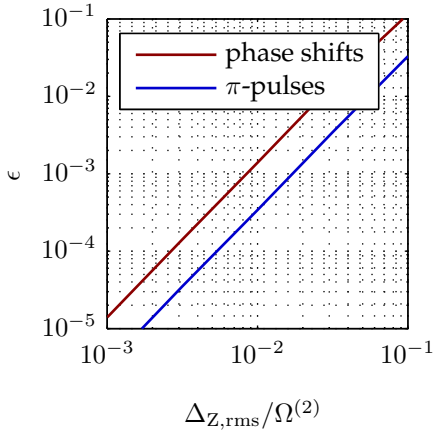

Figure 4.15: Gate error due to a fluctuating uncompensated AC-Zeeman shift, when modulated by $W_1(t/\Delta t)$ using a π-pulse, and a π-phase shift of $\bar{\phi}$. The latter case turns out to be identical to the unmodulated case. The simulation was set up as in section 4.3.5 for a $K = 4$ loop gate starting in $\bar{n} = 0$.

direction of the spin-dependent force. As a result, the phase space trajectories terminate closer to their point of origin by the end of the second loop, compared with the unmodulated scenario, thus improving the gate fidelity.

Global π-pulses are the inverse approach, in that the direction of the spin-dependent force remains unchanged, but the spins of both qubits are flipped (Fig. 4.13). This method has the same net effect as the phase shifts of $\bar{\phi}$ when it comes to timing and detuning errors. However, the spin flips have the additional benefit of providing some robustness against fluctuating AC-Zeeman shifts, as can be seen from the simulation results in Fig. 4.15. Although this discrepancy is, to a large degree, balanced out again when employing the dynamical decoupling scheme, which will be introduced in the next section, we nevertheless choose this method over $\bar{\phi}$ phase shifts in our experiment. The reason being that π-pulses also serve as refocussing pulses for partially complete carrier Rabi oscillations, which will prove useful for the dynamical decoupling scheme.

We will now quantify the error suppression capacity of the Walsh modulated

Mølmer-Sørensen (WMMS) scheme by deriving an expression for the associated gate fidelity. For convenience, we will stick to the method of phase shifting $\bar{\phi}$ as it is mathematically simpler and leads to the same end result. The corresponding gate Hamiltonian reads

$$\tilde{H}_{\mathrm{WMMS}}(t) = \frac{1}{2}\hbar\Omega^{(N)} W_k(t/\Delta t) S_{\bar{\phi}}\Big(a_{l,j} e^{i\delta t} e^{-i\frac{1}{2}\Delta\phi} + a_{l,j}^{\dagger} e^{-i\delta t} e^{i\frac{1}{2}\Delta\phi}\Big), \tag{4.59}$$

which is similar to Eq. 4.28, except that it follows the sign-swapping pattern given by the Walsh function. The corresponding gate propagator $\tilde{U}_{\mathrm{WMMS},\bar{\phi}}$ can be written as a series of propagators $\tilde{U}'_{\mathrm{MS},\bar{\phi}}$ for every individual loop. Each of them is similar to the basic Mølmer-Sørensen propagator $\tilde{U}_{\mathrm{MS},\bar{\phi}}$, given in Eq. 4.34, but with $S_{\bar{\phi}}$ replaced with $W_k(t/\Delta t)S_{\bar{\phi}}$ to take the sign alternation into account. The WMMS propagator reads

$$\tilde{U}_{\mathrm{WMMS},\bar{\phi}}(t_K, t_0) = \tilde{U}'_{\mathrm{MS},\bar{\phi}}(t_K, t_{K-1})\, \tilde{U}'_{\mathrm{MS},\bar{\phi}}(t_{K-1}, t_{K-2}) \ldots \tilde{U}'_{\mathrm{MS},\bar{\phi}}(t_1, t_0) \tag{4.60}$$

$$= \prod_{j=1}^{K} \mathcal{D}\big[S_{\bar{\phi}}\breve{\alpha}_{K-j}\big] e^{-iS_{\bar{\phi}}^2 \Phi(t_{K-j+1}, t_{K-j})} \tag{4.61}$$

$$= \mathcal{D}\left[S_{\bar{\phi}} \sum_{j=1}^{K} \breve{\alpha}_{K-j}\right] e^{-iS_{\bar{\phi}}^2 \breve{\Phi}}, \tag{4.62}$$

where we have used the definitions

$$\breve{\alpha}_j \equiv W_k\left(\frac{t_{j+1}+t_j}{2\Delta t}\right)\alpha(t_{j+1}, t_j), \tag{4.63}$$

$$\breve{\Phi} \equiv \sum_{j=1}^{K} \Phi(t_{K-j+1}, t_{K-j}) - \sum_{l=1}^{K-1} \mathrm{Im}\left\{\left(\sum_{m=1}^{l} \breve{\alpha}_{K-m}\right)\breve{\alpha}_{K-l-1}^{*}\right\}. \tag{4.64}$$

When stepping from Eq. 4.61 to Eq. 4.62 we made use of the identity

$$\mathcal{D}[\alpha] + \mathcal{D}[\beta] = \mathcal{D}[\alpha+\beta] e^{i\mathrm{Im}\{\alpha\beta^*\}}, \tag{4.65}$$

the proof of which rests on the Baker-Campbell-Hausdorff formula. The fidelity of the WMMS gate for general δ, $\Omega^{(2)}$, and Δt is given by

$$\mathcal{F} = \frac{1}{8}\left\{3 + e^{-16\left|\sum_{j=1}^{K}\breve{\alpha}_{K-j}\right|^2\left(\bar{n}+\frac{1}{2}\right)} + 4e^{-4\left|\sum_{j=1}^{K}\breve{\alpha}_{K-j}\right|^2\left(\bar{n}+\frac{1}{2}\right)} \sin\big(4\breve{\Phi}\big)\right\}, \tag{4.66}$$

and is equivalent to Eq. 4.55, when replacing $\alpha(\Delta t, 0)$ and $\Phi(\Delta t, 0)$ with $\sum_{j=1}^{K} \breve{\alpha}_{K-j}$ and $\breve{\Phi}$, respectively.

As already mentioned above, we arrive at the same result when implementing the modulation using global π-pulses. Here, the corresponding propagator can be pieced together from basic Mølmer-Sørensen propagators for each individual loop, with rotation operators $R(\pi, \bar{\phi} + \pi/2)$, defined in Eq. 4.50, inserted at every zero-crossing of the Walsh function. The propagators of both methods are equivalent up to a global spin flip.

The robustness of the gate against timing and detuning errors can, in theory, be optimised to the desired order when following two guidelines. First, for a given number of loops K, the Walsh modulation index needs to be $k = K - 1$. Second, the higher K, the higher the order of the error suppression [HCD+12]. However, increasing the number of loops and using the correct modulation index comes at the cost of having to use longer and more complex gate sequences, which will eventually perform worse than simpler ones due to control difficulties and an increase in sensitivity to other error sources. In our experiment, we use $K = 4$ and $k = 1$, i.e. two loops, followed by a π-pulse around an orthogonal axis to $\bar{\phi}$, followed by another two loops.

We compare the performance of such a gate against its unmodulated counterpart, using simulations based on Eq. 4.66. The results are presented in Fig. 4.16 and Fig. 4.17. The former shows the gate error in the presence of a timing miscalibration. For the parameters used in our experiment, the gate infidelity is brought down to $2(2) \times 10^{-4}$, when assuming an initial thermal occupation of the motional mode of $\bar{n} = 0.1$. This corresponds to an improvement by one order of magnitude from the error calculated in section 4.3.2 for an unmodulated single-loop gate. Fig. 4.17 shows the infidelity caused by normally-distributed shot-to-shot detuning fluctuations, simulated in the same way as described in section 4.3.2. The error is suppressed to $\sim 0.2\%$, which represents a reduction by

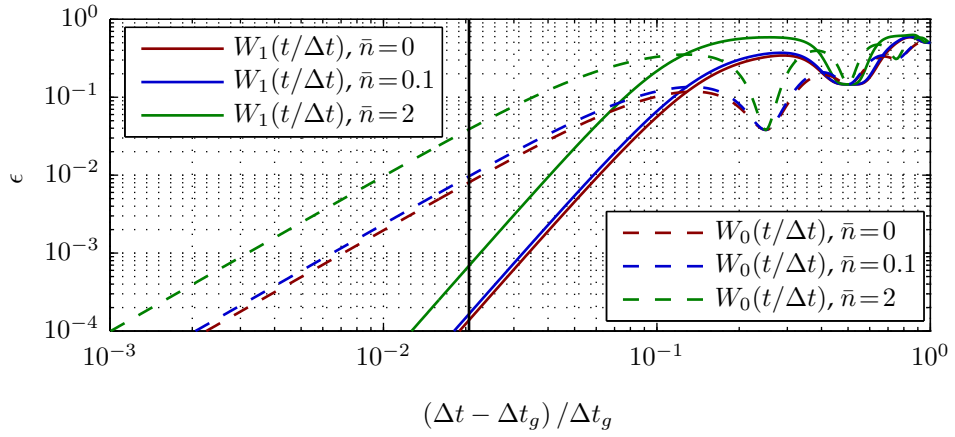

Figure 4.16: Errors due to a gate time miscalibration in a $K = 4$ loop sequence with a refocussing π-pulse after the first two loops (W_1), compared with its unmodulated counterpart (W_0). The graphs are calculated analytically from Eq. 4.55 and Eq. 4.66, for δ fixed at δ_g. Indicated by the black line is the systematic, fractional gate time miscalibration present in our experiments, amounting to $2.1(4) \times 10^{-2}$, and giving rise to a gate error of $2(2) \times 10^{-4}$ in the Walsh modulated case for $\bar{n} = 0.1$. This represents an improvement in fidelity by approximately two orders of magnitude compared with an unmodulated $K = 4$ loop gate, and about one order of magnitude compared with an unmodulated $K = 1$ loop gate.

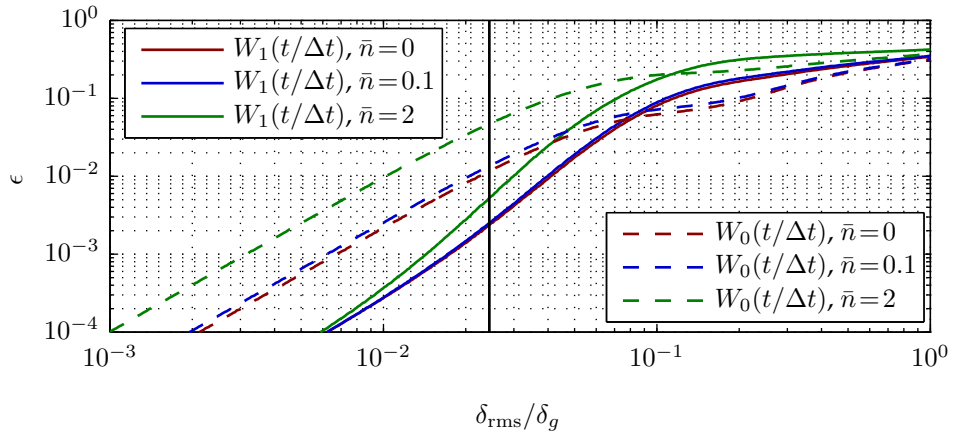

Figure 4.17: Comparison of errors due to gate detuning fluctuations in a $K = 4$ loop sequence with a refocussing π-pulse after the first two loops (W_1), against its unmodulated counterpart (W_0). The fluctuations are assumed to be normally-distributed around the ideal value δ_g, and to change δ from shot-to-shot of the experiment. For the simulation, the gate time and Rabi frequency were kept constant at $\Delta t = (\pi\sqrt{K})/\Omega^{(2)}$. The black line indicates fluctuations at $\delta_{\mathrm{rms}}/\delta_g \approx 30$ Hz r.m.s./ 1.2 kHz, in accordance with the values measured in this experiment. In the Walsh modulated case the gate error is $\sim 0.2\%$, while in the unmodulated case it is at the $\sim 1\%$ level for $K = 4$ loops, and at the $\sim 0.5\%$ level for $K = 1$ loop.

slightly better than a factor of two compared with the single-loop case without Walsh modulation.

4.4.2 Dynamical decoupling

A significant source of infidelity in the microwave driven Mølmer-Sørensen scheme are fluctuating uncompensated AC-Zeeman shifts. In our experiment, they contribute an error of order 1% if no countermeasures are employed, as previously discussed in section 4.3.5. In the following, we introduce a way to reduce this error by adding a weak carrier drive on the qubit transition to the red and blue sideband microwaves. We will refer to this modified gate scheme as the "dynamically-decoupled Mølmer-Sørensen" (DDMS) gate.

The full system Hamiltonian is

$$\tilde{H}_{\mathrm{DDMS}}(t) = \tilde{H}_{\mathrm{MS}}(t) + \tilde{H}_{\mathrm{Z}} + \tilde{H}_{\mathrm{C}}, \tag{4.67}$$

where $\tilde{H}_{\mathrm{MS}}(t)$ is the basic Mølmer-Sørensen Hamiltonian from Eq. 4.28, and where $\tilde{H}_{\mathrm{Z}}$ describes the undesired AC-Zeeman shift, given by Eq. 4.56. $\tilde{H}_{\mathrm{C}}$ is the carrier drive term of the form

$$\tilde{H}_{\mathrm{C}} = \frac{1}{2}\hbar\Omega_{\mathrm{C}} \sum_{n=1}^{N} \sigma_{+,n} e^{i\phi_{\mathrm{C}}} + \mathrm{h.c.}, \tag{4.68}$$

where Ω_{C} is the carrier Rabi frequency, and where ϕ_{C} is the phase of the carrier drive field. The carrier drive serves to dynamically decouple $\tilde{H}_{\mathrm{MS}}(t)$ from $\tilde{H}_{\mathrm{Z}}$. To understand this, we switch from the frame that rotates with $H_0 + H_{\mathrm{motion}}$, to one that co-rotates with $H_0 + H_{\mathrm{motion}} + H_{\mathrm{C}}$, so that

$$\hat{H}_{\mathrm{DDMS}}(t) = e^{(it\tilde{H}_{\mathrm{C}})/\hbar}\, \tilde{H}_{\mathrm{DDMS}}(t)\, e^{-(it\tilde{H}_{\mathrm{C}})/\hbar} = \hat{H}_{\mathrm{MS}}(t) + \hat{H}_{\mathrm{Z}}(t) + \tilde{H}_{\mathrm{C}}. \tag{4.69}$$

The circumflex in the above equation refers to the new interaction picture. The transformed Mølmer-Sørensen Hamiltonian reads

$$\hat{H}_{\mathrm{MS}}(t) = \frac{1}{2}\hbar\Omega^{(N)}\Lambda\Big(a_{l,j}e^{i\delta t}e^{-i\frac{1}{2}\Delta\phi} + a_{l,j}^{\dagger}e^{-i\delta t}e^{i\frac{1}{2}\Delta\phi}\Big), \tag{4.70}$$

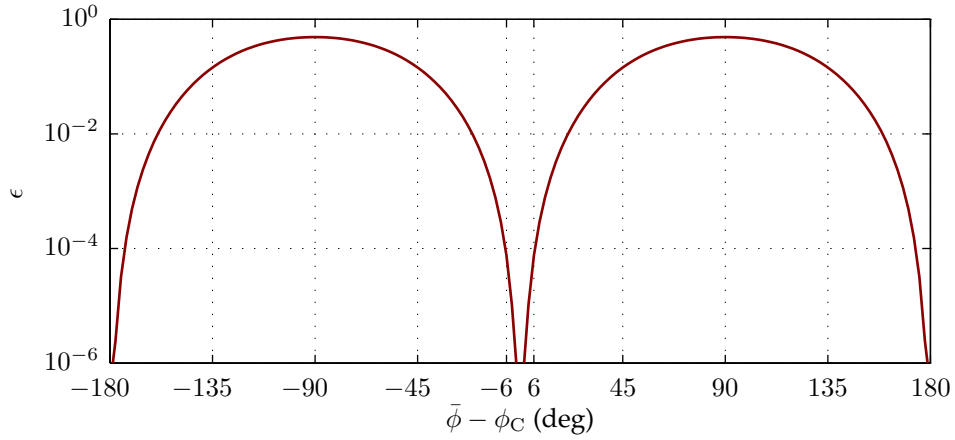

Figure 4.18: Simulated error of an ideal DDMS gate, as a function of the phase mismatch between $\bar{\phi}$ and ϕ_C. We assume similar parameters as have been used in our experiments, that is, Ω_C giving 12 complete carrier oscillations over a gate time of 3.25 ms. Introducing an uncompensated AC-Zeeman shift of $\Delta_Z = 2\pi \times 19.7(8)$ Hz, as present in the experiments, causes no significant alteration to the above graph.

where

$$\begin{aligned}\Lambda \equiv \sum_{n=1}^{N} \mathrm{sgn}(b_n^{(l,j)}) \Big\{ & \sigma_{x,n} \Big[\cos^2(\Omega_C t/2)\cos(\bar{\phi}) + \sin^2(\Omega_C t/2)\cos(2\phi_C - \bar{\phi}) \Big] \\ & -\sigma_{y,n} \Big[\cos^2(\Omega_C t/2)\sin(\bar{\phi}) + \sin^2(\Omega_C t/2)\sin(2\phi_C - \bar{\phi}) \Big] \\ & +\sigma_{z,n}\, \sin(\Omega_C t)\sin(\phi_C - \bar{\phi}) \Big\}. \end{aligned} \tag{4.71}$$

Once we manually match the spin phase $\bar{\phi}$ to the carrier phase ϕ_C, or to $\phi_C \pm \pi$, by shifting both ϕ_R and ϕ_B together, we find $\Lambda = S_{\bar{\phi}}$ and thus $\hat{H}_{MS}(t) = \tilde{H}_{MS}(t)$. We note that this phase matching needs to only be accurate to within $\pm 6^\circ$ to cause gate errors of less than 10^{-4}, as can be seen from Fig. 4.18. As a result, the carrier causes no disruption to the gate dynamics. It does however have an impact on the AC-Zeeman shift, which in the new interaction picture is given by

$$\hat{H}_Z(t) = \frac{\hbar}{2}\Delta_Z \sum_{n=1}^{N} \Big\{ \cos(\Omega_C t)\sigma_{z,n} + \sin(\Omega_C t)\Big(\cos(\phi_C)\sigma_{y,n} + \sin(\phi_C)\sigma_{x,n} \Big) \Big\}. \tag{4.72}$$

The stronger the carrier drive Ω_C compared with Δ_Z, the faster the oscillation of the summed terms in the above equation, and hence, the weaker their impact.

With the effect of the uncompensated AC-Zeeman shift suppressed in this manner, the gate Hamiltonian turns into $\hat{H}_{DDMS}(t) \approx \tilde{H}_{MS}(t) + \tilde{H}_C$. Although

$\tilde{H}_{\mathrm{C}}$ is responsible for this error suppression, it also introduces an unwanted dependence of the final state after the gate on Ω_{C}. This stems from the fact that the weak carrier drive causes resonant Rabi oscillations between the two qubit states of both ions. Unless Ω_{C} is tuned so that an integral number m of carrier flops is conducted over the course of the gate, that is, $\Omega_{\mathrm{C}}\Delta t = 2m\pi$, these flops upset the populations at the end of the gate operation.

The AC-Zeeman shift error reduction is not impaired when additionally employing Walsh modulation, that is, when replacing $\tilde{H}_{\mathrm{MS}}(t)$ with $\tilde{H}_{\mathrm{WMMS}}(t)$ in the above derivation. Furthermore, implementing the modulation using global π-pulses has the added benefit of refocussing partially complete carrier Rabi oscillations. This is a consequence of the π-pulses being 90° out of phase with the carrier drive and appearing symmetrically spaced about the half way point of the gate sequence. Thus, they act as spin-echo pulses, eliminating the above mentioned necessity to ensure that the ions undergo an integral number m of carrier flops over the course of the gate.

The AC-Zeeman shift fluctuations in this work are characterised to have an r.m.s. value of $\Delta_{\mathrm{Z,rms}} = 2\pi \times 19.7(8)\,\mathrm{Hz}$, which in section 4.3.5 was calculated to cause a gate error of 5.6(5)% for $K = 4$ loops when utilising the basic Mølmer-Sørensen scheme. Simulating by how much this error is reduced as a function of Ω_{C}, when also including a $W_1(t/\Delta t)$ Walsh modulation, is computationally intensive. To get an initial idea, we look at the simpler case of an uncompensated, static AC-Zeeman shift of $\Delta_{\mathrm{Z}} = \Delta_{\mathrm{Z,rms}}$. The resulting error as a function of the carrier strength is plotted in Fig. 4.19 for a $K = 4$ loop gate with, on the one hand, a refocussing π-pulse in the middle of the gate, and on the other hand, no Walsh modulation at all. Although not indicated in the plot, the unmodulated case gives, to a good approximation, similar gate errors as a modulation with a π phase shift of $\bar{\phi}$.

Three features are apparent from the figure. First, the previously mentioned

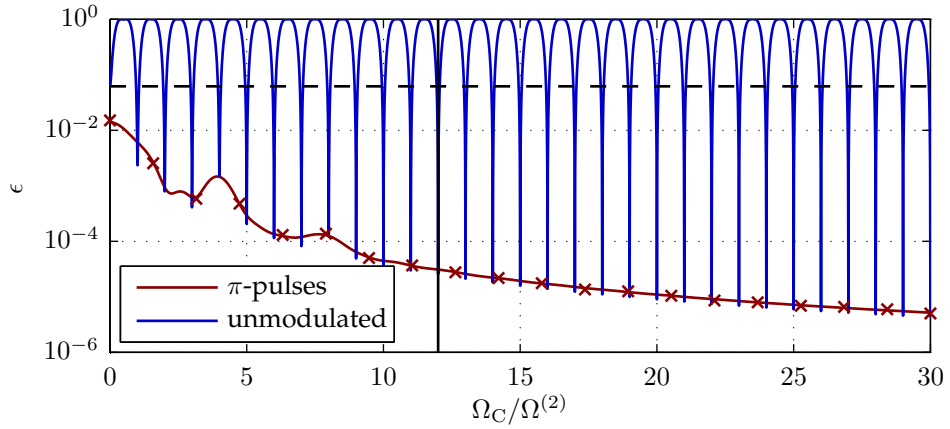

Figure 4.19: Simulated DDMS gate error as a function of carrier Rabi frequency in the presence of an uncompensated AC-Zeeman shift, with all other parameters set as for an ideal $K = 4$ loop gate at $\bar{n} = 0$. The red and blue lines represent the infidelity from a constant shift $\Delta_Z = 2\pi \times 19.7\,\mathrm{Hz}$, with and without the refocussing π-pulse, respectively. Indicated by the crosses is the gate error when utilising the refocussing π-pulse, but assuming normally distributed, shot-to-shot fluctuations of the uncompensated AC-Zeeman shift, with an r.m.s. value of $\Delta_{Z,\mathrm{rms}} = \Delta_Z$. The vertical black line corresponds to the parameters used in this work, that is, $\Omega_C/\Omega^{(2)} = 3.69\,\mathrm{kHz}/308\,\mathrm{Hz}$, producing 12 complete carrier flops. The horizontal, dashed, black line constitutes the gate error for the unmodulated, basic Mølmer-Sørensen gate.

advantage of implementing the Walsh modulation using a π-pulse. Its capability to refocus partially complete carrier oscillations, allows for low error levels to be achieved even if $\Omega_C \Delta t \neq 2m\pi$, for an integer m. Without it, or when using a π phase shift of $\bar{\phi}$ instead, the gate is more difficult to set up and loses its robustness against fluctuations in Ω_C. Second, the fast drop in gate error with the strength of the carrier. In this work, we apply 12 complete carrier flops, which are sufficient for an error suppression by more than three orders of magnitude. Third, at $\Omega_C = 0$ the DDMS curve is still lower than that of the basic Mølmer-Sørensen scheme. This is another manifestation of our earlier observation, made in section 4.4, that π-pulses provide some protection against AC-Zeeman shift errors, even without the carrier drive.

To get specific error values, we have to take into account the fluctuating nature of the uncompensated AC-Zeeman shift. We therefore perform a similar

simulation to that described in section 4.3.5, assuming normally distributed shot-to-shot variations with $\Delta_{\mathrm{Z,rms}} = \Delta_{\mathrm{Z}}$. The results for twenty different carrier strengths are indicated as crosses in Fig. 4.19 and show good agreement with the static case.

For the parameters used in our experiment, the AC-Zeeman shift error is suppressed to $3.0(2)\times10^{-5}$. When leaving out the refocussing π-pulse, the error is $3.7(3)\times10^{-5}$, confirming that the Walsh modulation has, to a good approximation, no effect on the gate error from uncompensated AC-Zeeman shifts. The converse is also true, meaning that the carrier oscillations have no significant impact on our $\sim 0.2\%$ estimate for the gate error due to detuning fluctuations. In fact, the DDMS scheme leaves all errors sources other than uncompensated AC-Zeeman shifts unaffected, with the exception of a small improvement of the (already negligible) error from imbalanced Rabi frequencies (section 4.3.3).

As a final remark, the DDMS gate is closely related to the "driven geometric phase gate" presented in [BSPR12], [LBB13], and experimentally realised in [TGB$^+$13]. Instead of two sideband tones and a weak carrier, that scheme involves one sideband tone and a strong carrier. Its advantage is most obvious when using lasers, as unlike the standard Mølmer-Sørensen gate, it eliminates the requirement for phase-stable, non-copropagating beams. However, this advantage is less relevant in an all-microwave approach, since in this case phase control is much easier to achieve.

The principal advantage of the DDMS scheme is that the carrier drive is used as a supplementary feature to reduce noise, rather than being an integral part of the gate mechanism. It is therefore possible to use weaker values of Ω_{C}, thus making the gate less prone to off-resonant carrier excitation or noise on the carrier drive. Furthermore, to achieve a given gate speed, the DDMS scheme requires only half the microwave power compared with the driven geometric phase gate, which is beneficial in terms of power dissipation in the trap chip.

5

Apparatus

In this chapter we provide an overview of the experimental setup and equipment used for the work presented in this thesis. Our experiments build on efforts presented in [All11, Har14], where much of the apparatus is described. To keep repetition to a minimum, we will mainly focus on components that have since been added or modified, and only briefly review some of the previously documented parts.

5.1 The ion trap

Our experiments were conducted using a room-temperature, microfabricated, planar ion trap of the "five wire" type [CBB$^+$05], the design and characterisation of which was done by D. T. C. Allcock and can be found in [All11, AHB$^+$13].

5.1.1 Design

Fig. 5.1 shows a to-scale schematic of the central portion of the trap, consisting of three parallel centre electrodes, and two outer DC control electrodes, each segmented into three parts. In order to create a Paul trap in the radial, i.e. $\{\hat{\mathbf{y}}, \hat{\mathbf{z}}\}$, plane, we put a 72 V radiofrequency signal at ~ 38.3 MHz on the outer two of the three axial electrodes, while holding the remaining sections of the trap at RF-ground. Axial confinement is provided through a static potential, generated by

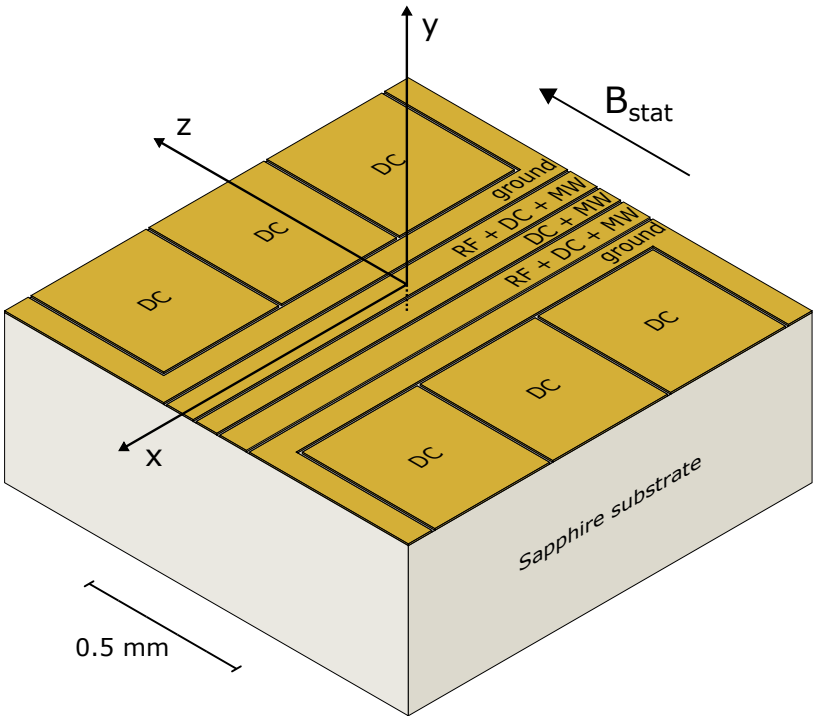

Figure 5.1: To-scale schematic of the central portion of the planar ion trap used in this work. The trap consists of $\sim 5\,\mu$m thick gold electrodes, electroplated on a 0.5 mm thick sapphire substrate. The three axial electrodes measure 91 µm - 70 µm - 91 µm with gaps of 8-10 µm between them, the uncertainty arising from fabrication tolerances. Both outer DC electrodes are segmented into three parts and set back from the centre electrodes by a 75 µm wide section of ground plane, providing a passageway for microwave return currents. The ions are held 75 µm above the surface, at the origin of the indicated coordinate system. We note that the axes of the radial secular modes lie at $\sim 40^\circ$ to $\{\hat{\mathbf{y}}, \hat{\mathbf{z}}\}$, as illustrated in Fig. 4.1.

putting appropriate DC voltages ($\sim 10\,$V) on the control electrodes. The resulting trap confines ions 75 µm above the surface of the middle electrode, and has a secular frequency of $\sim 3\,$MHz in the radial directions, and $\sim 500\,$kHz in the axial direction.

Since the electrode structure is axially symmetric, the principal axes of the trapping potential are aligned with $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$. This is problematic from the point of view of laser cooling, as in order for the 397 nm beams to not scatter off the trap's surface, they should lie parallel to the $\{\hat{\mathbf{x}}, \hat{\mathbf{y}}\}$ plane, which prohibits cooling of the ions' motion in the $\hat{\mathbf{y}}$ direction. Scattered 397 nm photons from the trap's surface generally give rise to a significant background signal on the PMT, making

it hard to distinguish the ions' fluorescence during readout, and can also cause significant charging of the trap substrate. We therefore follow the approach outlined in [ASS+10] and use the DC control voltages to create an additional, static quadrupole potential in the radial plane. The purpose of this potential is to tilt the radial principal trap axes away from $\{\hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ by $\sim 40°$. This allows for the Doppler beam to couple to both radial motional modes, while also breaking their degeneracy in frequency, separating them by $\sim 330\,\mathrm{kHz}$.

Apart from supporting the trapping RF and a DC bias, each of the three axial electrodes is also capable of simultaneously carrying microwave currents, which generate oscillating magnetic fields above the surface of the trap for driving quantum logic gates. Choosing to combine multiple signals on the same electrode, instead of having separate ones for RF and the gate microwaves, is technically more demanding and has its reason in the way the trap was originally intended to be operated.

The idea was to apply microwave signals of appropriate relative amplitudes and phases to the three axial electrodes at the same time, so that their respective magnetic fields cancel at the position of the ions, while retaining a large spatial gradient. In this way, the errors from off-resonant excitation can be greatly reduced when performing two-qubit gates [OLA+08]. Since the achievable sideband Rabi frequency, i.e. the gate speed, is usually what limits the fidelity of microwave driven gates (see section 3.2), the electrode layout was optimised for yielding the largest possible microwave field gradient at the position of the ions. As a result, combining the gate drive with the trapping RF turned out to be more advantageous than placing separate microwave electrodes around the RF electrodes [All11].

The three axial electrodes are microwave half-wave cavities that resonate near the clock qubit frequency of $^{43}Ca^{+}$ at 146 G and reduce the microwave power required to achieve a given current. Their Q-factors are ~ 5 and they are

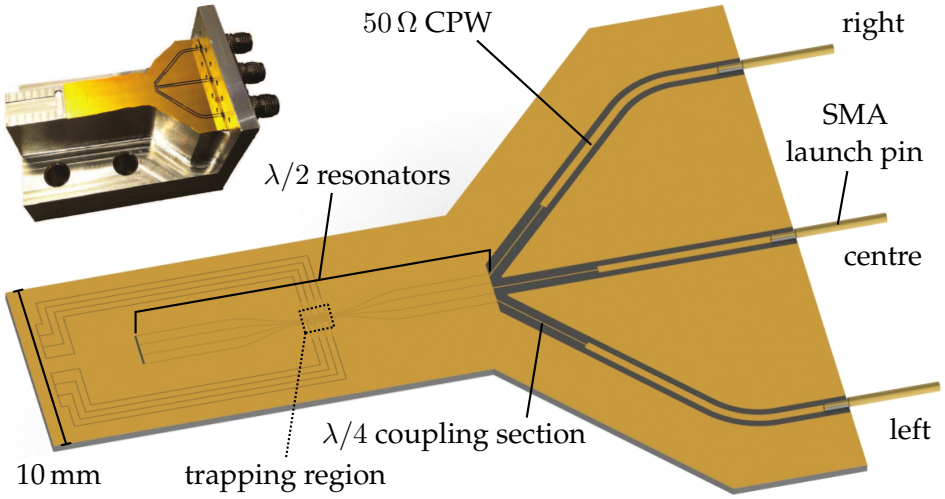

Figure 5.2: Model of the full ion trap's electrode structure, indicating integrated microwave elements, as well as the trapping region that was previously shown in Fig. 5.1. The top left inset is a photograph of the ion trap on its aluminium mount, which holds the three in-vacuum SMA connectors. Their launch pins reach through the mounting structure onto the trap surface, where they are attached to the co-planar waveguides (CPWs) with conductive epoxy. The DC control electrodes are routed towards the opposite end of the trap, where they are wire-bonded to copper wires (not shown), which are fixed inside a recess in the trap mount. To ensure RF grounding, each copper wire is additionally wire-bonded to an 820 pF capacitor, glued to the mount with conductive epoxy.

50 Ω impedance-matched to the microwave drive circuitry when on resonance. As illustrated in Fig. 5.2, microwave signals are fed through in-vacuum SMA connectors, one for each axial electrode, and guided to the trapping region via 50 Ω coplanar waveguides (CPW). There, the signals are matched into the resonators through quarter-wave coupling elements, i.e. a $\lambda/4$ section of CPW of higher impedance than the feed line.

5.1.2 Magnetic field characterisation

Although the trap is designed to carry microwave currents on all three axial electrodes, we did not attempt to null the field in this way. Instead, we apply the bichromatic gate drive to the centre electrode only, and rely on the DDMS scheme, explained in section 4.4.2, for protection against off-resonant excitation errors from a non-zero field at the position of the ion. To confirm that this electrode

electrode	pol. ang.	$B_{+1}^{(0)}$ (μT)	$B_0^{(0)}$ (μT)	$\partial B_{+1}/\partial q_z$ (T/m)	$\partial B_{+1}/\partial q_y$ (T/m)	$\partial B_0/\partial q_z$ (T/m)	$\partial B_0/\partial q_y$ (T/m)
centre	3.1°	0.681	17.92	0.128	0.136	0.25	0.23
right	62.9°	6.954		0.046	0.087		
left	58.4°						

Table 5.1: Properties of the magnetic fields produced by each of the three axial trap electrodes at the position of the ion, where 'left' and 'right' are assumed as seen from the microwave input end (see Fig. 5.2). The second column lists the polarization angles (pol. ang.), i.e. the angles that the magnetic field's polarization vectors span with the static magnetic field direction ($\hat{\mathbf{z}}$). The third and fourth column give the measured field components, while the remaining columns give the field gradients, all normalised to an input power of 1 mW into the diplexer. We performed measurements on the centre and right electrode, assuming that, by symmetry, the field values produced by the left electrode are similar to that of the right electrode. Blank entries in the second row of the table correspond to quantities that have not been sufficiently well determined to be reliable. Details about the measurements can be found in section 6.2 of [Har14].

is best suited for our gate, we characterised the magnetic fields generated by each axial electrode. Table 5.1 summarises the results, while details about the underlying measurements are presented in section 6.2 of [Har14].

The magnetic field from the centre electrode is predominantly π-polarized, which is beneficial from the standpoint of off-resonant excitation, as the closest π-polarized transitions are $\sim 50\,\mathrm{MHz}$ detuned from the radiation driving the gate, as opposed to the only $\sim 3\,\mathrm{MHz}$ detuned σ_+-polarized carrier transition. At the same time, the field has a strong σ-polarized gradient along $\hat{\mathbf{q}}_\mathbf{y}$ and $\hat{\mathbf{q}}_\mathbf{z}$, which is crucial for spin-motion coupling. In terms of resonant pulses, we use the left electrode for signals driving single-qubit gates and state preparation pulses, and the right electrode for the carrier drive required for the DDMS scheme, as these electrodes have a large σ-polarized field component.

5.1.3 Micromotion compensation

We detect the ion's micromotion using the method described in [ASS+10] and compensate for it by changing the DC control voltages accordingly. In the $\{\hat{\mathbf{x}}, \hat{\mathbf{z}}\}$-plane, we set the compensation field with an accuracy of $\sim 1\,\mathrm{V/m}$ using a

single $^{43}Ca^{+}$ ion. In the $\hat{\mathbf{y}}$-direction we switch to $^{40}Ca^{+}$ to get a detectable error-signal, but we are unable to reduce the micromotion below $\sim 30\,\mathrm{V/m}$. The most likely reason is RF pickup on the trap's axial electrode, which, because it needs to support DC and microwaves, can only be properly grounded at RF frequencies before the multiplexing circuitry, located outside of the vacuum chamber. However, the uncompensated micromotion is small enough to not cause any problems in our experiments.

5.2 Basic trapping equipment

In this section we introduce all basic pieces of apparatus needed for trapping $^{43}Ca^{+}$ at 146 G, with the exception of the laser setup and the microwave sources, which are described subsequently. Most of the equipment described below was assembled by D. T. C. Allcock and T. P. Harty.

5.2.1 Vacuum system

The ion trap on its aluminium base is mounted in the centre of a stainless steel "spherical octagon" vacuum chamber[1], with two 6 inch CF flanges at the top and the bottom of the chamber, and eight equally spaced 2.75 inch CF flanges around the sides. It is pumped by an ion pump[2] and a non-evaporable getter[3], sustaining a base pressure in the $10^{-12} - 10^{-11}$ Torr range, which is monitored by an ionization gauge[4].

As sketched in Fig. 5.3, the two CF flanges along the trap axis are used for attaching the vacuum pumps on one side, and SMA feedthroughs for the trapping RF and microwave signals on the other. The remaining six flanges on the side

[1] Kimball Physics, MCF600-SphOct-F2C8, 316L stainless steel.
[2] Varian (now Agilent Technologies), Diode VacIon Plus 20, 27 l/ s.
[3] SAES Getters, GP100 MK5, ST707 alloy, 240 l/ s H_2 pumping speed.
[4] Varian (now Agilent Technologies), UHV-24P, x-ray limit at 5×10^{-12} Torr, reduced performance at pressures lower than 5×10^{-11} Torr

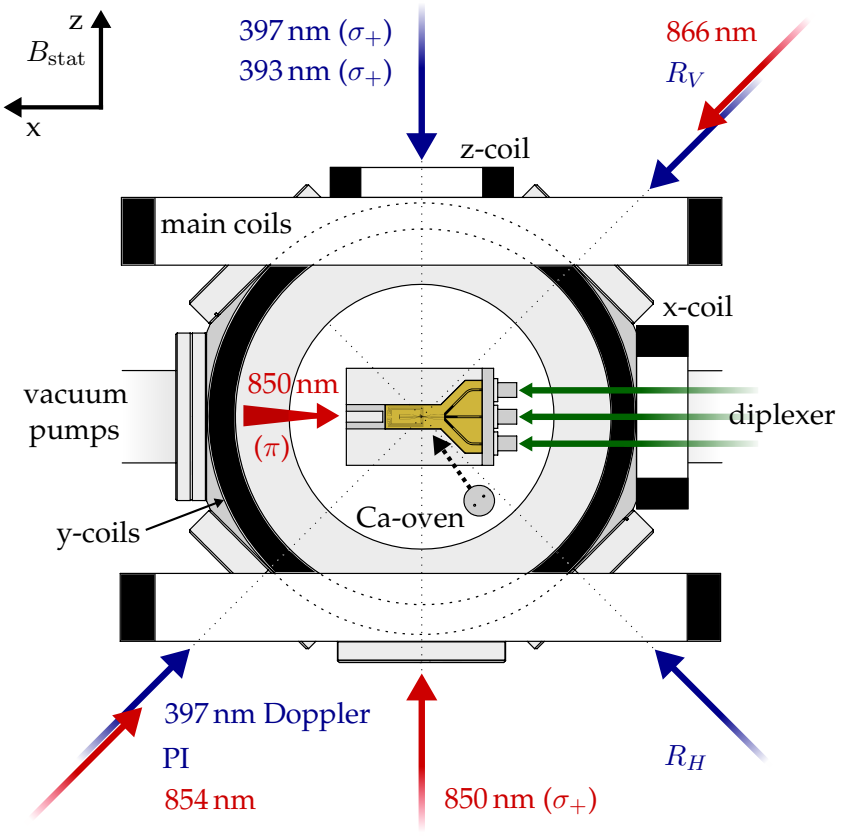

Figure 5.3: Sketch of the vacuum chamber, as seen from above, with the trap chip on its aluminium mount positioned in the centre of the octagon. Laser beam orientations, the location of magnetic field coils, vacuum pumps, and the diplexer are indicated. The three SMA connections from the diplexer to the trap mount, carrying the trapping RF, the microwave signals needed for the gate, and DC bias voltages are shown as green arrows. The 850 nm (π) beam enters the vacuum can through the top viewport, at $\sim 45°$ to the plane of the diagram, and hits the ion before being reflected from the trap surface and exiting through the top viewport. We use the same path for an 866 nm beam, when compensating micromotion along the $\hat{y}$-axis [ASS+10]. The beams labelled R_H and R_V are from the Raman laser and are used for sideband cooling. R_H is horizontally polarized with respect to the plane of the drawing, while R_V is vertically polarized. The neutral beam of calcium atoms from the oven is shown as the dotted, black arrow.

of the octagon are equipped with viewports[5] for laser access, with an additional viewport[6] attached to the top flange for collecting ion fluorescence as well as providing an entry point for an infrared laser beam in the $\{\hat{\mathbf{y}}, \hat{\mathbf{x}}\}$ plane. The bottom flange of the octagon is custom made and incorporates tapped mounting holes for the trap base and electrical feedthroughs for DC connections and the calcium oven.

The oven is a resistively heated stainless steel tube, constructed following the group's standard design [All11], and containing a granulated, isotopically-enriched calcium mixture with 12% $^{43}Ca^{+}$ and 88% $^{40}Ca^{+}$. For loading, we apply a 5 A current that heats the initially cold oven up within ~ 5 minutes. Once the oven is warm, a $^{40}Ca^{+}$ ion can be loaded approximately twice per minute while a $^{43}Ca^{+}$ ion can be loaded about once every two minutes. In order to reduce Doppler shifts on the 423 nm transition, the oven is placed so that the beam of neutral calcium is as close to perpendicular to the photo-ionization beams as possible ($\sim 83^{\circ}$) without obstructing the R_H Raman beam path.

5.2.2 Diplexer

In order to multiplex microwave signals and DC control voltages on each of the three axial electrodes, as well as to simultaneously put trapping RF voltages on the outer two, we use a diplexer. This design and testing of this filter network is described in [All11]. At its RF input, there is an impedance matching, toroidal resonant transformer[7], which has a resonance at 38.7 MHz, a Q of 65, and which provides a voltage step-up of 9.6. The diplexer is located outside of the vacuum chamber and its three outputs are connected to their respective trap electrodes

[5] Torr Scientific, VPZ38QBBAR-LN, fused silica viewport, 3 mm thick, broad-band anti-reflection coating on both sides, specified reflectance per surface is 0.5 % at 397 nm and 9.5 % at 866 nm.

[6] Torr Scientific, VPZ100-ITO/AR-LN, Kodial glass viewport, 4 mm thick, anti-reflection coating at 397 nm on air-side, thin layer of conductive indium tin oxide on vacuum side to prevent charging. Net transmission at 397 nm is 76 % at normal incidence.

[7] Micrometals T80-10

using short SMA cabling via SMA feedthroughs in one of the octagon's CF flanges. We use both passive insulation and active control to stabilise to protect the diplexer from temperature fluctuations.

5.2.3 DC source

The DC voltages for the trap are generated by a commercial digital to analogue converter (DAC) evaluation board[8], operated from the experimental control computer via USB. To establish a connection to the six control electrodes, we use screened cable to wire the DAC outputs to a low-pass filter board, attached directly to the D-type feedthrough in the vacuum can's base flange. For putting DC voltages on the three axial trap electrodes, we connect the DAC outputs to the designated diplexer inputs, using semi-rigid SMA cables.

5.2.4 Trapping RF

In order to create radial confinement, we apply an RF voltage to the two outer axial electrodes of the trap chip. The resulting radial secular frequency is proportional to the RF voltage amplitude and inversely proportional to its frequency [WMW+98]. Consequently, instabilities in both these parameters manifest themselves as motional mode frequency fluctuations, which in turn give rise to detuning errors when conducting two-qubit gates, as discussed in section 4.3.2.

We are more concerned about RF amplitude noise rather than frequency noise, as the latter has a weaker contribution to the resulting motional mode variation [Har14]. To reduce the former, we put our 38.7 MHz starting signal, generated from a HP 8656B frequency synthesiser, through a heavily saturated amplifier circuit, designed and built by T. P. Harty, and documented in [Har14]. The aim of the circuit is to stabilise the output amplitude of the synthesiser by referencing it

[8] Analog Devices inc., AD5372 DAC on an EVAL-ADUM4160 evaluation board.

coil	shape	N	dimensions (mm)	l (mm)	G/A
main	rectangular	2×16	137×266 ($\hat{\mathbf{y}} \times \hat{\mathbf{x}}$, inside)	52	1.2
x-coil	cylindrical	168	80/136 (diameter in/out)	133	0.9
y-coil	cylindrical	2×30	156.5/176 (diameter in/out)	38.3	3
z-coil	cylindrical	92	67/87 (diameter in/out)	110	0.46

Table 5.2: List of magnetic field coils used in our experiment, where N is the total number of turns, l is the distance from the coil centre to the ion, and the final column in the table gives the expected field (Gauss) produced at the ion per amp of current. For the main coils, this quantity was estimated from the coil dimensions and the Biot-Savart law, but experimentally we find that we require 135 A to produce the 146 G magnetic field. The location of all coils can be seen in Fig. 5.3.

to a high stability DC voltage source, which sets the amplitude of a square-wave signal. This is filtered to select the fundamental frequency component.

The output of the circuit is amplified with a low-phase-noise amplifier[9] and then stepped up by the toroidal resonant transformer, described in section 5.2.2. All components in the above chain are actively temperature stabilised to within ~ 1 mK, and screened from air currents. Radial motional mode stability measurements using a single ion will be presented in section 8.2.

5.2.5 Magnetic field coils

We make use of four sets of orthogonal magnetic fields coils around the vacuum chamber, in order to generate the static field of 146 G along $\hat{\mathbf{z}}$. The orientation of the coils is shown in Fig. 5.3, with their associated characteristics summarised in table 5.2.

The bulk of the field is produced by a pair of rectangular, liquid-cooled copper coils (main coils), operated at ~ 135 A and ~ 10 V using a pair of high-current power supplies[10] in series. Since these power supplies have a limited resolution of 125 mA, we employ a trim coil (z-coil), placed around the viewport through which the σ_+ polarized 397 nm and 393 nm beams enter the vacuum chamber, to fine-tune the field strength. Additional trim fields along $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are used

[9] Spectrum Microwave inc., TM3110, 15.5 dB gain
[10] Keysight Technologies (Agilent), 6671A, 2 kW system power supply

to align the main field onto the σ_+ polarised optical pumping beams, ensuring good polarisation purity for high-fidelity state preparation and readout. These fields are generated by a pair of coils bolted onto the top and bottom of the octagon (y-coils), as well as one coil fitted over the microwave feedthrough (x-coil). All trim coils are powered by commercial power supplies[11], connected to the experimental control computer via USB.

The main coil setup has not changed from the version described in [Har14], with the exception of the cooling system. Instead of drawing water from the building's mains water supply, we now employ a closed-cycle system cooled by a recirculating chiller[12]. The reason for the upgrade was gradual buildup of biological fouling within the hollow copper wires that make up the coils, leading to a reduction of the water flow rate and hence, an increase in coil temperature up to 40 °C. In the new system, we use a heat transfer fluid[13] that prevents biofouling and that is actively temperature stabilised to within 0.1 °C by the chiller unit, allowing us to operate the coils at a steady temperature of 34 °C.

We stabilise the current going into the main coils using a feedback system, designed and tested by J.E. Tarlton, and described in [Tar]. Its principle of operation is to shunt away the appropriate amount of current, based on the measurement of a fluxgate current sensor, placed at the coil input. We characterise the performance of the feedback system by means of coherence time measurements on the magnetic field sensitive stretch transition ($|M_m = +4\rangle \leftrightarrow |M_n = +3\rangle$) of a single $^{43}Ca^+$ ion. To this end, we use a Ramsey experiment with delay times of up to 600 µs, yielding a coherence time of 365(5) µs, which translates to magnetic field fluctuations of 261(3) µG r.m.s. around the mean. We note that our feedback system only reduces noise in the coil power supplies and does not help with ambient magnetic field variations. These variations occur predominantly at 50 Hz

[11] Aim-TTi - Thurlby Thandar Instruments Limited, QL series power supplies.
[12] Applied Thermal Control Ltd, K3 chiller
[13] Applied Thermal Control Ltd, Hexid heat transfer fluid

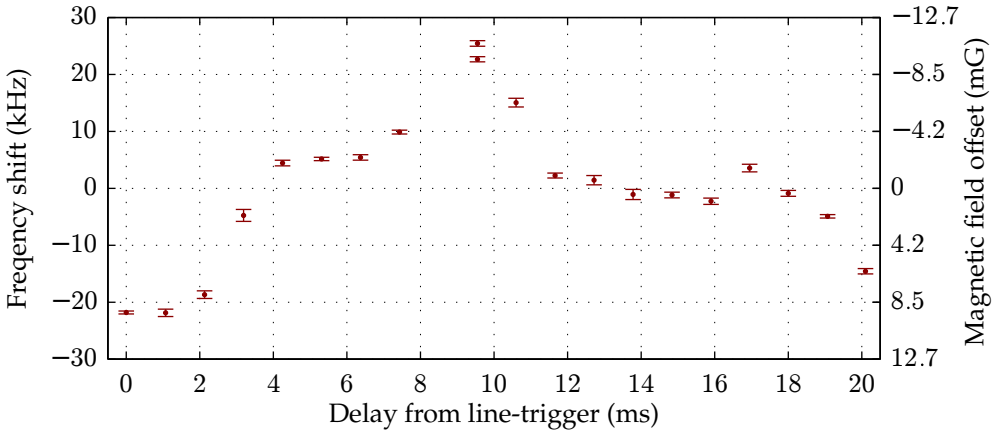

Figure 5.4: $|M_m = 4\rangle \leftrightarrow |M_n = 3\rangle$ transition frequency measured by means of Rabi spectroscopy, after a successively increasing delay from the line-trigger.

and at near DC, and are characterised in the following section.

5.2.6 Ambient magnetic field fluctuations and servo routine

Since we perform our gate experiments on a clock transition, we benefit from excellent protection[14] against uncontrolled frequency fluctuations due to magnetic field noise. However, during qubit initialisation and readout, we have to employ the microwave transfer pulses $\{A, B, C\}$ (introduced in section 2.5), which act on transitions that are vulnerable to magnetic field variations. Consequently, such variations impact the fidelity of the transfer pulses, and thus increase the state preparation and measurement (SPAM) error in our experiment.

As discussed in the previous section, magnetic field instabilities are dominated by noise from the main coil current supplies, which is suppressed with the help of a feedback system. What remains are ambient field fluctuations from the lab environment. These are generally hard to control and, in our labs, occur in two main forms.

First, repeatable field oscillations that are synchronous with the 50 Hz mains electricity cycle, stemming from running electrical equipment, e.g. power sup-

[14]The qubit frequency is first-order insensitive to changes in the magnetic field, while the second-order field dependence is $2.416\,\mathrm{mHz/mG^2}$.

plies. We characterise these oscillations at the ion by looking at the associated frequency change on a magnetic field sensitive transition as a function of time within a 20 ms mains cycle. To this end, we use microwave Rabi spectroscopy on the A-transition ($|M_m = 4\rangle \leftrightarrow |M_n = 3\rangle$), since A has a high linear magnetic field dependence ($-2.362\,\mathrm{MHz/G}$). The results of the measurement is plotted in Fig. 5.4.

To minimise the impact of 50 Hz noise, we run our experiments synchronised with the mains cycle, that is, we use a fixed point of the power line as trigger for our sequences (line triggering). This means that as long as the frequencies of all pulses addressing transitions other than the clock transition are independently calibrated for a given point in the 50 Hz cycle, there should be no increase in SPAM error due to field oscillations with mains frequency. In future, we plan on removing this additional calibration complexity by counteracting the 50 Hz noise by means of a feedforward mechanism on the main coil current.

The second form of ambient field fluctuations are slow magnetic field drifts caused by various sources within the building, including a nearby goods lift and a number of scanned superconducting magnets. Since our experiment is not equipped with any magnetic shielding, we cope with these drifts by employing an automated servo routine. This servo routine preforms Rabi spectroscopy on the A-transition to check for any magnetic field offsets from 146.0942 G, followed by an adjustment of the current in a trim coil (z-coil) to cancel it out. More specifically, we use a 22.2 µs long π-pulse to probe the 36 kHz FWHM wide resonance on both sides, that is, at ± 18 kHz from the transition frequency expected at 146.0942 G. In this way, any shifts of the line can therefore be determined from the imbalance in the measured populations. The servo routine keeps the static field within ≈ 1 mG of the desired value.

5.2.7 Imaging system

We collect and focus the ions' fluorescence onto a camera or photomultiplier tube[15] (PMT) with the help of an optical system, mounted vertically above the top viewport. The camera[16] is used for diagnostic purposes, such as beam alignment, while the qubits' state detection is performed with the PMT. We select between the two using a moveable beam-splitter. The objective is a custom 5-element lens with a numerical aperture of approximately 0.3 and a specified transmission of 95% at 397 nm. To minimise background counts, mainly due to photons scattered from the trap surface and room lights, we use an adjustable aperture, placed in the focal plane in front of the PMT, as well as a selection of optical filters. We open the aperture to facilitate beam alignment, allowing us to view an area of approximately 0.5 mm by 0.5 mm of the centre portion of the trap's surface on the camera. In this configuration, we can also determine the magnification of the imaging system, as the dimensions of the electrode structure are known. The magnification is 15.7, meaning that each 8 μm pixel on the camera corresponds to ≈ 0.5 μm in the focal plane.

We measure a typical absolute collection efficiency of the imaging system and PMT of $\approx 0.23\%$, that is, the probability of registering a count in the PMT per photon emitted by the ion. The exact value depends on the setting of the aperture and the filters used. A schematic of the imaging system and further details can be found in [Har14].

5.3 Laser systems

In order to cool, initialise and read out $^{43}Ca^{+}$, following the procedures described in sections 2.3-2.6, we drive five electric dipole transitions between 4S, 4P and 3D. Two further wavelengths are required for photo-ionization of neutral calcium.

[15] ET Enterprises Limited, P25PC-12, serial number 9584. At 400 nm, the specified quantum efficiency is $\approx 27\%$ and the typical dark count rate is $100\,s^{-1}$ at 20°.

[16] Andor Technology Ltd, Luca DL-406M-OEM electron multiplying CCD

The associated wavelengths fall into two ranges, the "blue", consisting of 389 nm, 393 nm, 397 nm and 423 nm, and the "red", consisting of 850 nm, 854 nm and 866 nm. These wavelengths are convenient for three reasons. First, they can be provided by commercially available diode lasers, second, these wavelengths are compatible with high quality optical fibres, and third, the two narrow wavelength ranges allow for standardisation of most optics.

Our lasers are located on separate optics tables to the main experiment, with the exception of the Raman laser system, which is assembled next to the ion trap apparatus for reasons of space availability. Each beam is transferred to the experiment using single-mode polarization-maintaining fibres, which turn beam pointing fluctuations in the laser source into straightforwardly correctable power fluctuations at the fibre output.

5.3.1 Photo-ionization lasers

For the implementation of the photoionization protocol, explained in section 2.3, we use a 423 nm Toptica DL100 external cavity diode laser (ECDL) in a Littrow configuration, and a 389 nm free-running diode in a home-made mount. We manually tune the 423 nm laser frequency to select the desired isotope of calcium to be loaded into the trap (the isotope shift between ^{40}Ca and ^{43}Ca is 612 MHz). The two beams are combined on a dichroic mirror and coupled into a single-mode fibre. They can be switched on and off with the help of a mechanical shutter, placed before the fibre.

5.3.2 Ca^+ lasers

The five lasers lasers involved in Doppler cooling, state initialisation, and readout are Toptica DL100 ECDLs, frequency stabilised to piezo-tuned reference etalons[17]. Stabilisation is either achieved by means of side-of-fringe locks or Pound-Drever-

[17]NPL "Low Drift Etalons", the specified drift rate is less than 0.5 MHz/hour.

Hall (PDH) locks [Bla01]. The former method has the disadvantage of being sensitive to intensity noise and having a capture range that is limited to the fringe width of the reference cavity. It is however significantly simpler to implement, which is why we employ the side-of-fringe technique on lasers that are less demanding in terms of frequency stability, i.e. on the 850 nm and 854 nm repump lasers.

The 397 nm, 393 nm and 866 nm lasers are equipped with PDH locks, which are more robust, but require the lasers to be phase modulated. In case of the blue lasers, this modulation is achieved with EOMs[18] placed in the beam paths going to the reference etalons. In case of the 866 nm repump laser, we resort to the simpler method of modulating the diode current directly, as the resulting sidebands on the experimental beam have no significant impact on the Doppler cooling performance.

To tune the laser frequencies, we vary the reference etalons' piezo voltages via the experimental control computer. We monitor the laser frequencies and output spectra by picking off a sample of each beam and feeding it to a multi-channel diagnostic system, which includes a wavemeter and a blue and red optical spectrum analyser (see section 2.1.3 in [All11] for further details).

We protect the lasers from optical feedback due to retro-reflected light from components further down the beam path, by placing optical isolators immediately after each laser's output. To remove amplified spontaneous emission, which may lead to unwanted shelving or deshelving of the ions, the 397 nm, 866 nm and 850 nm beams are filtered by blazed diffraction gratings. An additional benefit of using these gratings is that the lower intensity diffracted orders can be used for laser locking and diagnostics.

Each beam is put through a double-pass acousto-optic modulator (AOM) prior to being coupled into an optical fibre and transferred to the experiment. The

[18] Qubig GmbH, E0-F80M3 (running at 80 MHz in the 393 nm setup), E0-F50M3 (running at 70 MHz in the 397 nm setup).

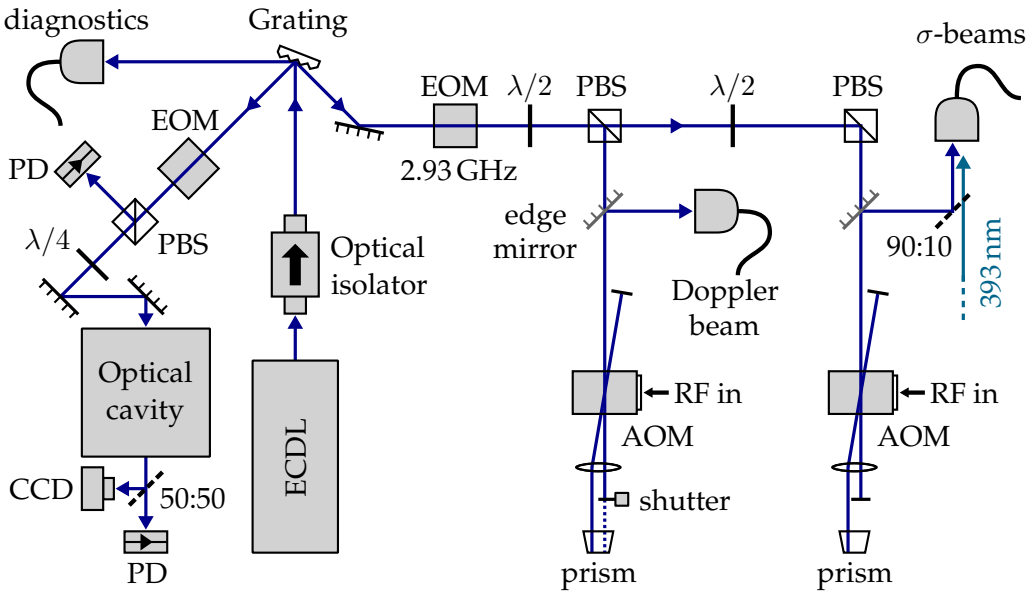

Figure 5.5: Schematic of the 397 nm laser setup, where some lenses, half-wave plates, and some mirrors were omitted for the sake of clarity. The laser output passes through an optical isolator onto a blazed diffraction grating, the lower intensity orders of which are used for PDH locking and diagnostics. The main diffracted order from the grating goes through a 2.93 GHz EOM and is subsequently split into two paths, each going through a double-pass AOM. Since blue AOMs diffract vertically polarized beams most efficiently, we use a prism and lens arrangement to vertically displace the second pass, which is then picked off with an edge mirror and coupled into an optical fibre. On the experimental table, the two fibre outputs become the Doppler cooling beam and the σ_+ polarized optical pumping beam, the latter of which shares its fibre with the 393 nm laser. We use a mechanical shutter in the AOM path of the Doppler beam, which when opened, allows for the 400 MHz red detuned 0^{th} diffracted order to be coupled into the fibre.

AOM is used to switch the beam on and off, as well as to stabilise its intensity via feedback onto the AOM's RF input, as explained in section 5.3.4. As an example of a typical laser setup, the 397 nm beam path is sketched in Fig. 5.5.

It is worth pointing out, that the 397 nm Doppler cooling beam is special, in that we couple both its $+1^{\text{st}}$ and the 400 MHz red detuned 0^{th} order of the associated double-pass AOM[19] into the corresponding fibre. The latter order is blocked off using a mechanical shutter and only switched on between experimental repeats for a few ms to support (high fluorescence) Doppler cooling. We find that this improves our two-ion trapping lifetime by approximately a

[19] IntraAction Corp., ASM-2001.5B8 (200 MHz centre frequency) with DE-2003EM26A driver

factor of two, allowing us to keep the ions for ~ 30 min during experimental operation, and up to a few hours in an idle state, when Doppler cooling only. Also present in the 397 nm laser setup is a 2.93 GHz EOM[20], which modulates sidebands of 85% relative intensity onto the carrier, creating the 397 nm-C and 397 nm-SB frequency tones, described in section 2.4.1, for repumping the $4S_{1/2}$ hyperfine structure during Doppler cooling and fluorescence detection.

The 866 nm laser output, after the isolator and diffraction grating, is split into two paths, each going through a separate double-pass AOM[21] to offset them by 360 MHz from each other, before getting recombined into the same fibre. We use these AOMs to switch between the frequency setting used for high fluorescence Doppler cooling and dark resonance cooling, as described in section 2.4.2.

5.3.3 Raman laser

The two 397 nm Raman beams used for resolved sideband cooling (section 2.4.3), are obtained from a single Toptica DL pro ECDL, detuned from the $4S_{1/2}$ to $4P_{1/2}$ transition by $\Delta \approx -30$ GHz. Since it is the difference frequency of the Raman beams that matters, rather than their absolute frequency, we do not lock this laser. The two beams, labelled R_{H} and R_{V}, are split from the main laser output using a polarizing beam splitter cube (PBS). R_{V} passes through a double-pass AOM[22] with an RF frequency input at ~ 76 MHz before it is coupled into a fibre, while R_{H} goes through a single-pass AOM running at ~ 95 MHz before it is coupled into its own fibre. Both AOMs are optimised for the same ($+1^{\mathrm{st}}$) order so that the resulting difference frequency of the two Raman beams equals the 57.480 MHz splitting between $|M_m = +4\rangle$ and $|M_m = +3\rangle$ of $^{43}Ca^{+}$ at 146 G, with additional small shifts to address radial mode sidebands at $\pm \sim 3$ MHz from this carrier transition.

[20] New Focus, 4431M resonant EOM (3.226 GHz centre frequency)
[21] IntraAction Corp., ATM-851A2 (85 MHz centre frequency) with DE-852EM26A driver
[22] IntraAction, ASM-851.5B8 (85 MHz centre frequency)

Beam	Diameter (μm)	$I_S/\mu W$	Polarisation		
			π	σ_+	σ_-
397 nm Doppler	40	1.70	$1/3$	$2/3$	0
397 nm (σ_+)	37	1.99	0	1	0
397 nm R_V	66	0.63	0	$1/2$	$1/2$
397 nm R_H	57	0.84	$1/2$	$1/4$	$1/4$
393 nm	37	1.88	0	1	0
866 nm	107	2.48	0	$1/2$	$1/2$
850 nm (σ_+)	66	5.97	0	1	0

Table 5.3: List of laser beam characteristics, including their $1/e^2$ beam intensity diameters at the position of the ion, their polarisations, as well as their saturation intensities [Szw09] per μW of beam power. The measurements were conducted by diverting the beams onto a beam profiler using a mirror on a flip mount, placed before the vacuum chamber viewport. The beam profiler was positioned carefully at the same distance from the final lens before the trap as the ions.

The RF drive for the AOMs is generated from two amplified[23] channels of a DDS source, similar to that used for our gate microwaves, described in section 5.4.1. By toggling between DDS frequency profiles, we can select various pre-programmed difference frequencies of the Raman beams, allowing us to pick which motional mode to sideband cool.

5.3.4 Trap optics

Figure 5.3 shows a sketch of all laser beam paths used in this work. The optical setup used for combining, polarisation cleaning and focussing the beams onto the ions did not significantly change from the version described in [Har14]. In brief, each fibre output beam on the main experimental table passes through a PBS, turning any polarisation variations from the fibre into power variations. We then pick off a sample of the beam onto a photo-diode, connected to an FPGA-based intensity stabilisation "noise-eater" [All11], which feeds back to the RF drive amplitude of an AOM earlier in the beam path.

The Raman beams enter the vacuum chamber at 45° to the static magnetic

[23] Mini-Circuits, ZHL-03-5WF, 5 W amplifier

field direction, so that their difference **k**-vector lies parallel to it, allowing the cooling of radial modes of motion. R_{V} is vertically polarised and therefore contains equal contributions of σ_+ and σ_-, while R_{H} is horizontally polarised and contains all polarisation components. The 866 nm beam is combined with R_{V} by passing through the back of the final blue mirror used in the path, and counter-propagates along the 397 nm Doppler beam for more efficient dark resonance cooling [AHS+16].

All beams are focussed onto the ions using 200 mm achromatic doublet lenses. The corresponding beam diameters at the trap centre are listed in table 5.3.

5.4 Microwave drive system

For quantum logic operations on the clock transition in $^{43}Ca^+$ at 146 G, we need four different types of microwave signals to be delivered to the appropriate trap electrodes at the right time, namely, the high power red and blue sidebands for driving the gate, a weak carrier for the dynamical decoupling, and a stronger carrier for single qubit operations, e.g. the analysis pulse. The microwave network used for generating the above signals is sketched in Fig. 5.6.

At its beginning, there is a DDS source, which after up-conversion by 3.5 GHz using a single-sideband (SSB) system, provides signals at the red sideband, blue sideband, and carrier frequency. Each of the signals is output on a separate channel. Furthermore, the DDS source allows for full control over their amplitudes, frequencies and phases.

The red and blue sideband signals are then amplified and combined onto a single line, before going through a cavity filter, to improve spectral purity, and then going through a microwave noise-eater, to improve power stability. Subsequently the, now bichromatic, signal is fed onto the centre axial trap electrode via the diplexer. The carrier channel follows a separate path, going through an amplification and filtering stage before connecting to the diplexer.

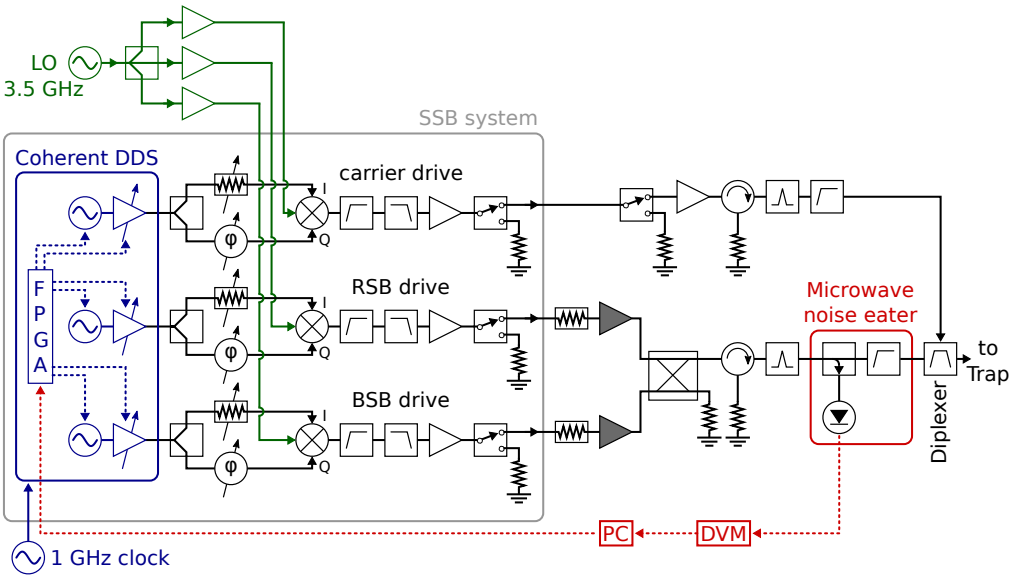

	90° hybrid coupler (Narda Microwave, 4033C)
	Cavity filter - 3.196 to 3.204 GHz (API Technologies corp., BOM22633)
	High pass filter - 1.7 GHz to 5 GHz (Mini-Circuits, VHF-1320+)
	Low pass filter - DC to 3.9 GHz (Mini-Circuits, VLF-3800+)
	Diplexer
	Directional coupler (Pulsar microwave corp., CS20-04-436/13)
	Power splitter/combiner - 2 way, 0° (Mini-Circuits, ZFSC-2-1-S+)
	Power splitter/combiner - 3 way, 0° (Mini-Circuits, ZA3PD-4-S+)
	Amplifier - 16 dB gain (Mini-Circuits, ZX60-4016E+)
	High power amplifier - 46 dB gain (Mini-Circuits, ZHL-16W- 43+)
	Voltage variable attenuator (Mini-Circuits, ZX73-2500-S+)
	SSB switch (Mini-Circuits, ZASWA-2-50DR+)
	180° voltage variable phase shifter (Mini-Circuits, JSPHS-42+)
	IQ-mixer (Eclipse Microwave inc., IQ2040MR4)
	Circulator (Sierra Microwave Technology L.L.C., 1CY63)
	IC-power detector (Analog Devices inc., ADL5511)

Figure 5.6: Microwave drive system used for generating the red and blue sideband (bottom two channels) as well as the carrier for the dynamical decoupling and analysis pulse (top channel). The red, dotted line represents the microwave power feedback loop. It consists of digital multimeters (DVM), which pass the measured voltages from the IC-power detector to the control computer (PC), which in turn instructs the FPGA on board the Coherent DDS to adjust its output amplitudes.

All components, except the DDS source and the diplexer, are housed in SMA connector packages, joined together with short, flexible SMA cabling or I-pieces. The DDS source together with the parts involved in the single-sideband system are mounted on an aluminium baseplate, which in turn is rigidly attached to the optical table, close to the ion trap.

In the following, we will describe the constituents of the above network in further detail. The hardware for generating the transfer pulses for qubit initialisation and enhanced optical pumping is presented in [Har14] and will therefore only be briefly summarised here.

5.4.1 DDS source

The DDS source[24] provides four RF output channels from four Analog Devices AD9910 chips, each of which has eight programmable frequency / phase / amplitude profiles, and an associated pulse-shape. The source is clocked at 1 GHz by an HP 8656B frequency synthesiser, phase-locked to the 10 MHz output of an SRS FS725 rubidium atomic frequency standard. Two DDS channels are used for generating the red and blue sidebands, one channel is used for producing both the strong and weak carrier drive, and the last channel is unused. During an experiment, the DDS receives TTL signals from an FPGA-based pulse sequencer, prompting it to switch between profiles and triggering the playback of a rising or falling pulse-shape. Details about our gate sequence are given in section 7.2, but in short, we use the DDS source to implement four main features. First, we pulse-shape the gate microwave signals to suppress off-resonant excitation transients[25]; second, we linearly ramp the red and blue sideband frequency over the duration of the pulse to suppress an observed motional mode frequency chirp; third, we implement a pre-programmed pulse envelope to reduce the impact of

[24] Enterpoint Ltd, Milldown DDS1 ChannelCard

[25] As explained in section 4.3.4, off-resonant excitation does happen during the "steady-state" portion of the gate, at quite a significant level (10% coherent oscillation of state amplitude).

thermal transients; and fourth, we use the DDS source as part of a feedback loop on the amplitude of the sideband microwaves to improve the long-term stability of the gate Rabi frequency.

5.4.2 Single-sideband system

The function of the single-sideband system is to up-convert the $\omega_{\mathrm{IF}} \approx 2\pi \times 300\,\mathrm{MHz}$ outputs from the DDS source to 3.2 GHz, using an IQ mixer network for single-sideband mixing. To this end, we split each DDS channel into two paths. The first passes through a voltage variable attenuator and connects to the I-port of an IQ mixer, while the second passes through a voltage variable phase shifter and connects to the Q-port of the same mixer. We drive the local oscillator (LO) inputs of the mixers with 10.5 dBm signals at 3.5 GHz, derived from a single signal generator[26], the output of which is split by a three-way power splitter. By properly adjusting the voltages going to the attenuators and phase shifters, we null the upper sideband appearing at $\omega_{\mathrm{LO}} + \omega_{\mathrm{IF}}$ for each channel, leaving the principal outputs at $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$. With the correct DDS frequencies ω_{IF}, this produces the desired signals for the gate.

The final component of each channel on the single-sideband system is a microwave switch (SSB switch), which serves to improve extinction during parts of the sequence where microwaves are not required. These switches also serve to protect the ion trap from accidental, continuous high-power microwave input, as they are wired to a safety interlock which restricts the amount of time they can stay open to a 20 ms period.

5.4.2.1 Bichromatic microwave drive

The red and blue sideband outputs from the SSB system are separately amplified by two high-power (16 W) amplifiers[27], which are cooled by tower CPU coolers

[26] Rohde & Schwarz, SMA100A, 9 kHz-6 GHz RF Analog Signal Generator
[27] Minicircuits, ZHL-16W-43+, $\sim$ 46 dB gain at 3.2 GHz.

with attached fans[28]. The attenuators before the amplifiers are used to roughly set the desired power level and balancing of the two microwave signals. After the amplification stage, the sidebands are combined onto a single line using a $90°$ hybrid, with its unused output ports connected to a high-power terminator. To protect the amplifiers from reverse power, the signal is then passed through a circulator, which diverts retro-reflected microwaves to another high-power terminator.

The microwave line subsequently goes through a cavity filter, to suppress frequency components that might lead to off-resonant excitation. Figures 5.7 and 5.8 show the red and blue sideband power spectrum, measured before and after the cavity filter, demonstrating the improvement in spectral purity. As a final step, the bichromatic signal is sent through a directional coupler, which picks off -20 dB of power for the microwave noise-eating circuit, while the uncoupled portion is fed onto the centre axial trap electrode via the diplexer.

5.4.2.2 Carrier drive

The carrier output from the SSB system goes through an additional microwave switch for enhanced extinction, a 16 dB gain amplifier, a circulator, and a cavity filter before it reaches the diplexer. Within the diplexer, the signal is combined with the trapping RF and DC bias, and passed on to the right axial trap electrode. The DDS source switches between two power setpoints, one for generating short π-pulses on the qubit transition, such as the analysis pulse, and one for the weaker (3 µW into diplexer) carrier drive used for dynamical decoupling.

5.4.3 Microwave amplitude stabilisation

The function of the microwave noise-eater is to stabilise the sideband powers going into the trap, in order to avoid gate errors due to fluctuating or unbalanced

[28] Akasa, Freedom Tower CPU Cooler. Two towers are mounted on home-built aluminium structure, which is in good thermal contact with both high power amplifiers.

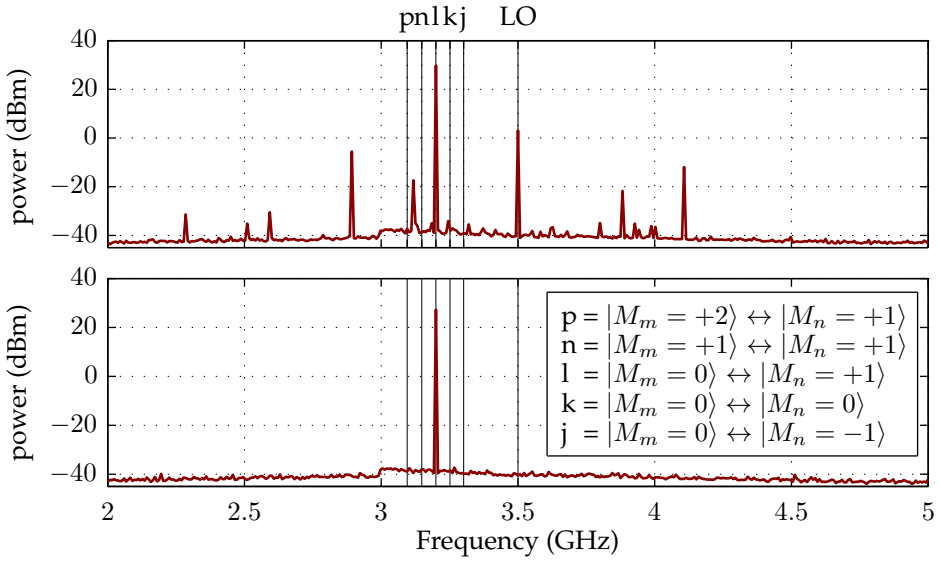

Figure 5.7: Red sideband power spectrum measured before (top) and after (bottom) the cavity filter using a spectrum analyser. Indicated by the black lines are the frequencies for the transitions connecting to the clock qubit (p,n,k,j), the clock qubit itself (l), as well as the local oscillator (LO) at 3.5 GHz. The above transition labelling is in accordance with Fig. A.1 of appendix A. Before the cavity filter the local oscillator leakage is suppressed by ~ 26 dB while all remaining spurs, including the upper sideband at 3.8 GHz are even weaker. After the cavity filter all spurs are reduced to the noise floor of the measuring device, which is ~ 65 dB below the red sideband signal.

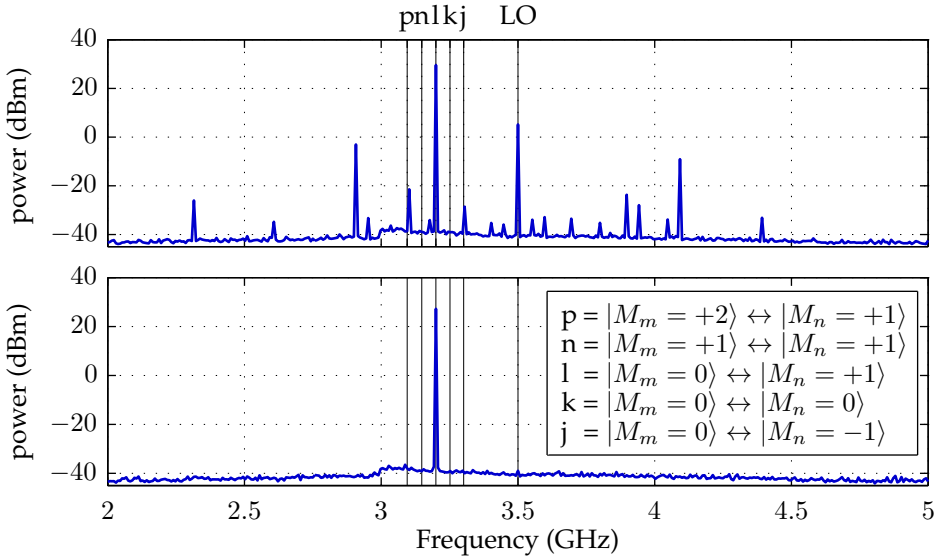

Figure 5.8: Blue sideband power spectrum measured before (top) and after (bottom) the cavity filter. The results are comparable to those obtained for the red sideband, presented above.

sideband Rabi frequencies, which are described in sections 4.3.1 and 4.3.3, respectively. Its principle of operation is to measure the power of a small fraction of each sideband signal, and based on the outcome, adjust the amplitude of the corresponding channel at the DDS source.

The feedback mechanism utilises a directional coupler, located before the diplexer and after the cavity filter, to divert -20 dB of the main signal to an IC power detector[29]. The DC output voltage of the power detector is proportional to the r.m.s. voltage of its input. Towards the end of each shot of the gate sequence, after qubit readout, the sideband microwaves are pulsed on, one sideband at a time, for a duration of 120 μs each. These are the two acquisition periods during which we sample the detector's output for our power measurements. To this end, we use two digital multimeters [30], one for each sideband. However, since 120 μs is shorter than their integration time, we interpose two sample-and-hold amplifiers[31] between the power detector and the multimeters. During the first acquisition period, the detector's output voltage is then captured by one of the sample-and-hold amplifiers, and during the second acquisition period, it is captured by the other. The sampled voltages are held by the amplifiers to bridge the multimeter integration time of 4 ms.

The measured values from the multimeters are registered by the experimental control computer, which takes the average over all shots (typically 512) of one repeat of the experiment and compares it to the set point voltages, the calibration of which is described in section 7.3. It then makes appropriate adjustments to the DDS source's output amplitudes before starting the next repeat of the experiment.

[29] Analog Devices inc., ADL5511, DC to 6 GHz envelope and r.m.s. detector.
[30] Hewlett-Packard, 34401A
[31] Analog Devices inc., AD783, High Speed, Monolithic Sample-and-Hold Amplifier.

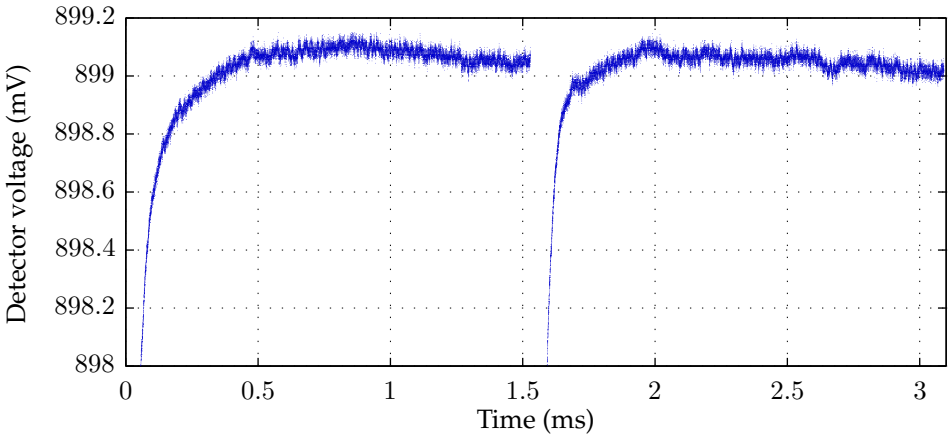

Figure 5.9: Blue sideband pulse power, as measured by the IC power detector for a 3.1 ms gate with a refocussing π-pulse mid-way through, and pre-distortion parameters of $\tau = 2.8$ ms and $k = 0.0214$ (for both gate halves). Although the gate at hand is configured differently than the final version with which we achieved our best fidelity, the resulting pulse quality is consistent with that obtained with our optimal sequence. The pulse in the first gate half is stable to within 0.01% of its mean value after ~ 0.3 ms, while in the second half, after the system had more time to warm up, it takes ~ 0.15 ms.

5.4.4 Microwave amplitude pre-distortion

We use the DDS source described in section 5.4.1 to pre-distort the gate pulse amplitude, in order to reduce the effect of slow thermal power transients on the timescale of ms. We model the shape of the transient by

$$f(t) = f_0 \left(1 - ke^{-t/\tau}\right), \tag{5.1}$$

where τ is the time constant of the exponential growth, and where k is a scale factor. The DDS outputs a pulse with the corrected envelope

$$f'(t) = \frac{1-k}{1-ke^{-t/\tau}}, \tag{5.2}$$

so that the resulting pulse going into the trap has a constant envelope of $f(t) \times f'(t) = f_0(1-k)$. Since our gate is divided into two halves, separated by a refocussing π-pulse, we introduce a separate scale factor for the second half, but use the same τ.

The above parameters are calibrated experimentally using the microwave noise-eater, following the procedure outlined in section 7.3. Fig. 5.9 shows the

typical pulse quality we achieve with optimised settings. We remark that an important additional measure for reducing power transients is a warmup pulse, applied as shortly before the gate as possible. For a 3.25 ms gate we use a 1.7 ms warmup pulse, as described in section 7.2.

5.4.5 Initialisation drive

For the $\{A, B, C, D, E, Q\}$ transfer π pulses used in the qubit initialisation and enhanced optical pumping procedure, we use a frequency octupled AD9910 DDS, running at 400 MHz and clocked at 1 GHz from the same HP 8656B frequency synthesiser as the DDS source of the SSB system. The octupler is assembled from standard Mini-Circuits components and did not change from the version described in [Har14].

6

Laser cooling of $^{43}Ca^{+}$ at 146 G

In this chapter, we discuss the experimental implementation and optimisation of the dark resonance Doppler cooling scheme, introduced in section 2.4.2. Using Raman sideband thermometry, we systematically explore the parameter space spanned by the intensities and detunings of the beams involved, to find robust settings that minimise temperature and can readily be adapted to give high fluorescence for qubit readout. Despite the difficulties introduced by the large number of widely separated states in $^{43}Ca^{+}$ at 146 G, we were able to demonstrate cooling to a temperature of 0.3 mK, slightly below the Doppler limit for the corresponding two-level system. We also show how, with the addition of a continuous Raman sideband cooling step, the ion's radial motional modes can be cooled further, close to their motional ground state.

The theoretical analysis of the data presented below, using time-dependent optical Bloch equation simulations, is given in [AHS+16, Jan15]. A preliminary experimental investigation of dark resonance Doppler cooling in $^{43}Ca^{+}$ at 146 G, using microwave sideband thermometry, can be found in [Har14].

6.1 Sideband thermometry

We measure the temperature of a particular mode in a single $^{43}Ca^{+}$ ion with the help of sideband thermometry [LBMW03], using our Raman beams. This

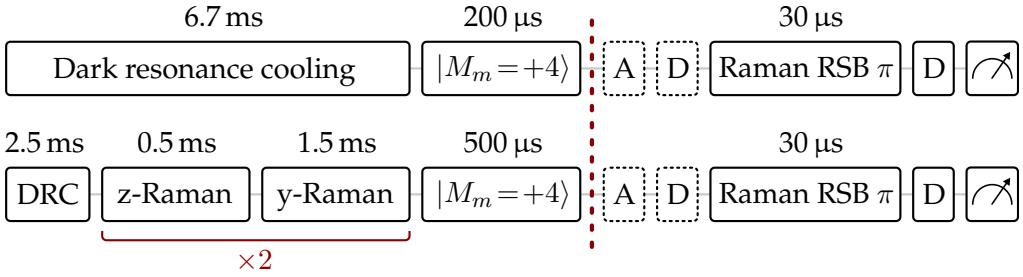

Figure 6.1: Pulse sequences used for sideband thermometry after dark resonance cooling (DRC, top figure) and resolved sideband cooling (bottom figure). The latter proceeds via two repeats of alternating Raman first red sideband cooling pulses, first addressing the z-mode (z-Raman) and then the y-mode (y-Raman). The length of the dark resonance cooling period was chosen so that pulses involved in the thermometry part of the sequence, after the red dotted line, appear at the same point in the 20 ms line-triggered experimental cycle as they did in the resolved sideband cooling case. The difference in optical pumping duration to $M_m = +4$ has no impact on the quality of the state preparation. The dashed A and D pulses are used to probe the blue sideband (see text).

method relies on comparing the spin flip probabilities after exciting the ion via the first red and blue sideband transition of the mode of interest. We deduce the temperature from the ratio of the spin-flip probabilities

$$\frac{P_\mathrm{r}}{P_\mathrm{b}} = \frac{\bar{n}}{\bar{n}+1}, \tag{6.1}$$

where P_r (P_b) is the spin flip probability for the first red (blue) sideband. This ratio is independent of the Rabi frequency and Lamb-Dicke parameter, and can be inferred from either detuning or time scans of both sidebands. In this work, we chose the former, as Rabi flopping time scans are sensitive to drifts in the motional mode or in the light shifts caused by the Raman beams.

The pulse sequences underlying the thermometry experiments are illustrated in Fig. 6.1. They consist of an initial period of cooling, using the method to be investigated, followed by preparing $|M_m = +4\rangle$ according to the recipe given in section 2.5. For a red sideband scan, we apply a red Raman π-pulse, which takes population from $|M_m = +4\rangle \otimes |n\rangle$ to $|M_m = +3\rangle \otimes |n-1\rangle$. For a blue sideband scan, we additionally insert the microwave transfer π-pulse sequence $\{A, D\}$ before the red Raman π-pulse, which now takes population from $|M_m = +3\rangle \otimes |n\rangle$

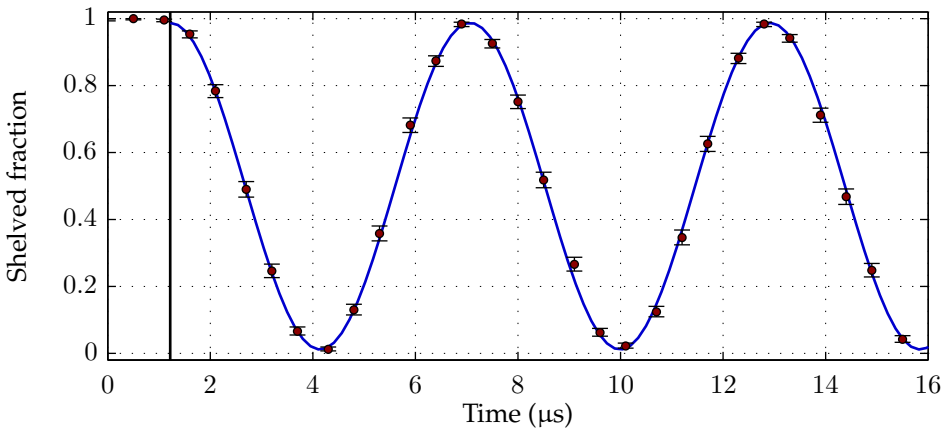

Figure 6.2: Raman carrier time scan with both radial modes cooled close to the ground state. The Rabi frequency is fitted to be 171(2) kHz, while the AOM dead time, indicated by the black line, is 1.2 μs.

to $|M_m = +4\rangle \otimes |n+1\rangle$. In this way we are capable of driving either red or blue sidebands without having to change the frequency of the Raman beams. This reduces potential errors due to, for example, variations in the thermal response of the AOMs with frequency. Finally, we apply another D pulse before switching to our high fluorescence laser parameters, to be introduced shortly, for readout.

Two calibration experiments are required in preparation for the thermometry experiments. First, we measure the single-ion motional mode frequencies using a tickle experiment, explained in section 8.2. The axial mode appeared at 477 kHz, while the radial modes were at 3.270 MHz and 3.630 MHz, which we will from now on refer to as y- and z-modes. Unless stated otherwise, all thermometry is carried out on the 3.270 MHz y-mode.

Second, we characterise the effect of our Raman beams on the ion. Recalling section 2.4.3, they operate $\Delta = -30$ GHz detuned from the $2S_{1/2} \leftrightarrow 4P_{1/2}$ transition, connecting the $|M_m = +3\rangle$ and $|M_m = +4\rangle$ Zeeman states, separated by 57.480 MHz. To measure the carrier π-time, we cool both radial modes near the ground state using the sequence shown in Fig. 6.1, but leaving out the thermometry. We then take a time scan with the Raman lasers on resonance with the carrier transition, as shown in Fig. 6.2. The fit gives a Rabi frequency of

171(2) kHz and an AOM dead time of 1.2 µs, which is a consequence of the AOMs switching the two Raman beams on and off at slightly different times, due to acoustic propagation delay. Unless stated otherwise, this dead time is included whenever we quote a pulse length.

In a similar way, we calibrate the Raman first blue sideband π-time to be 30 µs. Since we want to use this pulse for thermometry after dark resonance cooling as well as after resolved sideband cooling, we have to make sure that, in both cases, it appears at the same point in the (line-triggered) experimental cycle. The 50 Hz magnetic field noise, explained in section 5.2.6, may otherwise lead to a reduction in the quality of the pulse, as the $|M_m = +3\rangle \leftrightarrow |M_m = +4\rangle$ transition is sensitive to magnetic field fluctuations. We therefore match the length of the dark resonance cooling period to that of the resolved sideband cooling, as shown in Fig. 6.1.

6.2 Experimental Doppler cooling optimisation

We reiterate that there are three laser frequencies, bundled into two beams, used for Doppler cooling $^{43}Ca^{+}$ at 146 G. The two 397 nm beam frequencies 397 nm-C and 397 nm-SB have a fixed separation of 2.930 GHz and a fixed intensity ratio of $1/0.85$, where 397 nm-C is the component with the higher intensity and the lower frequency. The polarisations of both above beams as well as that of the 866 nm beam are given in table 6.1.

Although we use the first Raman sidebands for our temperature measurements, we found that it was easier to minimise the temperature when looking at the height of the second red Raman sideband. To this end, we used a pulse length of ~ 50 µs, which is significantly shorter than the second sideband π-time. We searched the parameter space, spanned by the 397 nm and 866 nm detunings and intensities, to minimise the spin-flip probability on the second sideband transition, which thus minimised the temperature. The best parameters we found

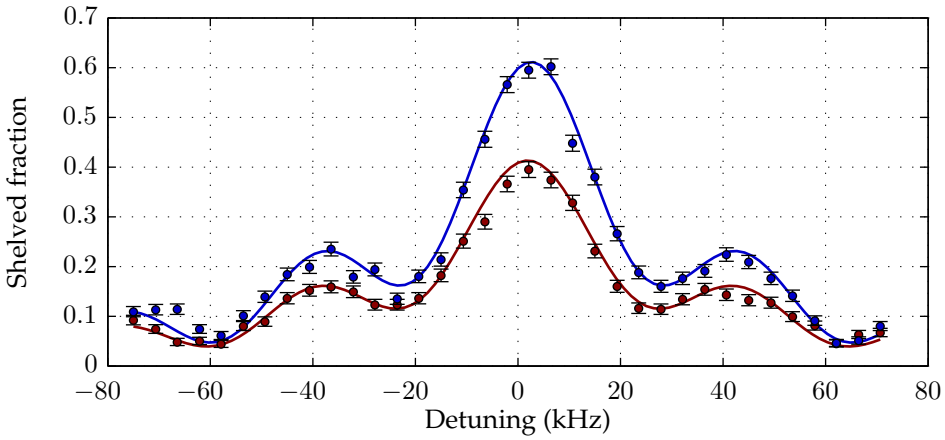

Figure 6.3: Detuning scans of the Raman laser over the first red and blue sideband resonances of the radial y-mode, using our optimised dark resonance Doppler cooling parameters. The fitted temperature is $\bar{n} = 1.97(5)$, corresponding to an effective temperature of $T \approx 0.3\,\mathrm{mK}$.

were 5 μW of total 397 nm power in a 40 μm $1/e^2$ intensity diameter beam, and 40 μW of 866 nm power in a 107 μm $1/e^2$ intensity diameter beam. In terms of detunings, we operate the 397 nm laser near resonance, ~ 5 MHz below the point where we can no longer reliably hold the ion in the trap due to Doppler heating. The 866 nm laser is set roughly half way up the red side of the dark resonance, as illustrated in Fig. 6.4. With these optimised parameters, we took a pair of red and blue first Raman sideband scans, shown in Fig. 6.3, from which we deduced a y-mode temperature of $\bar{n} = 1.97(5)$.

Using these settings as a starting point, we map out the dependence of the temperature reached by dark resonance Doppler cooling on the 397 nm and 866 nm powers and detunings. We generally scan one parameter at a time, while keeping the others at their optimised values. The only exception is the 397 nm detuning, which we treat last, and for which we re-optimise the 866 nm detuning at each step of the scan, using the second red sideband. This allows us to vary 397 nm frequency in order to find its best setting, without moving away from the optimal point on the dark resonance, as this has already been investigated separately, when looking at the 866 nm detuning. At each value of the varied

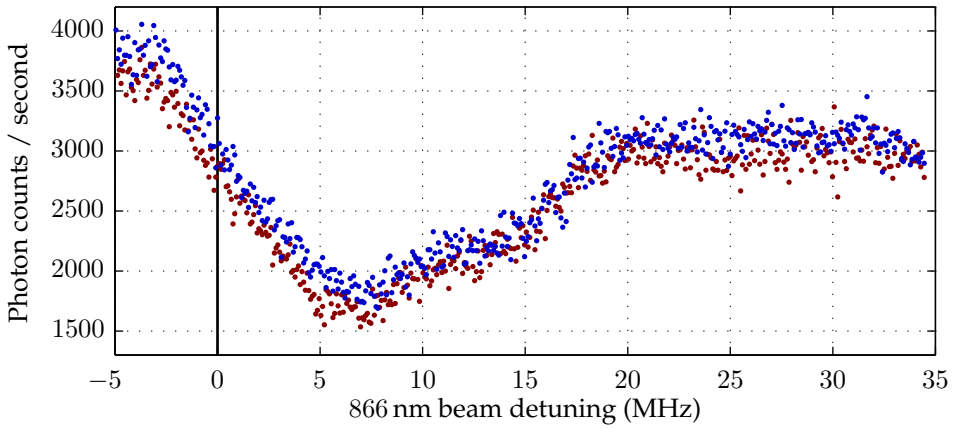

Figure 6.4: Measurement of the 397 nm fluorescence as the frequency of the 866 nm beam is scanned across the dark resonance in the positive direction (red), and reversed (blue), with all remaining parameters set to their optimised values. The detuning used in Fig. 6.3 is at 0 MHz, highlighted by the black line. We subtracted the background signal of 1028 counts/s from this data.

parameter, we take a red and blue first sideband scan for fitting the corresponding temperature. The results of all four scans are shown in Fig. 6.5. We can see that the dark resonance cooling gives temperatures below the two-state Doppler limit over a range that is wide enough for it to not be exceeded through typical experimental fluctuations in each parameter.

We believe that the main systematic error in the above data comes from changes in detuning, caused by drifts in the lasers' frequency reference cavities. To estimate these, we tracked the position of the dark resonance using 866 nm frequency scans, similar to that of Fig. 6.4, interleaved between the data runs. The drifts are below ~ 2 MHz over the duration of each of the four experiments.

To switch between dark-resonance and high-fluorescence cooling, we raise the 397 nm power from 5 μW to 75 μW, while leaving its detuning unchanged, and the 866 nm power from 40 μW to 330 μW, while increasing its frequency by 360 MHz. We use an additional AOM to achieve this frequency shift. The resulting fluorescence amounts to $\sim 25\,000$ counts/s, with a net photon detection efficiency of 0.17%. This is comparable to the count rate from a single, saturated,

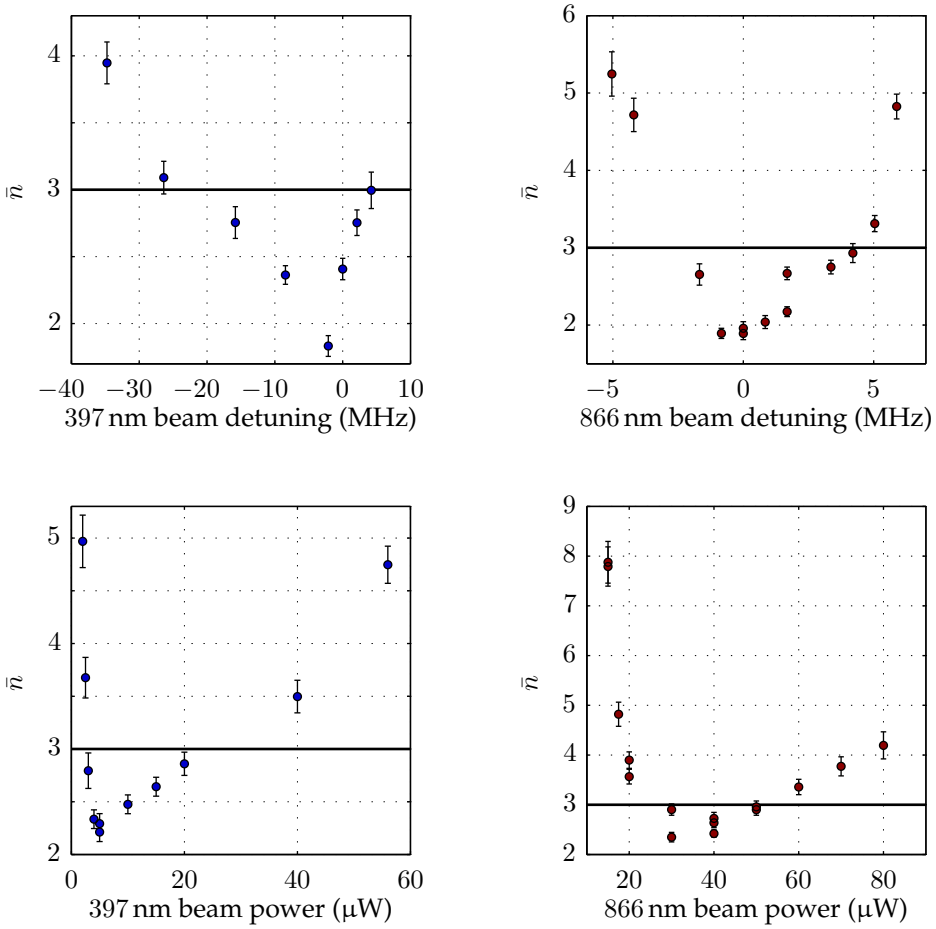

Figure 6.5: Radial y-mode temperature reached by dark resonance Doppler cooling as a function of 397 nm and 866 nm powers and detunings, with the frequency separation and intensity ratio between 397 nm-C and 397 nm-SB fixed. The error bars are the statistical uncertainty from the sideband fits. Indicated by the black line is the Doppler limit for a two level system with the same 22 MHz natural linewidth as the Ca^+ $S_{1/2}$ to $P_{1/2}$ transition. The frequency offset in the detuning scans is set so that 0 MHz corresponds to the optimal settings used in Fig. 6.3. The visible deviation of the best 397 nm detuning from 0 MHz in the above graph is caused by drifts in the lasers' frequency reference cavities. A comparison with a theoretical numerical model can be found in [AHS+16, Jan15].

	Dark res. cooling		High fluorescence		Polarisation		
	Intensity (I_S)	Detuning (MHz)	Intensity (I_S)	Detuning (MHz)	π	σ_-	σ_+
397 nm-SB	1.05(20)	1042(10)	15.8(30)	1042(10)	$1/3$	0	$2/3$
397 nm-C	1.24(20)	-1888(10)	18.6(30)	-1888(10)	$1/3$	0	$2/3$
866 nm	161(20)	-159(10)	1330(170)	201(10)	0	$1/2$	$1/2$

Table 6.1: Optimal parameters for dark-resonance cooling and for obtaining high fluorescence, taken from [AHS+16]. Intensities, detunings and their uncertainties are derived from fits to 866 nm fluorescence spectra, using theoretical analysis presented in [Jan15]. The saturation intensity I_S is defined as $I_S = (2\pi hc)/(3\lambda^3\tau)$, where τ is the lifetime of the $4P_{1/2}$ level. Detunings are given with respect to the separation of the centres of gravity of the $2S_{1/2}$ and $4P_{1/2}$ levels, and (for column 3) mark the position of the bottom of the dark resonance. In case of the 866 nm frequency, this corresponds to ~ 6.5 MHz in Fig. 6.4.

$^{40}Ca^{+}$ ion, which is sufficient to allow for fluorescence detection fidelities of $\sim 99.98\%$ [Bur10].

Using our optimised settings for dark resonance Doppler cooling and high fluorescence readout, we conduct a series of ~ 1 GHz wide 866 nm fluorescence scans. Based on fits to the resulting data, using time dependent optical Bloch equations [Jan15], it is possible to independently derive the associated beam intensities and detunings. The obtained values are given in [AHS+16] and reproduced in 6.1.

6.3 Resolved sideband cooling characterisation

To reach the motional ground state, we employ the above described, optimised dark resonance Doppler cooling scheme, and follow it by resolved Raman sideband cooling. The underlying sequence for measuring the resulting temperature is shown in Fig. 6.1. We take a pair of first sideband detuning scans on the two modes of interest, shown in Fig. 6.6 and Fig. 6.7. The fitted temperatures are $\bar{n} = 0.08(1)$ for the y-mode, and $\bar{n} = 0.16(2)$ for the z-mode. The Lamb-Dicke parameters can be deduced by comparison with the carrier Rabi frequency, giving $\eta = 0.098(2)$ and $\eta = 0.082(2)$ for the y- and z-mode, respectively. This enables

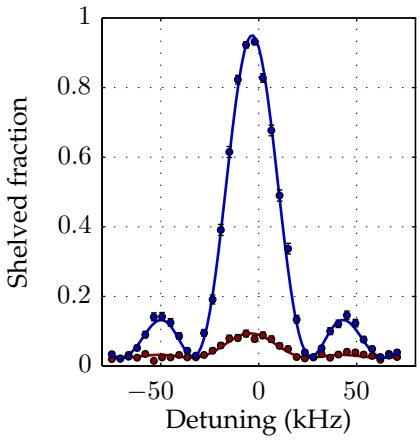

Figure 6.6: Red and blue first Raman sideband detuning scans of the radial y-mode. The fitted Lamb-Dicke parameter is $\eta = 0.098(2)$, while the fitted temperature is $\bar{n} = 0.08(1)$.

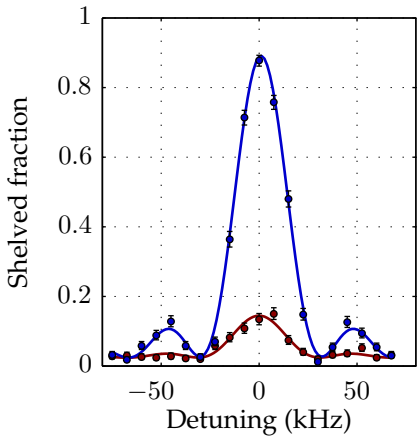

Figure 6.7: Red and blue first Raman sideband detuning scans of the radial z-mode. The fitted Lamb-Dicke parameter is $\eta = 0.082(2)$, while the fitted temperature is $\bar{n} = 0.16(2)$.

us to determine the direction of the y-mode axis to be $40°$ with respect to the static magnetic field direction.

For the two-qubit gate experiments, we use a similar continuous sideband cooling sequence on the two radial rocking modes, but with slightly modified pulse lengths (see section 7.2), allowing us to achieve $\bar{n} \sim 0.1$ on both modes. With larger Raman power and detuning Δ, we should be able to achieve even lower temperatures.

7

Two-qubit gate implementation

In the following, we present the results of our two-quit gate experiment, which is a DDMS gate with 99.7(1)% fidelity. We discuss the pulse sequence used, along with all necessary calibration experiments for setting up the two-qubit operation. The chapter concludes with a presentation and discussion of the results.

7.1 Gate parameters

Although already mentioned in previous chapters, we briefly recap the most important settings used for our highest fidelity DDMS gate. Of the two-ion radial y- and z-rocking modes at 3.255 MHz and 3.588 MHz, respectively, we choose the y-mode for our gate. The radial y- and z-COM modes were measured to be at 3.286 MHz and 3.615 MHz, while the axial x-COM mode is at 450 kHz.

We apply 2 W in each sideband, measured before the diplexer, to the centre axial trap electrode, producing a gate Rabi frequency of $\Omega^{(2)} = 2\pi \times 308$ Hz. The resulting differential AC-Zeeman shift is measured to be $\omega_{\mathrm{Z}} = 2\pi \times 20.78$ kHz. Both sideband phases ϕ_{R} and ϕ_{B} are equal when generated by the DDS source, but since no special precautions were taken to exactly match their corresponding signal's path lengths, we assume that they are not equal any more once they reach the ion trap. To ensure and maintain the correct gate pulse area, we pre-distort the sideband pulse envelope to suppress slow power transients, and employ the

microwave noise eater to reduce long term power drifts. Further improvements to the long term thermal stability are achieved by maintaining a constant duty cycle of the gate microwaves during the dead times between gate experiments. We also apply a 200 Hz linear ramp to both sideband frequencies, in order to minimise the impact of an observed chirp of the motional mode frequency.

The gate time is $\Delta t_g = 3.25$ ms and the gate detuning is $\delta_g = 2\pi \times 1.23$ kHz, producing four loops in phase space. A 3.2 µs π-pulse with a $\bar{\phi} + \pi/2$ phase is applied mid-way through the gate, rotating the spin state around an orthogonal axis to that set by $\bar{\phi}$. It serves to implement a $W_1(t/\Delta t)$ Walsh modulation as well as to refocus any partially complete carrier oscillations. For the dynamical decoupling, we apply a 3 µW signal to the right axial trap electrode, which is resonant with the AC-Zeeman shifted qubit transition and gives 12 complete carrier flops with $\Omega_{\mathrm{C}} = 2\pi \times 3.69$ kHz.

7.2 Gate sequence

The full gate sequence starts with a 2 ms period of high fluorescence Doppler cooling. During this time, as well as between sequences, we employ the additional 400 MHz red detuned 397 nm cooling beam to improve the two-ion lifetime. Subsequently, we apply 2 ms of dark resonance cooling, using the optimised parameters found in section 6.2. We then apply two cycles of continuous resolved sideband cooling, each cycle consisting of a 750 µs Raman pulse addressing the radial z-rocking mode, followed by a 1.5 ms Raman pulse on the radial y-rocking mode. This concludes the cooling stage of the sequence, leaving both radial rocking modes in a thermal state of $\bar{n} \sim 0.1$, and the COM modes at temperatures of $\bar{n} \sim 2$.

We then employ 300 µs, i.e. two cycles, of microwave enhanced optical pumping to prepare the ions in $|M_m = +4\rangle$, as described in section 2.5. At this point, the only microwave transition connecting the population to the rest of

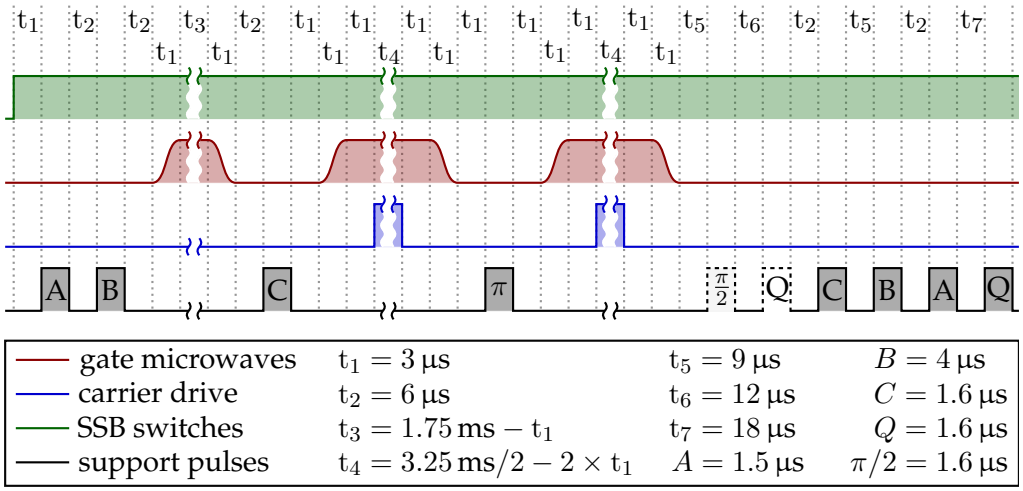

Figure 7.1: Timing diagram (not to scale) of the central portion of the gate sequence, where t_3 is the gate warm-up pulse and t_4 is one half of the gate. The gate microwaves are smoothly switched on and off over a period t_1. Dead times arise due to the time it takes to toggle DDS profiles. The first, optional, Q pulse is used for measuring the readout levels. The refocussing π-pulse consists of two 1.6 µs long $\pi/2$-pulses, separated by a dead time of 1 µs, the latter of which can be disabled to facilitate calibration experiments. In a similar way, the analysis $\pi/2$-pulse can be extended to to a π-pulse.

the ground state manifold is the $|M_m = +4\rangle \leftrightarrow |M_n = +3\rangle$ stretch transition at 2.87 GHz. Thus, we are safe to open the microwave sideband switches (SSB switches, see section 5.4.2) without causing off-resonant excitation due to spectral impurities from e.g. any remaining local oscillator leakage at 3.5 GHz. We recall that the purpose of these switches is not to turn the gate microwaves on and off, which is done by the DDS using amplitude pulse shaping, but rather to serve as a safety feature in case of DDS failure, as well as to improve extinction during parts of the sequence other than the two-qubit gate.

With the SSB switches open, we apply the microwave transfer pulses A and B, leaving the ions in $|M_m = +2\rangle$, and then turn the gate microwaves on for 1.75 ms in order for the system to warm up. We then use a C pulse to take the population to $|1\rangle$, and subsequently switch on the gate microwaves, which are interrupted at the half-way point by the refocussing π-pulse. Thanks to the warm-up, the power transients during the gate are less severe, and can be compensated by

pre-distorting the sideband pulse envelope. The carrier drive turns on after the sideband pulse rise time and turns off before the fall time begins.

After the gate is completed, we transfer the population in $|1\rangle$ back to $|M_m = +4\rangle$ and $|0\rangle$ to $|M_n = +1\rangle$ by executing the π-pulse series $\{C, B, A, Q\}$. To enable read-out level measurements (section 7.5) from either qubit state, we have the option of inserting another Q pulse before the $\{C, B, A, Q\}$ transfer sequence. We then use one cycle of the shelving sequence given by Eq. 2.6, followed by 3 ms of fluorescence detection using our high fluorescence Doppler parameters, and a 250 µs long 854 nm deshelving pulse. As the last step before closing the SSB switches and re-enabling Doppler cooling with the 400 MHz red-detuned beam, the microwave sidebands get pulsed on once more, one at a time for a period of 120 µs each, back-to-back. During this time, the microwave noise eater measures their respective powers.

The pulse sequence is illustrated in Fig. 7.2, while a detailed timing diagram of the centre portion is given in Fig. 7.1. We line-trigger the sequence in order to minimise the impact of the 50 Hz ambient magnetic field fluctuations (see section 5.2.6) on our state preparation and measurement (SPAM) error. These field fluctuations are one of two reasons for why we generally keep all pulses at fixed positions, and replace those that are not needed for a particular experiment with delays. The second, equally important, reason for why we keep the timing constant, is to maintain control over thermal effects. Consequently, we are not able to study reliably the gate's dynamics by means of time scans, and use detuning scans instead.

7.3 Gate set-up procedure

In the following section we describe the necessary steps to set up our DDMS gate. The main advantage of this procedure is that it can, to a large extent, be performed on a single ion. When setting the gate up from scratch, we usually

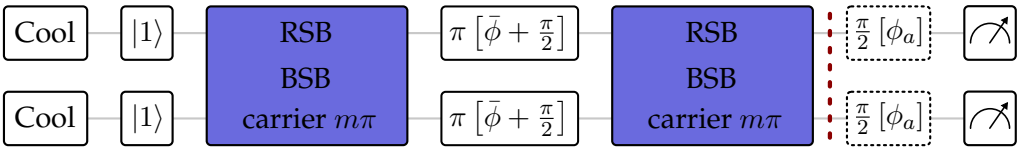

Figure 7.2: DDMS gate sequence with a refocussing π-pulse mid-way through. Uncoloured pulses are executed on the bare qubit transition while blue colouring indicates an AC-Zeeman shift due to the gate microwaves. The carrier drive induces m complete flops in each half of the gate, where m is an integer. At the red, dotted line the sequence has produced the Bell state given in Eq. 4.42, at which point the associated populations can either be measured directly, or we can append an analysis $\pi/2$-pulse with a variable phase ϕ_a, in order to produce a parity fringe.

Figure 7.3: Single-ion pulse sequence used for measuring the AC-Zeeman shift generated by the sideband microwaves. We apply the carrier drive in the first half of the gate and scan its frequency to find the shifted qubit transition.

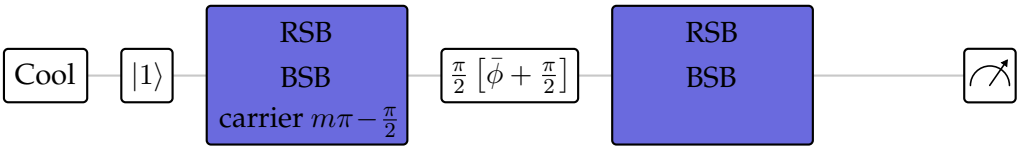

Figure 7.4: Single-ion pulse sequence used for calibrating the phase of the refocussing π-pulse. The carrier is only applied in the first gate half, with its amplitude set to produce $(m-1/2)$ Rabi flops. We reduce the refocussing π-pulse to a $\pi/2$-pulse and scan its phase to find the value that gives equal populations in $|0\rangle$ and $|1\rangle$.

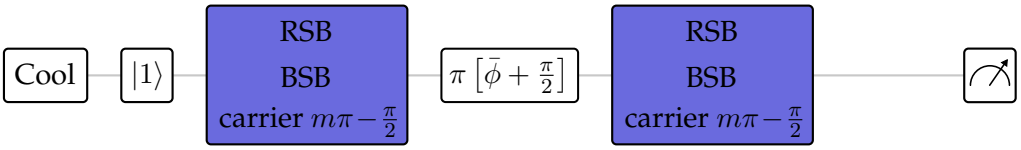

Figure 7.5: Single-ion pulse sequence used for calibrating the phase offset between the carrier pulses in the two gate halves, both of which are configured to give $(m-1/2)$ Rabi flops, where m is an integer. We scan the phase of the carrier pulse in the second gate half and look at the value that minimises the population in $|1\rangle$.

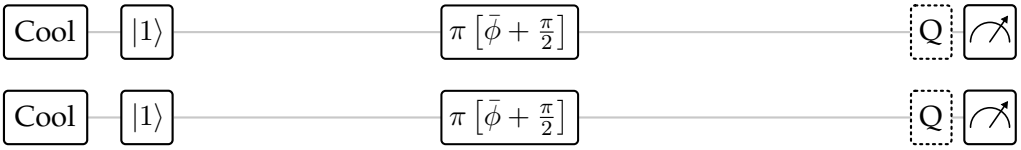

Figure 7.6: Pulse sequence used for measuring the SPAM error. We measure the populations in $|00\rangle$ and $|11\rangle$ with both the gate and the carrier drive disabled. For the latter state, we add an additional Q transfer π-pulse at the end of the sequence.

have to go through this procedure multiple times, as some parameters are inter-dependent, while others require the gate to be fully set up before it becomes evident that further adjustments are needed.

Step 1 - Setting the microwave sideband power: In order to set the red and blue sideband Rabi frequencies, we first roughly tune the associated microwave powers to within 1 dB of the target value for the gate time we are aiming for. We do this by monitoring the signal going into the diplexer using a handheld power meter, and manually swapping attenuators before the 16 W amplifiers. Next, we balance the the sideband powers up to the 0.1 dB resolution limit of the above power meter, by adjusting the output amplitudes of the associated DDS channels.

For fine-tuning, we change over to measuring sideband powers while running the gate sequence, using the IC power detector from the microwave noise-eater, described in section 5.4.3. We set the gate time to its calculated value for the expected Rabi frequency, and execute the gate sequence, with only one sideband enabled at a time. We then let the sample-and-hold amplifiers sample continuously while monitoring the output of one of them directly on a fast oscilloscope, bypassing the multimeters. In this way we are able to see the IC power detector's output voltage over the course of the two gate halves, allowing us to not only fine tune the DDS output amplitudes and improve the sideband power balancing to within 0.03 dB, but also to optimise the pulse envelope pre-distortion parameters.

We conclude the power calibration procedure by setting the microwave noise eater set point voltages. To this end, we repeatedly run the sequence as for normal gate operation, while recording the detector's output voltage, sampled during the two 120 µs noise eating periods after readout. The mean of the measured voltages over typically $\sim 30\,000$ runs is then used as our microwave noise eater set point.

Step 2 - Setting the carrier drive Rabi frequency: In this step we tune the carrier drive strength for the dynamical decoupling to roughly set the desired number of Rabi flops over the course of the gate duration. To this end, we simplify our gate sequence by disabling the refocussing π-pulse, disabling the gate microwaves to avoid AC-stark shifts, but retaining the carrier drive, which we use to measure Rabi flopping over the course of the gate time on a single ion. Although not strictly required for the proper operation of the gate, we aim at setting a whole number of carrier flops, as this will prove convenient for future calibration steps. Fine-adjustment of the carrier amplitude will take place in step 4. We note that in order for the dynamical decoupling to be effective, $\Omega_{\mathrm{C}} \gg \Delta_{\mathrm{Z,rms}}$ must be satisfied.

Step 3 - Compensating the static AC-Zeeman shift: We measure the static AC-Zeeman shift ω_{Z} caused by the gate drive, by means of Rabi spectroscopy on a single ion, using the pulse sequence illustrated in Fig. 7.3. It involves applying the carrier drive, which has previously been configured on the unshifted qubit transition, during the first gate half. The sideband microwaves, which are off-resonant with the one-ion motional modes, are kept on during both halves to retain similar thermal behaviour as during the actual gate, and to produce the AC-Zeeman shift under investigation. By simultaneously scanning the frequencies of both the microwave sidebands and the carrier by the same amount, we produce data as shown in Fig. 7.7, which we then fit to determine ω_{Z}. From this step onwards, we take the AC-Zeeman shift into account for both the carrier and the gate drive.

Step 4 - Setting pulse amplitudes: There are a total of four amplitudes in need of calibration, namely, the carrier drive amplitude in the first gate half, the carrier drive amplitude in the second gate half, the refocussing π-pulse amplitude and the analysis $\pi/2$ pulse amplitude. Each uses a separate DDS amplitude profile

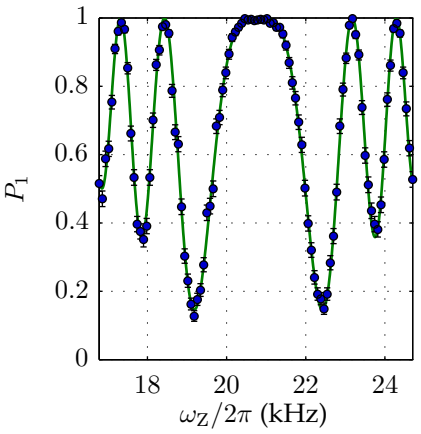

Figure 7.7: Rabi spectroscopy on the AC-Zeeman shifted qubit transition of a single ion. The fit gives $\omega_Z = 2\pi \times 20.781(3)$ kHz.

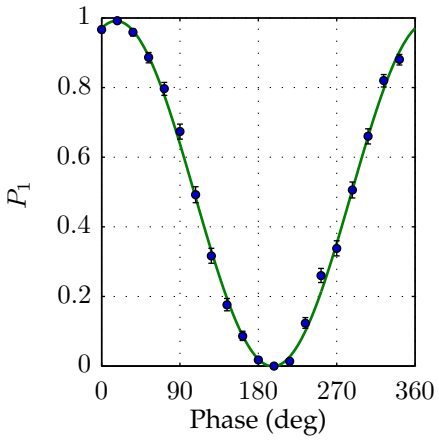

Figure 7.8: Refocussing pulse phase calibration scan, using the sequence illustrated in Fig. 7.4. The desired phases occur at $P_1 = 0.5$.

that can be used to make adjustments. The experiments described below are all conducted on a single ion.

To begin with, we fine-tune the carrier drive strength in the first gate half, for which we utilise the same sequence as shown in Fig. 7.3. We recall that in step 2, we have already roughly configured the carrier amplitude in both halves to give a whole number of flops over the course of the full gate. This proves to be useful now, as we can simply conduct a narrow DDS amplitude scan of the carrier drive in the first gate half, to maximise (minimise) the resulting population in $|1\rangle$, provided we have set an even (odd) number of flops in step 2. Next, we add the refocussing π-pulse to the above sequence and, as before, scan its DDS amplitude to minimise (maximise) the resulting population in $|1\rangle$.

Subsequently, we correct for potential over- or under-rotations of the spin state by the carrier drive in the second half of the gate, with respect to the first. This is important, as the refocussing π-pulse only protects from errors due to partially complete carrier Rabi oscillations, if the pulse areas preceding and succeeding it are equal. The underlying calibration experiment employs the full gate sequence (Fig. 7.2) and, similar to above, relies on scanning the

corresponding DDS amplitude. If we have set an even (odd) number of whole Rabi flops over the course of the full gate, we now look to minimise (maximise) the populations in $|1\rangle$ after the second gate half. In this work, the discrepancy between the carrier strengths in the two gate arms is non-zero, but less than 0.2%, and stems from thermal effects following the reactivation of the sideband microwaves after the short cool-down period due to the refocussing π-pulse.

Finally, we configure the analysis $\pi/2$-pulse. It is easiest to set amplitudes if the pulses involved in the sequence leave the qubit in one of its two states, but not in a superposition, since then the calibration amounts to finding the population minimum or maximum, instead of a specific value. With this in mind, we use the full gate sequence, including the 1.6 µs long analysis pulse, which we temporarily extend by appending another 1.6 µs long pulse after a 1 µs delay. We scan the joint DDS amplitude to make them act as a π-pulse and subsequently remove the second pulse again, leaving us with a calibrated $\pi/2$ analysis pulse.

Step 5 - Setting pulse phases: There are two phases that can be set with the help of a single ion, namely, the phase of the refocussing π-pulse, and the phase of the carrier drive in the second half of the gate. The former needs to be configured so that it induces spin state rotations around an axis perpendicular to that of the carrier drive. However, prior to applying the refocussing π-pulse, the gate microwaves are ramped down over the course of 3 µs, causing the qubit to acquire a phase offset due to the gradual decrease of the AC-Zeeman shift.

To calibrate out this offset, we use the sequence illustrated in Fig. 7.4, which includes two modifications to our standard sequence. First, we transform the refocussing π-pulse into a $\pi/2$-pulse. This is straightforward, since the refocussing pulse, in fact, consists of two $\pi/2$-pulses, separated by a 1 µs dead time, and we can simply replace one of them with a delay. Second, we tweak the carrier drive amplitude in the first half of the gate to effectively give a $\pi/2$-pulse, while leaving it out altogether in the second half. We then scan the refocussing $\pi/2$-pulse from

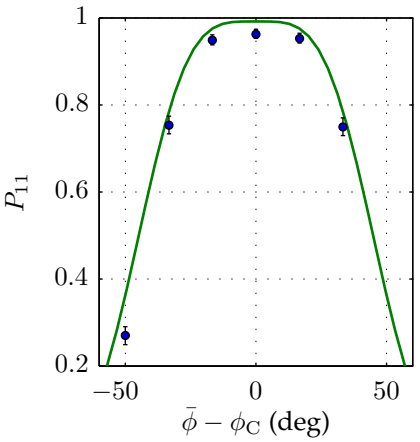

Figure 7.9: Two-ion populations as a function of the phase mismatch between $\bar{\phi}$ and ϕ_C, measured using the full DDMS gate sequence. The solid theory curve assumes our usual gate parameters. Its deviation from the data is due to an uncompensated motional frequency chirp.

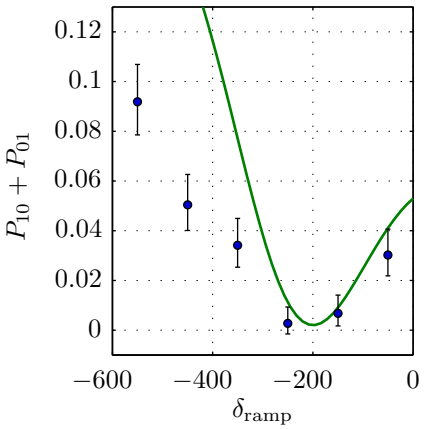

Figure 7.10: Single-ion populations as a function of a linear gate detuning ramp, used to compensate a motional mode frequency chirp. The solid theory curve assumes our usual gate parameters, with the addition of a 200 Hz linear chirp, which is cancelled using $\delta_{ramp} = -2\pi \times 200\,\text{Hz}$.

0 to 360°, inducing the population in $|1\rangle$ to trace out a fringe, as shown in Fig. 7.8, with the sought-after phase occurring at $P_1 = 1/2$.

Next, we set the phase of the carrier drive in the second half of the gate, which differs from that in the first half due to changes in the qubit transition frequency, resulting from the interruption of the gate microwaves by the refocussing π-pulse. The discrepancy is calibrated out using the sequence shown in Fig. 7.5. It is equivalent to the full gate sequence, except that the amplitudes of the carrier drive in both gate arms are set to effectively give a $\pi/2$-pulse each. As before, we then scan the phase of the pulse under investigation from 0 to 360°, with the sought-after phase occurring at the minimum of the population fringe.

Step 6 - Calibrating the spin phase: In order for the carrier drive to dynamically decouple the gate from AC-Zeeman shift fluctuations, we have to match the spin phase $\bar{\phi}$ to the phase of the carrier drive ϕ_C, or to $\phi_C \pm \pi$, as explained in section 4.4.2. To this end, we employ the full gate sequence on two ions, and use

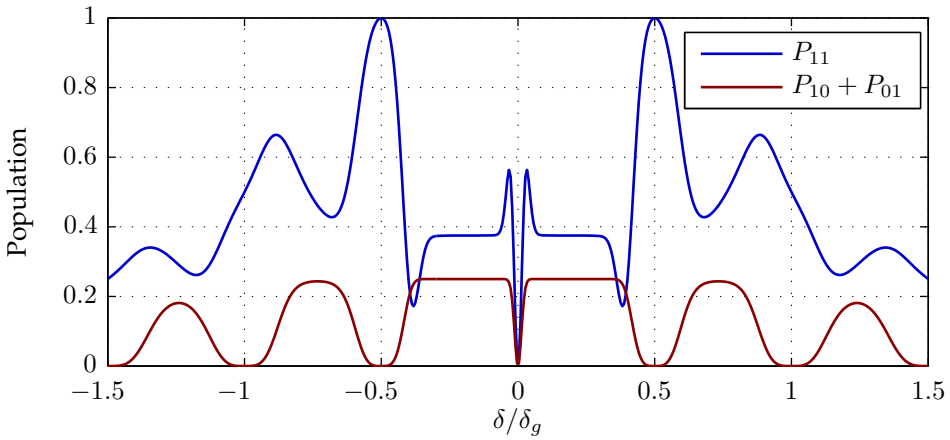

Figure 7.11: State populations as a function of fractional gate detuning δ/δ_g, for an ideal $K = 4$ loop DDMS gate with a refocussing π-pulse mid-way through. At $\delta = 0$, the gate microwaves are resonant with the motional sidebands and cause a linear displacement in motional phase space, which is reversed by the Walsh pulse. When $\delta = \pm\delta_g/2$, we trace out two loops in phase space, with twice the intended radius and the spin flip in-between. In this case, we acquire double the desired amount of geometric phase. For $\delta = \pm\delta_g$, the gate is configured properly, giving $P_{10} + P_{01} = 0$ and $P_{11} = P_{11} = 1/2$.

the two-ion populations P_{11} as our diagnostic for steering $\bar{\phi}$ towards its correct value. We work at $\delta = \pm\delta_g/2$, since at this gate detuning P_{11} has a distinctive maximum, shown in Fig. 7.11, which makes it less sensitive to motional mode drift. The height of the feature depends on the deviation of $\bar{\phi}$ from ϕ_{C}. When scanning $\bar{\phi}$, we produce data as shown in Fig. 7.9, from which we fit the value of $\bar{\phi}$ that maximises the two-ion populations to typically within $\pm 1^{\circ}$. According to the simulations carried out in section 4.4.2, this accuracy is sufficient for the associated gate error to be below 10^{-6}.

Step 7 - Compensating the motional mode frequency chirp: In the final step, we apply a linear ramp to both sideband frequencies over a range δ_{ramp}, the aim of which is to reduce the observed chirp in the motional mode frequency. We believe that this chirp is a result of thermal transients in the trap, caused by the high-power gate microwaves. The ramp is configured so that at the beginning of the gate $\delta = \delta_g + \delta_{\mathrm{ramp}}/2$, and at the end $\delta = \delta_g - \delta_{\mathrm{ramp}}/2$, leaving the average

gate detuning at δ_g.

To set the optimal value of δ_{ramp}, we employ the full gate sequence on two ions, but this time we use $P_{10} + P_{01}$ as our diagnostic. We work at $\delta = \delta_g$, where the single ion populations $P_{10} + P_{01}$ should ideally be zero for a perfect gate, as can be seen from Fig. 7.11. The calibration procedure relies on scanning δ_{ramp} while minimising $P_{10} + P_{01}$, as shown in Fig. 7.10. We see the data does not match the theory curve, implying that the actual frequency chirp may have nonlinear components. However, what matters is that with the help of our linear ramp we are still capable of reaching the minimum, which in out experiment appears at $\delta_{\mathrm{ramp}} = -2\pi \times 200\,\mathrm{Hz}$.

Step 8 - Adjusting the gate time: In principle, the DDMS gate is fully configured after step 7, provided the gate time we assumed at the very beginning of the set-up procedure is correct. To check whether this is the case, we scan the gate detuning and measure the associated state populations produced by the gate, as shown in Fig. 7.12. We use the resulting data to fit the Rabi frequency, using the analytic propagator from Eq. 4.62, and asses whether the deviation of the fitted value from $\Omega^{(2)}\Delta t = \pi\sqrt{K}$ is large enough to justify an adjustment of Δt. Based on this assessment, we either change the gate time and work through the entire set-up procedure again, or we move on to other experiments.

7.4 Fidelty measurement

The highest gate fidelity we achieve, using the settings laid out in section 7.1, is 99.7(1)%, characterised with the help of the reduced process tomography procedure, described in section 4.2.4. The underlying data consists of five parity scans, building up Fig. 7.13, interleaved with measurements of the populations and finished with a measurement of the SPAM error. Each data point consists of 512 shots of the experiment, after which the microwave noise eater makes

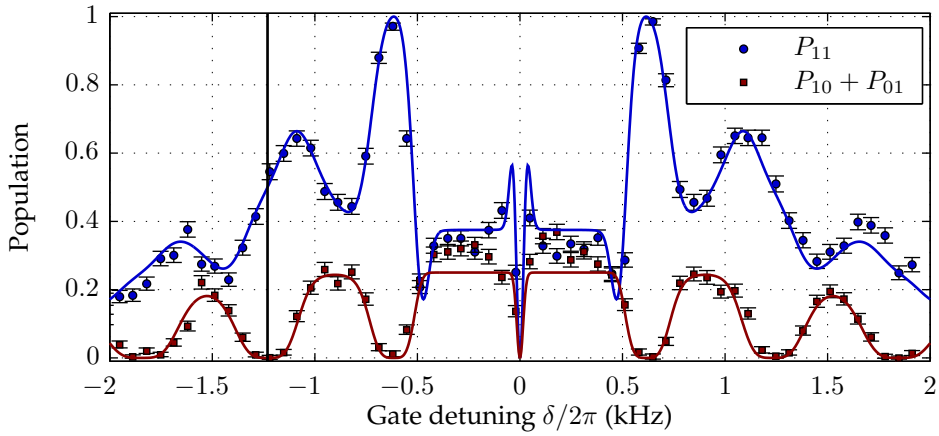

Figure 7.12: State populations as a function of gate detuning δ. The solid red and blue line represent numerical simulations of the ideal gate with $\Delta t_g = 3.25\,\text{ms}$, in the absence of any dissipation mechanisms and assuming a perfectly ground state cooled motional mode. At the black line the detuning conforms to Eq. 4.38, resulting in populations that are consistent with the maximally entangled Bell state given in Eq. 4.42, that is $P_{10} + P_{01} = 0$ and $P_{11} = 1/2$.

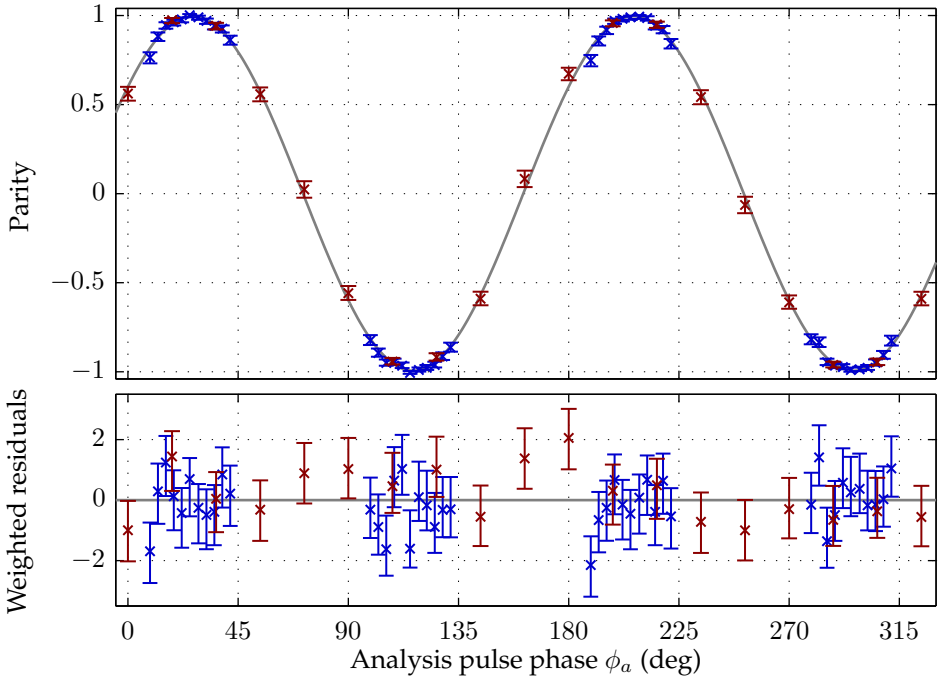

Figure 7.13: Parity signal generated by our best gate, built up of five separate experimental runs, interleaved with population measurements. Every point consists of 512 shots of the experiment, where the associated error bars indicate 68% confidence intervals. All data have been corrected for SPAM errors and fitted using maximum-likelihood estimation, assuming a binomial distribution, resulting in an oscillation amplitude of 0.9953(23). Shown below are the weighted residuals, giving a reduced $\chi^2 = 0.68$.

adjustments to the sideband powers. For reasons of thermal stability, the gate microwaves were fired into the trap throughout the whole experimental run, with an approximately constant duty cycle, even when no sequence was being executed. The full data set was acquired over a period of ∼1 hour, during which the motional mode drift was consistent with the "flat" intervals shown in Fig. 8.5.

Using the sequence illustrated in Fig. 7.2, we measure the $P_{11} + P_{00}$ populations after the gate to be $0.9980(8)$. To obtain the magnitude $\left|\rho_{11,00}\right|$, we scan the phase of the analysis pulse ϕ_a and map out a parity fringe. Since the parity measurement is a binomial experiment, that is, the parity is either even or odd, we can use a maximum-likelihood method assuming binomial statistics to fit the data. This is described in more detail in the supplementary material of [BHL+16], as well as in section 8.3.4 of [Bal14].

We fit the parity fringe with the function

$$\mathcal{P}(\phi_a) = C_0 + C_1 \sin(2\phi_a + \phi_0), \tag{7.1}$$

where the fitted amplitude offset C_0 is consistent with zero, and the fitted amplitude $C_1 = 0.9953(23)$. The phase offset ϕ_0 of the fringe is obtained from an independent scan, taken at the very beginning of the experimental run, and is not floated in the fit. This exemplifies the stability of the phase of the Bell state produced by the gate, as explained in section 4.2.4.

7.5 State preparation and measurement error

The above data have been corrected for SPAM errors, which we measure using the same gate sequence as presented in section 7.2, but with both the carrier drive and the gate microwaves turned off, as illustrated in Fig. 7.6. We repeat the sequence to give 2×10^4 trials, in half of which we expect to measure all population in $|00\rangle$, and in the other half of which we insert an additional Q

microwave transfer π-pulse before the readout procedure, so that we expect to find all population in $|11\rangle$. Detecting the ions in a state other than the one that was prepared is registered as an error. We define the average SPAM error of a single ion as $1/2\,(\epsilon_D + \epsilon_B)$, where ϵ_B represents events in which the ion was prepared in the bright state $|0\rangle$ but detected in the dark state $|1\rangle$, and similarly for ϵ_D.

The average SPAM error measured in this experiment is 0.34(3)% per qubit, which is about five times higher than 0.068(5)% figure reported in [HAB+14] for single-ion work in the same trap. We suspect that this discrepancy arises for three main reasons.

First, the separation between the $P_{01} + P_{10}$ and the P_{11} histograms when using photon-count thresholding on two-ions is much smaller than that between the P_0 and P_1 histograms on a single ion. This makes above-threshold events in the $P_{01} + P_{10}$ histogram, as well as below-threshold events in the P_{11} histogram more likely, which increases ϵ_D.

Second, providing sufficient separation between the $P_{01} + P_{10}$ and the P_{11} histograms relies on long fluorescence detection periods, giving the metastable $3D_{5/2}$ shelf more time to spontaneously decay. In this experiment, we count photons for 3 ms, resulting in a histogram distribution illustrated in Fig. 2.9 of section 2.6. Consequently, due to shelf decay during this time span, ϵ_D is $3\,\text{ms}/1168\,\text{ms} = 0.26\%$, which contributes 0.13% to our average SPAM error, assuming ϵ_B negligible. In comparison, this error contribution was 1.5×10^{-4} in [HAB+14] due to a shorter fluorescence detection time, as well as due to the utilisation of the time-resolved photon counting method [MSW+08].

Third, the length of the microwave transfer π-pulses is less carefully tuned compared with [HAB+14], resulting in a higher error per pulse due to off-resonant excitation of spectator transitions and magnetic field noise, as explained in section 5.2.2 of [Har14].

It is worth pointing out that the first two of the three above reasons can be attributed to the fact that we use a PMT for fluorescence detection, which only allows us to resolve how many ions are fluorescing, but not which ones they are. This problem can be remedied by using a CCD camera instead, as it would provide spatial resolution and allow the states of each of the ions in the string to be determined with high fidelity, as explained in [BSWL10] and in chapter 5 of [Bur10].

8

Two-qubit gate characterisation

In the following chapter, we describe a selection of purpose-designed experiments, used for characterising the most important sources of infidelity that affect our DDMS gate. This complements the theoretical error analysis of section 4.3, in that we now provide details about how some of the quantities used therein were measured.

Single-qubit results that have previously been obtained with our current ion trap, and are presented in [Har14, HAB+14], will not be repeated here. We merely recall that the 146 G clock qubit coherence time is $T_2^* = 50(10)$ s, the combined state preparation and measurement (SPAM) error is $6.8(6) \times 10^{-4}$, and the average single-qubit gate error, measured by means of randomised benchmarking, is $1.0(3) \times 10^{-6}$.

8.1 Heating rate

We measure the heating rate of the motional mode under investigation using sideband thermometry, explained in section 6.1. To this end, we employ the sequence illustrated in Fig. 8.1, which starts with cooling the mode of interest near the ground state, and subsequently preparing the $|M_m = +4\rangle$ state. We then turn off all laser beams to let the mode heat due to ambient electric field noise for a period of time, after which we apply a red or blue sideband pulse, the detuning

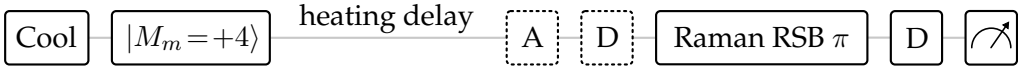

Figure 8.1: Pulse sequence used for measuring the heating rate, applicable to one or two ions. After ground state cooling, we prepare the $|M_m = +4\rangle$ state, and subsequently wait for a varying amount of time to allow for heating. We then apply either a red or a blue Raman sideband pulse, the latter in the form of a red Raman sideband pulse preceded by the microwave transfer π-pulses A and D, which eliminates the need to change the Raman laser frequency. Lastly, we apply another D pulse and read out from the stretch qubit, i.e. $|M_n = +3\rangle$ and $|M_m = +4\rangle$.

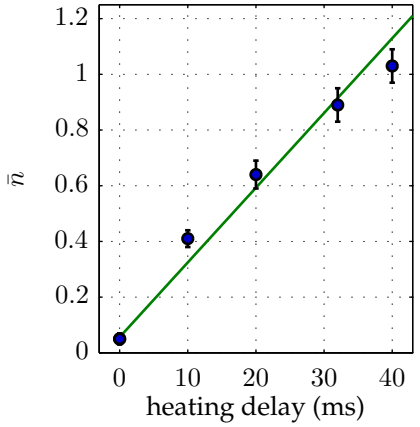

Figure 8.2: Heating of the single-ion radial COM mode at 3.286 MHz as a function of time. A linear fit gives $\dot{\bar{n}} = 27(2)$ quanta/s.

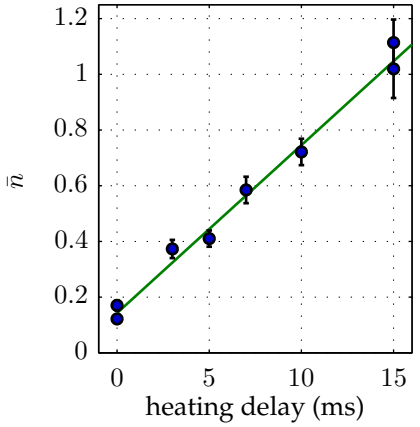

Figure 8.3: Heating of the two-ion radial COM mode at 3.286 MHz as a function of time. A linear fit gives $\dot{\bar{n}} = 60(15)$ quanta/s.

of which is scanned by 40 kHz as part of the thermometry. Acquiring and fitting a pair of sideband scans for several different delays allows us to map temperature versus heating time. A linear fit to the data yields the heating rate $\dot{\bar{n}}$.

In this trap, the heating rate of the single-ion 3.286 MHz radial COM mode is $\dot{\bar{n}} = 27(2)$ quanta/s, as shown in Fig. 8.2. For the two-ion radial COM mode with the same frequency we measure $\dot{\bar{n}} = 60(15)$ quanta/s (Fig. 8.3), which is approximately twice the single ion heating rate, as expected. The heating rate of the two-ion radial rocking mode at 3.255 MHz is lowest, and even with delays of up to 80 ms there was negligible measurable effect. Leaving two ions in the dark for such long times drastically increases the probability of ion loss, making the measurement challenging. Furthermore, the nearby COM mode heats significantly faster, so that the side lobes of its sinc^2 sideband feature, when

radial mode	frequency (MHz)	$\dot{\bar{n}}$ (quanta/s)	S_E $\left(10^{-13}\,\mathrm{V}^2/(\mathrm{m}^2\mathrm{Hz})\right)$
one-ion COM	3.286	27(2)	6.5(5)
two-ion COM	3.286	60(15)	15(4)
two-ion rocking	3.255	$\lesssim 5$	$\lesssim 1.2$

Table 8.1: Summary of the measured heating rates expressed in terms of the linear increase of the motional mode occupation $\dot{\bar{n}}$, as well as the spectral density of the electric field noise $S_E = (4m\hbar\omega\dot{\bar{n}})/e^2$ [BKRB15], where e is the charge of the ion, m is its mass and ω is the mode frequency.

scanning the Raman laser frequency, overlap with the much smaller feature of the rocking mode, even at short heating delays. This makes accurate fitting of the latter mode's temperature difficult. Based on the data, which can be found plotted in appendix C, we have estimated the two-ion radial rocking mode heating rate to be $\dot{\bar{n}} \lesssim 5$ quanta/s, which according to simulations carried out in section 4.3.6 amounts to a gate error of $\lesssim 0.2\%$. We summarise our measurements in table 8.1.

These resulting values are very low for a room-temperature surface trap and rather resemble heating rates observed in cryogenically cooled setups (see [HCW+12] for a review). As anomalous heating is not yet fully understood, it remains unclear as to precisely which part of the trap fabrication procedure is responsible for the low heating rates.

8.2 Radial mode frequency stability

Optimal gate performance requires the bichromatic field to be precisely tuned with respect to a selected motional mode, as otherwise the ions' internal and external degrees of freedom are left entangled at the end of the gate operation. Consequently, any changes of the mode's frequency after we have set the gate detuning relative to it, lead to errors. We estimate the significance of this effect in our system by monitoring the frequency of one of the radial COM modes of a single ion over time. The gate rocking mode frequency fluctuations are expected to be very similar, as $\omega_{\mathrm{rocking}} = \sqrt{\omega_{\mathrm{radial}}^2 + \omega_{\mathrm{axial}}^2}$ and our axial mode frequency

Figure 8.4: Pulse sequence used for the tickle experiment, in which we locate a motional mode's resonance by scanning the frequency of the tickle pulse.

ω_{axial} at ~ 500 kHz is much smaller than our radial COM mode frequencies ω_{radial} at ~ 3 MHz.

To this end, we employ the "tickle" method, which relies on the pulse sequence illustrated in Fig. 8.4. It starts with sideband-cooling the mode under investigation to the ground state, followed by initialising the $|M_m = +4\rangle$ state using our standard preparation techniques, outlined in section 2.5. We then apply a 20 ms long, -61 dBm, ~ 3 MHz tickle pulse[1] to the left trap electrode. The associated field acts as an oscillating force, exciting the ion's motion to a coherent state, when on resonance with a vibrational mode. Subsequently, we probe the ion's red motional sideband using a Raman π-pulse, the duration of which was appropriately tuned beforehand to take into account the change in Rabi frequency due to the increased motional state of the ion after tickling. Finally, we apply a D microwave π-pulse to transfer population from $M_m = +3$ and $M_n = +3$, so that we can read out from $M_m = +4$ and $M_n = +3$ using our usual detection scheme, explained in section 2.6.

There are two possible outcomes of this experiment for a given tickle frequency. On the one hand, we may detect the ion in its initial spin state ($P_1 = 1$), from which we conclude that the Raman sideband π-pulse had no effect, and that the ion's motion therefore remained in the ground state. This means that the tickle force is off resonant at the frequency used. On the other hand, we may detect spin flips ($P_1 \neq 1$), from which we infer then the tickle frequency is close to resonance and excited the motion, which caused the Raman π-pulse to induce a population transfer between the qubit states. Repeating this procedure for a

[1] The tickle pulse is generated using an arbitrary waveform generator, outputting a 180 mV_{pp} sinusoid that goes through a 20 dB attenuator, a microwave switch, an RF isolation transformer and another 30 dB attenuator, before being connected to the diplexer.

series of tickle frequencies, and recording the resulting population change in the qubit, gives rise to data as shown in the inset of Fig. 8.5. It can be fitted using a Lorentzian function, which for the parameters used in this work has on average a FWHM of 50(10) Hz, localising the position of the resonance to within 6(2) Hz. Conducting a series of these experiments back-to-back allows us to track changes in the motional mode frequency over time. We classify the mode fluctuations measured in this way into three groups, depending on the timescales they appear on in our experiment.

First, slow drifts on the timescale of hours, which can be taken out by manual adjustments of the gate detuning. As an example, consider Fig. 8.5, showing tickle data that was acquired over a period of three consecutive days and exhibits peak drift rates of $\sim 120\,\mathrm{Hz/h}$, as well as flat intervals, into which we generally try to fit our experiments. We believe that the cause for this behaviour is thermal, as we see drifts on similar timescales when operating the calcium oven or changing the high field coil currents. More specifically, passing 5 A through an initially cold oven, induces an exponential motional mode frequency decay with a time constant of 12 hours, and initial gradients of order $\sim 3\,\mathrm{kHz/h}$. A probable explanation for this effect is radiative heating of the ion trap, as there is a direct line of sight between the chip and the resistively heated oven filament. Likewise, the thermalisation process after powering up the high field coils, which represent a $\sim 34\,^{\circ}\mathrm{C}$ heat source in contact with the vacuum chamber, causes the motional mode frequency to drift towards lower values, with an initial rate of $\sim 2\,\mathrm{kHz/h}$. However, at the time the data presented in Fig. 8.5 were measured, we were aware of both these problems. It was therefore common practice to leave the high field coils turned on continuously and keep the ion loading times as short as possible. Thus, in the hope of gaining more insight into the cause of the long term fluctuations, we monitored the ambient air temperature using a sensor, mounted above the imaging system. Its data output is also plotted in Fig. 8.5, but does

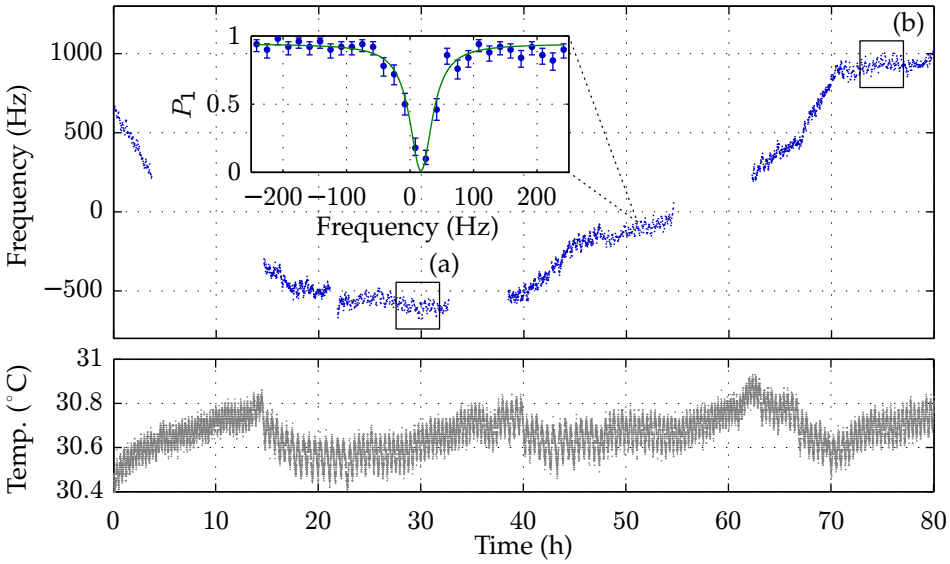

Figure 8.5: Frequency drift of the single ion COM-mode around its average of 3.297 MHz, monitored over a period of three days. The top plot shows the full data set, with gaps due to ion loss. Each point is generated from a Lorentzian fit to a tickle scan, a typical one of which is shown in the inset. The data within the two frames, labelled (a) and (b), can be found magnified in Fig. 8.6. Illustrated below the main figure is the ambient air temperature, measured above the imaging system, inside the thermal enclosure of the optical table.

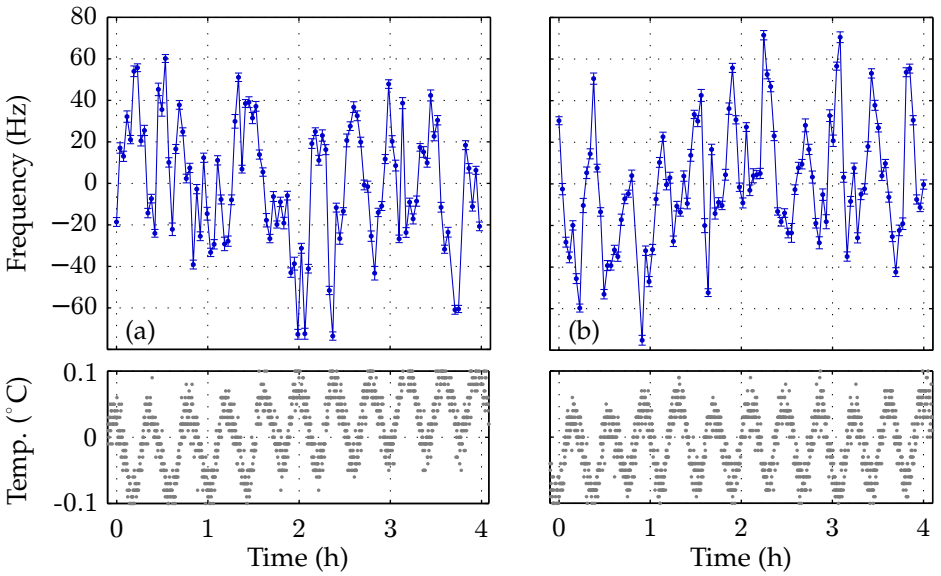

Figure 8.6: Magnified view of the radial mode drift from within frames (a) and (b) of Fig. 8.5, around their respective means. Data points are spaced by ~ 2 min and show average fluctuations of ~ 30 Hz r.m.s. Plotted below are the associated ambient air temperature variations, also around their means (30.6 °C in (a) and 30.7 °C in (b)). The oscillations by ± 0.1 °C are related to the lab's air-conditioning cycle, with a periodicity of ~ 23 min.

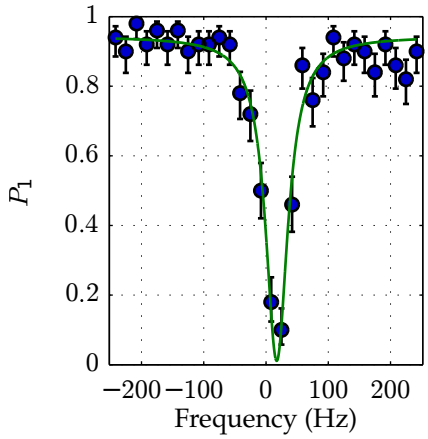

Figure 8.7: Example of tickle scan capturing the single-ion radial mode while approximately stationary.

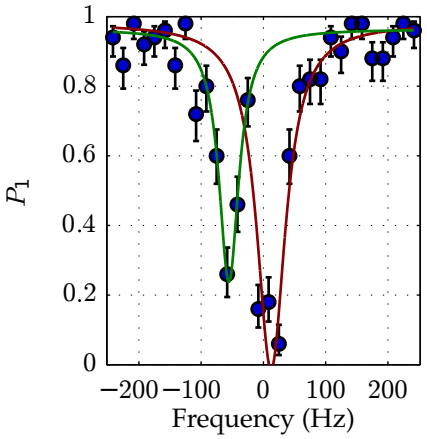

Figure 8.8: Example of tickle scan capturing an apparent abrupt jump of the radial mode by 68 Hz during a single scan.

not show any obvious common features with the long term radial mode drift, implying that the source of the perturbation might originate from within the vacuum chamber.

Second, motional mode fluctuations on the minute timescale, which we can take to be stationary over the course of the gate operation, but changing between experimental runs. To estimate their size, we examine the two approximately drift-free periods (a) and (b) from Fig. 8.5, shown enlarged in Fig. 8.6. The data points are acquired every ~ 2 min, which is comparable to typical durations of e.g. a parity scan. There is a visible correlation between the frequency variations and the lab's air-conditioning cycle of ~ 23 min periodicity. However, we surmise that further contributing factors, of potentially acoustic nature, are present, as we also observe sporadic mode frequency jumps by up to 70 Hz within a single tickle scan. An example of such a sudden shift can be found in Fig. 8.8, while Fig. 8.7 illustrates a stationary signal, for comparison.

Both data sets in Fig. 8.6 exhibit average fluctuations of ~ 30 Hz r.m.s. around their respective means. We model the associated gate error for a basic Mølmer-Sørensen gate in section 4.3.2, and for a Walsh modulated gate in section 4.4.1,

assuming for simplicity that the frequency variations are normally distributed. For our optimal, $K = 4$ loop gate sequence, the simulated error is $\sim 0.2\%$, while without the refocussing π-pulse, this error would be $\sim 0.5\%$.

Finally, motional mode frequency jitter on even shorter timescales, that is, during the gate pulse. We expect fast fluctuations of this sort to have a broadening effect on the lineshape of the tickle resonances, which are Fourier limited to 50 Hz by the tickle pulse duration of 20 ms. Recalling that the average FWHM of our Lorentzian fits to these resonances is 50(10) Hz, we can bound the excess broadening to $\lesssim 10$ Hz. The impact of fast motional mode frequency jitter on the gate fidelity is best captured using the motional coherence formalism, which is the subject of the subsequent section.

8.3 Motional coherence

We attempted to measure the motional coherence time using a Ramsey experiment on a superposition state between $|n = 0\rangle$ and $|n = 1\rangle$ of the single-ion radial COM mode. This involved ground state cooling and preparing the $|M_m = +4\rangle$ state, then applying a 2.5 µs Raman carrier $\pi/2$-pulse to create the spin superposition $(|M_m = 4\rangle + |M_m = 3\rangle) \otimes |n = 0\rangle$, as illustrated in Fig. 8.9, and subsequently using a 37.5 µs Raman red sideband π-pulse to reach the motional superposition $|M_m = 4\rangle \otimes (|n = 0\rangle + |n = 1\rangle)$, as shown in Fig. 8.10. We then let this state evolve freely before mapping back to the spin superposition with another Raman red sideband π-pulse, followed by a Raman carrier $\pi/2$-pulse, the phase of which was scanned to produce a Ramsey fringe.

The idea of the experiment was to observe the decay in fringe contrast for different periods of free evolution and associate it to the motional coherence time τ_{C}. However, we found that this decay was overwhelmed by the reduction in fidelity of the red sideband π-pulses. The reason for this were shifts of the associated transition's frequency, mainly due to slow fluctuations of the laboratory

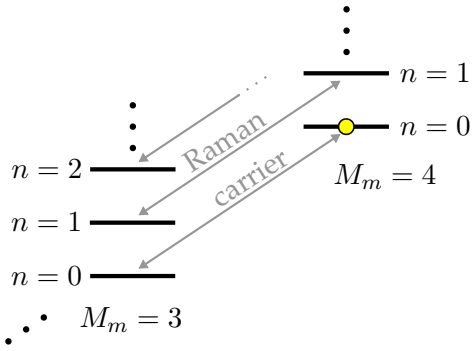

Figure 8.9: Simplified energy level diagram of the ground state of $^{43}Ca^{+}$, including the lowest few motional states $|n\rangle$. In order to generate a motional superposition state from the initial $|M_m = 4\rangle \otimes |n = 0\rangle$ state (yellow circle), we first create a spin superposition $(|M_m = 4\rangle + |M_m = 3\rangle) \otimes |n = 0\rangle$ using a Raman carrier $\pi/2$-pulse. The states coupled by this pulse are indicated by the grey arrows.

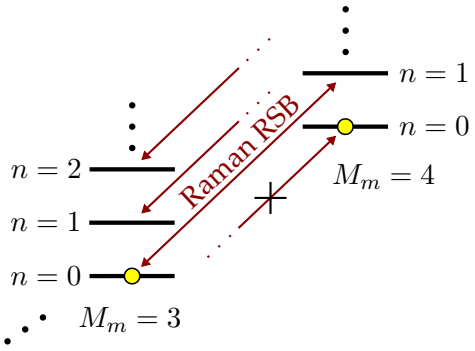

Figure 8.10: Building on Fig. 8.9, the spin superposition state $(|M_m = 4\rangle + |M_m = 3\rangle) \otimes |n = 0\rangle$ (yellow circles) is transformed into the motional superposition state $|M_m = 4\rangle \otimes (|n = 0\rangle + |n = 1\rangle)$ using a Raman red sideband π-pulse, which has no effect on $|M_m = 4\rangle \otimes |n = 0\rangle$. The states that are coupled instead are indicated by the red arrows.

magnetic field. These come from various sources within the building, including a nearby goods lift and a number of scanned superconducting magnets. The result were non-statistical features in the Ramsey fringe data, which made accurate fitting, and hence, the deduction of τ_C difficult.

It would have been possible to overcome this problem by preparing the motional superposition through a magnetic field insensitive spin superposition of the clock states $|M_m = 0\rangle$ and $|M_n = 1\rangle$. More specifically, the microwave transfer pulse sequence $\{A, B, C\}$ (see section 2.5) could be used to pass pupulation from $|M_m = 4\rangle$ to $|M_n = 1\rangle$, followed by a microwave $\pi/2$-pulse on the clock transition to create $(|M_m = 4\rangle + |M_n = 1\rangle) \otimes |n = 0\rangle$. A subsequent red sideband pulse on the same transition, which can in principle be performed using lasers or microwaves, would then create the desired motional superposition state $|M_n = 1\rangle \otimes (|n = 0\rangle + |n = 1\rangle)$.

At the time of writing this thesis, we are unable to implement the above experiment, since it is not possible to bridge the 3.2 GHz hyperfine splitting

with our current Raman laser setup, described in section 5.3.3. Likewise, the alternative approach of using microwaves is, at this stage, unfeasible due to the poor quality and long duration of the associated sideband pulses.

We can therefore only make a rough estimate of the motional coherence time, based on our previous measurement of the heating rate, described in section 8.1. If we assume (pessimistically) the worst case heating rate of $\dot{\bar{n}} \lesssim 5$ quanta/s for the gate mode (section 8.1), we could upper-bound the motional coherence time as $\tau_{\mathrm{C}} \leq 2/\dot{\bar{n}}$. This results from the general decay time relationship $T_2 \leq 2T_1$ for a two state quantum system [Fox06]. According to the simulations in section 4.3.7, this would give a gate error of 1.3×10^{-3}. The coherence time could of course be shorter than 400 ms, but presumably by not more than a factor of ~ 2 given our measured gate error of 3×10^{-3} (section 7.4).

8.4 Off-resonant excitation

We attempted to independently measure the off-resonant excitation caused by the gate microwaves, by diagnosing their impact on a single qubit's populations. To this end, we apply the DDMS gate sequence to a single ion, with all parameters set up as for an operation on the two-ion rocking mode. In this configuration, the drive field is approximately 30 kHz detuned from the nearest single-ion motional mode and therefore does not excite any motion, allowing us to observe its impact on the populations, primarily via off-resonant excitation. Additionally, this experiment captures errors that may arise from imperfections in the dynamical decoupling field or the π-pulses.

The measurement consists of four experiments, the pulse sequences of which are illustrated in Fig. 8.11. In the experiment (a), we measure the shelved population after preparing the qubit in $|1\rangle$ and executing the DDMS sequence, followed by a carrier π-pulse. This latter pulse, together with the refocussing Walsh pulse, induces the qubit population to be shelved, so that we ideally expect

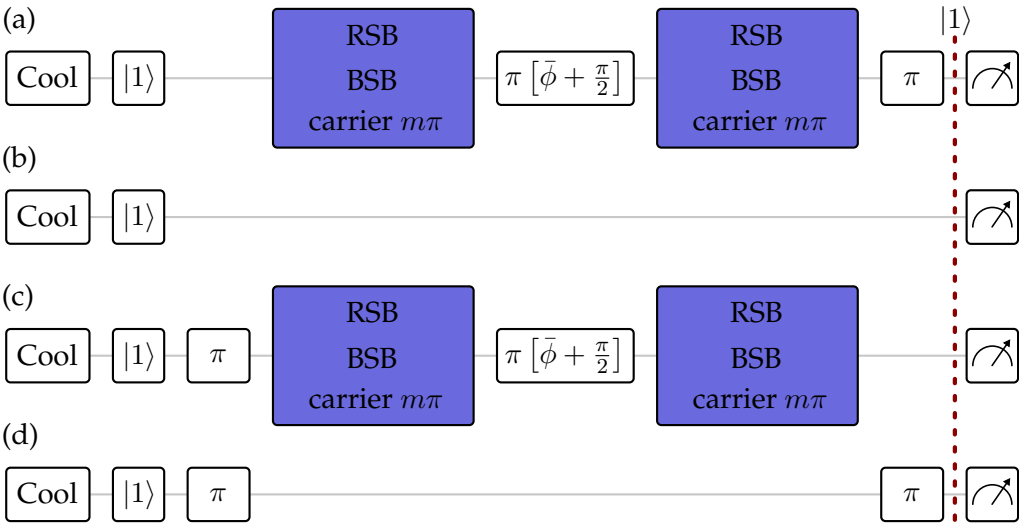

Figure 8.11: Pulse sequences used for measuring off-resonant excitation caused by the gate microwaves using a single ion in each case. In all four experiments we read out from $|1\rangle$ (red dotted line), to retain sensitivity to errors due to population escaping the qubit subspace.

to measure the P_1 readout level. However, spin flips and population escaping the quibt subspace due to off-resonant excitation, causes a detectable deviation from the readout level, which we measure independently using sequence (b).

In experiment (c) we conduct the same experiment as in (a), but with the carrier π-pulse moved in front of the DDMS sequence, so that the qubit starts in $|0\rangle$ and the qubit ends up shelved again. We stick to the shelf since we want the measurement to remain sensitive to errors due to population escaping the quibt subspace. Experiment (d) determines the corresponding readout level, using a π-pulse before and after the delay that replaces the gate sequence.

After $50\,000$ repeats of each of the four experiments, we measure the deviation of the population in (a) from the readout level (b) to be $\epsilon_1 = -3(4) \times 10^{-4}$, the deviation of the population in (c) from the readout level (d) to be $\epsilon_2 = -5(4) \times 10^{-4}$, and the average deviation to be $-4(3) \times 10^{-4}$. When transforming the measured values into a gate error following [Bal14]

$$\epsilon = \bar{\epsilon} - \frac{1}{2}\left(\epsilon_1^2 + \epsilon_2^2\right) + 2\bar{\epsilon}\left(1 - \bar{\epsilon}\right) \approx 3\bar{\epsilon}, \tag{8.1}$$

where $\bar{\epsilon} = \frac{1}{2}\left(\epsilon_1 + \epsilon_2\right)$, we obtain $\epsilon = -1.3(8) \times 10^{-3}$. The negative sign comes

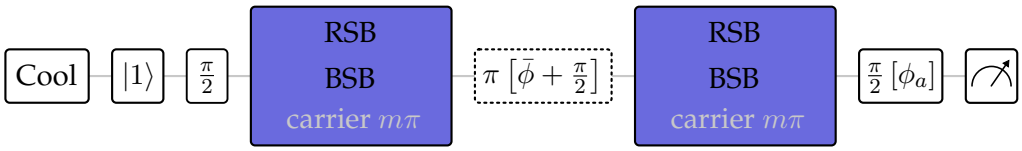

Figure 8.12: Ramsey experiment on a single ion, with the optimal DDMS gate sequence from section 7.2 inserted into the Ramsey delay, in order to estimate the accompanying AC-Zeeman shift fluctuations. We scan the phase ϕ_a of the second $\pi/2$-pulse to generate a fringe. Blue colouring indicates the presence of the AC-Zeeman shift, the constant part of which is taken into account in the carrier drive frequency. When both the carrier drive and the refocussing π-pulse are switched off, the sequence tests for AC-Zeeman shift fluctuation during the basic Mølmer-Sørensen gate.

from the fact that the readout level measurements (b) and (d) yielded poorer populations than those obtained from (a) and (c), indicating that the statistical sensitivity of this measurement was not sufficient to detect any effect from off-resonant excitation. We assume any gate error is therefore $\lesssim 10^{-3}$, i.e. comparable to the statistical sensitivity of these measurements.

8.5 Uncompensated AC Zeeman shift fluctuations

In order to probe the uncompensated AC-Zeeman shift fluctuations on the qubit, as well as to demonstrate the effectiveness of the DDMS scheme, we conduct a Ramsey experiment on a single ion. To this end, we embed our optimal gate sequence, described in section 7.2, between two $\pi/2$-pulses, while scanning the phase of the second one to generate a fringe, as illustrated in Fig. 8.12. Since all gate parameters are set up as for an operation on the two-ion rocking mode, the bichromatic drive field is approximately 30 kHz detuned from the nearest single-ion motional mode. Consequently, its only effect is to create the same AC-Zeeman shift (including fluctuations) as during the real gate operation, but without coupling to the ion's motion.

Fig. 8.13 illustrates the Ramsey fringes obtained from both the DDMS gate, as well as the basic Mølmer-Sørensen sequence, without the refocussing π-pulse or the dynamical decoupling carrier drive. We can calculate the r.m.s. deviation of

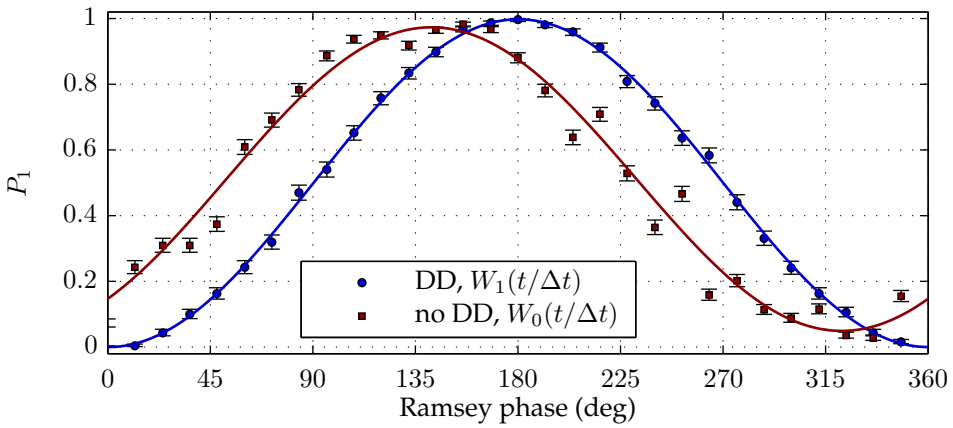

Figure 8.13: Ramsey experiment on a single ion, when using the sequence in Fig 8.12. For the blue data set (circles), we use our optimal DDMS gate sequence whereas for the red data set (squares) we repeat the measurement without Walsh modulation (the π pulse) and dynamical decoupling. All data have been fitted using maximum-likelihood estimation, the results of which are plotted as solid lines and give sinusoidal fringe contrasts of 0.998(5) and 0.924(6) after correcting for a SPAM error of 0.0026. We attribute the discrepancy between the two to AC-Zeeman shift fluctuations.

the AC-Zeeman shift fluctuations from the measured fringe contrast, following [Har14]

$$(2P_1^{\max} - 1) = e^{-\frac{1}{2}\left(\Delta_{\mathrm{Z,rms}}\Delta t_g\right)^2}, \tag{8.2}$$

assuming the variations are normally-distributed and slow (shot-to-shot). For the basic Mølmer-Sørensen scheme we measure a contrast of 0.924(6), which corresponds to $\Delta_{\mathrm{Z,rms}} = 2\pi \times 19.7(8)$ Hz r.m.s. According to the simulations described in section 4.3.5, this translates to a gate error of $5.6(5) \times 10^{-2}$. When applying the same fluctuation characteristics to simulations of a DDMS gate, as was done in section 4.4.2, the error is predicted to drop to $3.0(2) \times 10^{-5}$, which in practice would mean the effective elimination of the error. An important caveat here is that the above prediction assumes that there are no AC-Zeeman shift fluctuations during the gate pulses. The measurement at hand is not sensitive enough to verify this, as the measured Ramsey fringe contrast for the DDMS sequence is 0.998(5), which is consistent with one.

To improve on the statistical precision of the measurement as well as to get an

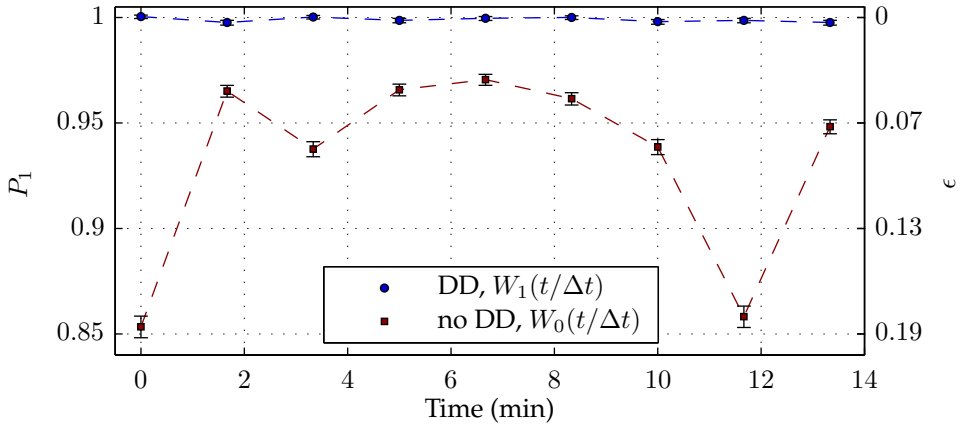

Figure 8.14: Ramsey experiment on a single ion, when enclosing the gate sequence in a pair of $\pi/2$-pulses, setting the phase of the second one to the peak of the fringe, and monitoring the fluctuations of P_1 over time. Similar to Fig. 8.13, the blue data set (circles) corresponds to our optimal DDMS gate sequence whereas the red data set (squares) is generated without Walsh modulation and dynamical decoupling. All data have been corrected for a 0.0034(3) SPAM error. The average of the blue data points gives $P_1 = 0.9994(4)$, which is consistent with the fringe contrast from Fig. 8.13. On the right-hand ordinate we indicate the simulated two-qubit gate error for the P_1 population given by the red points, that is, for the basic Mølmer-Sørensen scheme.

idea for long term stability of the AC-Zeeman shift fluctuations, we perform the same Ramsey experiment as before, but with the phase of the second $\pi/2$-pulse fixed to the peak of the fringe. We then monitor the evolution of P_1 over time, as shown in Fig. 8.14. For our optimal DDMS sequence, the averaged P_1 over 9×5000 repeats of the experiment, taking ~ 14 min, is 0.9994(4). We infer from this that the gate fields induce a loss in qubit coherence of $\lesssim 0.1\%$, which includes contributions from AC-Zeeman shift fluctuations, off-resonant excitation as well as imperfections in the dynamical decoupling field or the π-pulses. We assume that consequent errors in the two-qubit gate are at a similar level, $\lesssim 0.1\%$. For the gate sequence without the dynamical decoupling and Walsh modulation, the contrast varies between 0.71(1) and 0.941(5), where the associated gate errors are indicated in Fig. 8.14.

As a side note, these experiments imply that further developing the current trap to a version containing multiple trapping zones along the centre electrodes,

including the option to shuttle ions between them, might be worth considering. Such a system would allow for DDMS gates to be performed on selected pairs of neighbouring ions, while only imposing the above determined error of $\lesssim 0.1\%$ on single ions trapped in other zones, which can be used as e.g. memory qubits.

8.6 Systematic miscalibrations of the gate area

The leading systematic miscalibration in our gate is that of the gate time, as it is a laborious parameter to tune (see section 7.3). In order to estimate its deviation from the optimum in our best gate, given by $\Delta t_g = \pi\sqrt{K}/\Omega^{(2)}$ with $\Omega^{(2)}$ fixed, we fit the detuning scan shown in Fig. 7.12 using the analytic propagator from Eq. 4.62 and with only the Rabi frequency $\Omega^{(2)}$ floated. We then compare the fitted to the ideal value of $\Omega^{(2)}$, and convert the discrepancy to a gate time error, which turns out to be 0.07(1) ms on 3.25 ms. Using the theory presented in section 4.4.1, we calculate the corresponding gate error to be $2(2) \times 10^{-4}$. We have no reason to believe that Δt is changing over time, meaning that, in principle, this error contribution can be eliminated with more careful calibration.

Systematic deviations of the gate detuning δ are not expected to contribute a significant additional error, other than that due to fluctuations already treated in section 8.2. This is because during a data run, systematic shifts in δ can be read off directly from the populations after the gate, and straightforwardly corrected via the experimental control program. We check this regularly between experimental scans, as we are on constant lookout for drifts of the motional mode.

8.7 Microwave power stability

In the first instance, we estimate the gate error from Rabi frequency fluctuations and imbalance by conducting microwave power measurements using the microwave noise eater. To this end, we run the gate sequence 60×512 times while

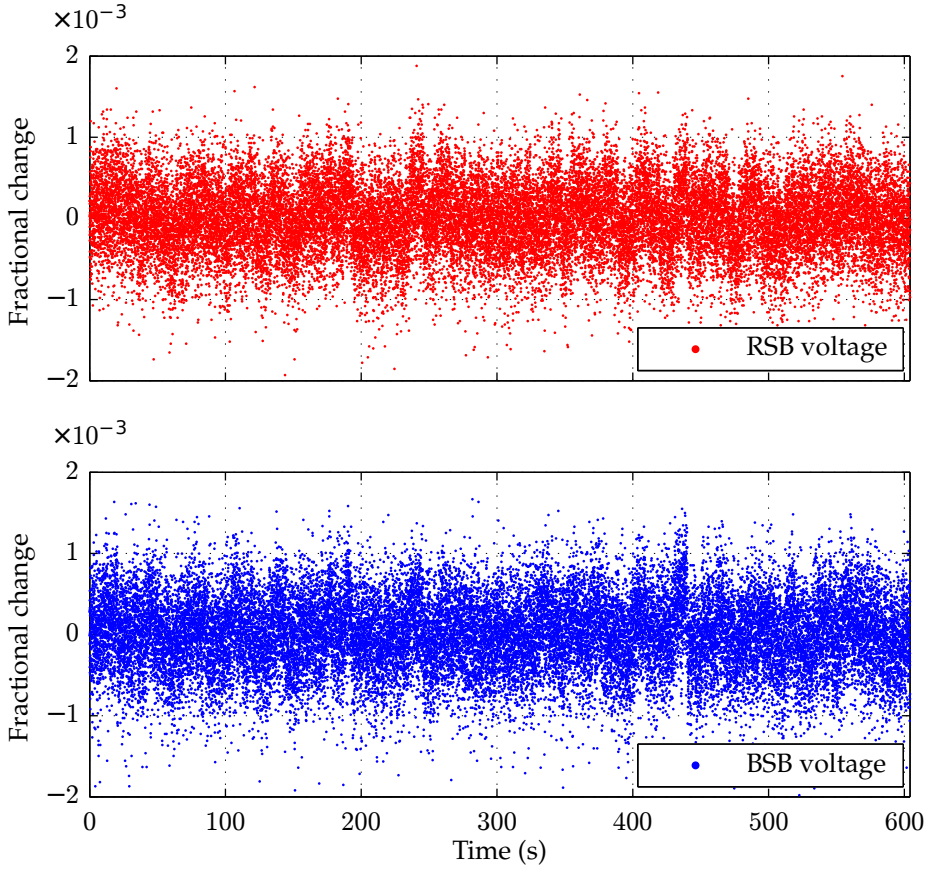

Figure 8.15: Fractional change of the red and blue microwave sideband's r.m.s. voltage, while running the gate sequence 60×512 times, each repetition taking 20 ms. Normalised r.m.s. deviations of 3.6×10^{-4} and 3.9×10^{-4} can be deduced for the red and blue sidebands, respectively.

recording the IC power detector's output voltage, sampled individually for each sideband during the two 120 µs noise eating periods after readout. A comparison of the data from both sidebands (Fig. 8.15) yields a power difference, which can be converted to a Rabi frequency imbalance. At the time we took our best gate data, we were able to infer a fractional difference between the sideband Rabi frequencies of $3.4(6) \times 10^{-3}$. According to section 4.3.3, this corresponds to a negligible gate error at the 10^{-7} level.

From the same data, we can also extract an estimate for the power stability. The fractional change of the detector's output is 3.6×10^{-4} and 3.9×10^{-4} for the red and blue sidebands, respectively, translating to a negligible gate error at

the 10^{-6} level when using the simulations from section 4.3.1. For completeness, we quantified the inherent noise from the detector by conducting the same measurement, but without any input power, resulting in an r.m.s. deviation of 1.8×10^{-4}. Although this represents a significant portion of the above measured fluctuations, taking them into account would only correct the already insignificant gate error further downwards.

Since the microwave power detector used in the above measurements receives its input signals from a directional coupler located before the diplexer, noise introduced after that pickoff point, that is, within the diplexer or the trap itself, is not captured by the detector and is therefore a minimum estimate. Thus, we use the previously measured AC-Zeeman shift fluctuations $\Delta_{\mathrm{Z,rms}}$ to get an estimate of the Rabi frequency fluctuations at the ion, as a consistency check. The AC-Zeeman shift ω_{Z} on the qubit transition due to off-resonant excitation of a spectator transition by the sideband microwaves, is proportional to Ω^2 and therefore to the microwave power P, where Ω is the Rabi frequency on that spectator transition. From section 8.5 we know that that $\Delta_{\mathrm{Z,rms}} = 2\pi \times 19.7(8)$ Hz r.m.s. for slow fluctuations, while $\omega_{\mathrm{Z}} = 2\pi \times 20.78$ kHz. We then deduce $\Delta_{\mathrm{Z,rms}}/\omega_{\mathrm{Z}} \propto P_{\mathrm{rms}}/P \approx 10^{-3}$, which when translated into Rabi frequency fluctuations yields $\sim 5 \times 10^{-4}$. According to simulations (section 4.3.1) this gives a gate error at the 10^{-6} level, which is consistent with our previous estimate.

8.8 Error summary

We summarise our measurement outcomes and associated gate error predictions from simulation in table 8.2. It is important to note that this list is not exhaustive and misses several errors that are expected to have a non-negligible contribution to the gate's error budget, but were not within our measurement capabilities (e.g. motional decoherence).

The three leading contributions are frequency fluctuations of the motional

error source	measured value	gate error (10^{-3})
motional mode frequency fluctuations	$\sim 30\,\mathrm{Hz}$ r.m.s.	~ 2
heating rate	$\lesssim 5\,\mathrm{quanta/s}$	$\lesssim 2$
gate fields	$\lesssim 0.1\%$ contrast loss	$\lesssim 1$
gate area miscalibration	$\lesssim 2 \times 10^{-2}$ fractional	$\lesssim 0.2$
Slow Rabi frequency fluctuations	$\sim 5 \times 10^{-4}$ fractional	$\lesssim 0.001$
Rabi frequency imbalance	$\sim 3 \times 10^{-3}$ fractional	$\lesssim 0.0001$
Measured total gate error		$3(1)$

Table 8.2: Summary of measured values and corresponding predicted DDMS gate errors, assuming our best gate parameters given in section 7.1. We separate the three leading errors from additional contributions using the horizontal line. The "gate fields" error source encompasses infidelities from the AC-Zeeman shift, off-resonant excitation as well as imperfections in the dynamical decoupling carrier drive or in the π-pulses used in the measurement. We note that this list is not complete and misses several error sources that are expected to have a non-negligible contribution to the error budget, but were not within our measurement capabilities (e.g. motional decoherence).

mode, motional heating and errors induced by the gate fields. In the case of the frequency fluctuations of the motional mode, the measurement precision was limited by the difficulty to pin down a drifting, stochastic quantity. In the case of the heating rate, the limitation was the demanding measurement process that required two ions to remain trapped for long periods of time without laser cooling. Lastly, the contribution from the gate fields encompasses several sources of error, and is therefore difficult to factor into the gate fidelity with good precision without additional information and a more detailed noise model.

9

Conclusions

9.1 Summary

The main result of this thesis is the experimental implementation and characterisation of a high-fidelity two-qubit quantum logic gate, driven by near-field microwaves in a room-temperature microfabricated surface ion trap. The gate is applied to two $^{43}Ca^{+}$ hyperfine clock qubits using the oscillating magnetic field gradient produced by a microwave electrode, integrated into the trap structure. We measure a fidelity of 99.7(1)%, after accounting for state preparation and measurement errors of 0.34(3)% per qubit. The gate scheme used is a variation of the standard Mølmer-Sørensen gate, with the addition of dynamical decoupling and Walsh modulation, which stabilise the qubits against fluctuations of AC-Zeeman shift and the motional mode frequency, respectively.

In previous work, using the same qubit states in the same apparatus, single-qubit performance beyond the fault-tolerant threshold has been demonstrated [HAB+14]. Combined with the results from this work, we have therefore implemented all elementary qubit operations required for quantum computing in one trap. We have not demonstrated individual qubit addressing, but we note that it was already demonstrated using microwave methods in a similar trap, using $^{25}Mg^{+}$ qubits, and we would expect similar performance in our apparatus. Our results are summarised in table 9.1 and published in [HSA+16].

operation	error	reference
memory ($[\Delta t_g = 3.25\,\mathrm{ms}]$ / $[T_2^* = 50\,\mathrm{s}]$)	7×10^{-5}	[HAB+14]
state preparation	2×10^{-4}	[HAB+14]
global single-qubit gate (randomised benchmarking)	1×10^{-6}	[HAB+14]
single-shot readout (single qubit)	5×10^{-4}	[HAB+14]
single-shot readout (per qubit for two qubits)	3×10^{-3}	this work
two-qubit gate (reduced process tomography)	3×10^{-3}	this work

Table 9.1: Summary of errors in elementary qubit operations achieved in our present experimental setup using the "atomic clock" qubit in $^{43}Ca^+$ at 146 G. The readout error for two qubits could be reduced to the 5×10^{-4} level measured in [HAB+14] by using spatially-resolved fluorescence detection [BSWL10]. This should work well due to the relatively large axial ion separation (8 μm).

We have also experimentally demonstrated a simple and robust method for Doppler cooling $^{43}Ca^+$, by making use of a two-photon dark resonance. Despite its complex level structure at 146 G, this method does not require any more laser frequencies than are needed at low magnetic field. The lowest temperature reached was 0.3 mK, which is slightly below the Doppler limit for the corresponding two-level system. The cooling parameters we found are readily adaptable to allow the fluorescence of the ion to be raised to a similar level to that typically observed for $^{40}Ca^+$, or ions with metastable D states, in order to enable qubit readout with high precision. We have further shown that the application of a resolved sideband cooling step after dark resonance cooling allows preparation of both radial motional modes sufficiently close to the ground state for the implementation of high-fidelity two-qubit gates.

9.2 Comparison with other results

Trapped ion quantum computing experiments using near-field microwaves are a relatively young field of research compared with efforts based on lasers. The reason mainly comes from the necessity for microfabricated surface ion traps, which only recently became available and reliable.

We plot the history of two-qubit gate errors achieved in trapped ion systems

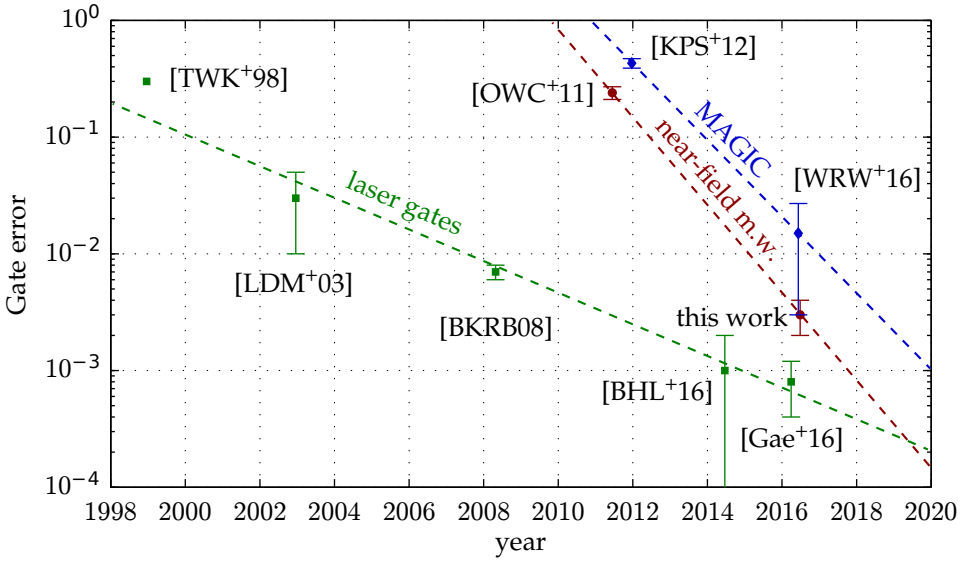

Figure 9.1: Progression of two-qubit gate errors in trapped-ion systems using lasers (green), the MAGIC scheme involving a static magnetic field gradient (blue, see section 3.2.3), and near-field microwaves (red). The results have been dated by means of their corresponding arxiv preprint appearance date.

in Fig. 9.1. Our work represents the highest two-qubit gate fidelity attained using microwave technology, but the error we measure is still about a factor of four worse than that of the best laser gate, which is $8(4) \times 10^{-4}$ [Gae16]. We note that an implementation of a two-qubit laser gate in a surface trap has recently been published in [HFM+16] and yielded a gate error of 16(2)%. In unpublished work [Mau16], gate set tomography of two-qubit laser gates in a surface trap has yielded a fidelity of $\approx 99.58(6)\%$.

9.3 Outlook

The three leading errors in our gate are believed to be frequency fluctuations of the motional mode, motional heating, and errors introduced by the gate fields. All three can be reduced in future experiments by means of purely technical improvements.

An important step is to increase the gate's speed, which decreases its sensitiv-

ity to heating and motional mode frequency fluctuations. Faster gates require larger magnetic field gradients, which can be obtained by moving the ions closer to the trap electrodes that generate the field, or by using a higher microwave power P. The field gradient scales as $\sqrt{P}$ and as d^{-2}, where d is the ion to trap surface separation. However, three caveats have to be noted here.

First the magnetic field strength increases as d^{-1} and $\sqrt{P}$, which worsens the off-resonant excitation error. Second, we are likely to run into problems with laser scatter from the blue cooling and detection beams for too small values of d, which may lead to charging effects on the trap chip or to an unacceptably high background signal during fluorescence detection. Third, and most importantly, the heating rate is expected to scale as d^{-4} [TKK+00]. Thus, a reduction in the ion to trap surface separation will, in practice, have to be accompanied by further technical measures to mitigate the impact of the above issues.

As an example, heating rates can be decreased using surface cleaning techniques with laser pulses [AGH+11] or argon ion bombardment [HCW+12]. Cryogenic operation has also been shown to reduce electric field noise from the trap surface [LGA+08] and comes with the additional benefit of reducing the electrical resistance of the current-carrying electrode. This allows for the application of higher microwave powers that might otherwise damage the trap chip due to excessive heat dissipation.

Off-resonant excitation and A.C. Zeeman shifts, and their fluctuations, can be reduced by nulling the microwave field at the position of the ions. This can either be done by carefully balancing currents in multiple nearby electrodes [OWC+11, WOC+13a], or with the help of improved trap geometries, such as the meander-electrode design presented in [CKDO14] and investigated in [WHZ+16]. The dynamical decoupling method demonstrated in this work should prove to be particularly effective when used in combination with these more advanced trap designs, as it significantly reduces the level of field suppression that must be

achieved.

Finally, with sufficient microwave field gradients, one may reduce some of the laser complexity by sideband cooling the ions using microwaves [OWC+11, WRW+15]. The problem of scattered light may be overcome by appropriately timed trap parameter changes, moving the ions closer to the surface for the gate operation, and further away from it for cooling and readout (e.g. by using two connected trap zones with different electrode widths).

At the time of writing this thesis, our group has already begun work on the next iteration of this experiment, which will employ liquid helium cooling, argon ion surface cleaning and a meander trap design with an ion height of 50 μm. We will use the $^{43}Ca^{+}$ clock transition at 288 G, which is π-polarised and therefore couples more efficiently to the magnetic field gradients produced by the trap. The hope is that these improvements will allow us to reach two-qubit gate errors at the 10^{-4} level. A similar apparatus is already operational at the NIST Ion Storage Group, using $^{25}Mg^{+}$ qubits.

Our group is in parallel working on spatial addressing with near-field microwaves. Recent results [ACLS+16] imply an addressing error of 1.3×10^{-6}, measured using $^{43}Ca^{+}$ at a field of 2.8 G in a multi-zone surface trap. The aim is to, in future, combine both experiments into one microwave QCCD prototype, paving the way for scalable devices based on microwave-driven quantum logic.

Appendices

A

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

In chapter 2.2 we have seen that the $|M_n\rangle$ and $|M_m\rangle$ states in the $4S_{1/2}$ ground level of $^{43}Ca^{+}$, when subjected to a static magnetic field $\mathbf{B}_{\mathrm{stat}}$ of arbitrary strength, can be expressed as linear combinations of at most two high field states $|I, J, M_I, M_J\rangle$ in the following way

$$\begin{aligned} |M_n\rangle &= \alpha_n \,|I, J, M_n + 1/2, -1/2\rangle + \beta_n \,|I, J, M_n - 1/2, 1/2\rangle \,, \\ |M_m\rangle &= \alpha_m \,|I, J, M_m + 1/2, -1/2\rangle + \beta_m \,|I, J, M_m - 1/2, 1/2\rangle \,, \end{aligned} \tag{A.1}$$

where $|M_n\rangle$ are the states originating from the low field $|I, J, F = 3, M_F\rangle$ states, and $|M_m\rangle$ are the states originating from the low field $|I, J, F = 4, M_F\rangle$ states, when ramping the static magnetic field up into the intermediate field regime. The expansion (Clebsch-Gordan) coefficients

$$\begin{aligned} \alpha_n &= \langle I, J, M_I = M_n + 1/2, M_J = -1/2 \,|I, J, F = 3, M_F = M_n\rangle \,, \\ \alpha_m &= \langle I, J, M_I = M_m + 1/2, M_J = -1/2|I, J, F = 4, M_F = M_m\rangle \,, \\ \beta_n &= \langle I, J, M_I = M_n - 1/2, M_J = +1/2 \,|I, J, F = 3, M_F = M_n\rangle \,, \\ \beta_m &= \langle I, J, M_I = M_m - 1/2, M_J = +1/2|I, J, F = 4, M_F = M_m\rangle \,, \end{aligned} \tag{A.2}$$

are given in tables A.1 and A.2 for magnetic fields of 146.0942 G and 287.7827 G, respectively.

Based on these values we can determine the magnetic dipole transition matrix elements for σ_+ and σ_- polarized transitions with $M_n - M_m = \pm 1$

$$\begin{aligned} \langle M_n|\mu_{\pm 1}|M_m\rangle = \pm \frac{1}{\sqrt{2}}\Big\{ &\mu_B g_J \alpha_n \beta_m \\ &\mp \mu_N g_I \left(\beta_m \beta_n \sqrt{16 - M_n^2} + \alpha_m \alpha_n \sqrt{15 - 2M_n - M_n^2}\right) \Big\}, \end{aligned} \tag{A.3}$$

and for π polarized transitions with $M_n - M_m = 0$

$$\begin{aligned} \langle M_n|\mu_0|M_m\rangle = &\frac{1}{2}\mu_B g_J \left(\alpha_n \alpha_m - \beta_n \beta_m\right) \\ &+ \mu_N g_I \left\{\alpha_n \alpha_m (M_n + 1/2) - \beta_n \beta_m (M_n - 1/2)\right\}, \end{aligned} \tag{A.4}$$

the numeric values of which are listed in tables A.3 and A.4 for the two magnetic fields of interest.

M_n	α_n	β_n
-3	0.3175	-0.9483
-2	0.4548	-0.8906
-1	0.5645	-0.8255
0	0.6611	-0.7503
+1	**0.7504**	**-0.6610**
+2	0.8354	-0.5497
+3	0.9181	-0.3964

M_m	α_m	β_m
-4	1	0
-3	0.9483	0.3175
-2	0.8906	0.4548
-1	0.8255	0.5645
0	**0.7503**	**0.6611**
+1	0.6610	0.7504
+2	0.5497	0.8354
+3	0.3964	0.9181
+4	0	1

Table A.1: Expansion coefficients to be inserted into Eq. A.1 for a valid representation of $|M_n\rangle$ and $|M_m\rangle$ states in the $4S_{1/2}$ ground level of $^{43}Ca^{+}$ at 146.0942 G. The states involved in the clock transition at this field are highlighted in bold.

M_n	α_n	β_n
-3	0.2880	-0.9576
-2	0.4155	-0.9096
-1	0.5202	-0.8541
0	0.6154	-0.7882
+1	**0.7071**	**-0.7071**
+2	0.7992	-0.6011
+3	0.8953	-0.4455

M_m	α_m	β_m
-4	1	0
-3	0.9576	0.2880
-2	0.9096	0.4155
-1	0.8541	0.5202
0	0.7882	0.6154
+1	**0.7071**	**0.7071**
+2	0.6011	0.7992
+3	0.4455	0.8953
+4	0	1

Table A.2: Expansion coefficients to be inserted into Eq. A.1 for a valid representation of $|M_n\rangle$ and $|M_m\rangle$ states in the $4S_{1/2}$ ground level of $^{43}Ca^{+}$ at 287.7827 G. The states involved in the clock transition at this field are highlighted in bold.

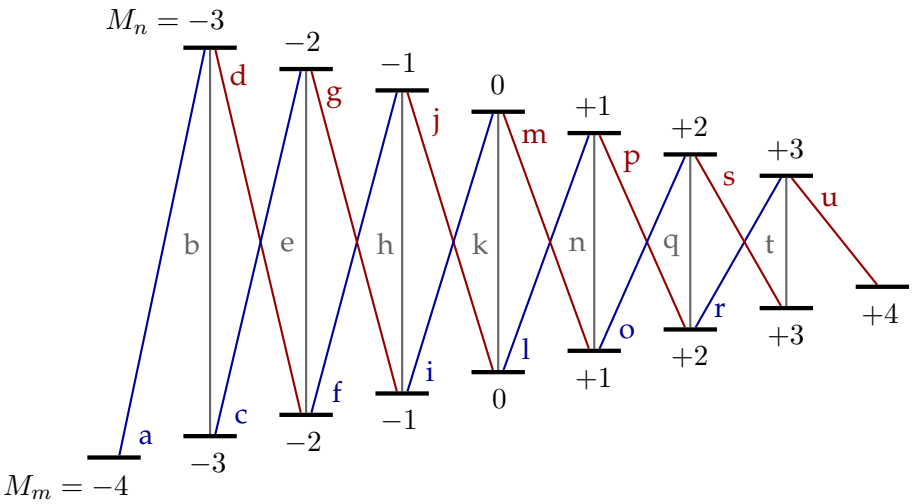

Figure A.1: Labelling of magnetic dipole transitions in the ground state of $^{43}Ca^{+}$.

	M_m	M_n	Frequency (GHz)	$\mathrm{d}f/\mathrm{d}B_{\mathrm{stat}}$ [$\mathrm{d}^2f/\mathrm{d}B_{\mathrm{stat}}^2$] (MHz/G)	$\langle M_n\|\mu_Q\|M_m\rangle$ (μ_B)
a	-4	-3	3.589 033 180	2.519 894	+1.342 420
b	-3	-3	3.543 000 725	2.237 093	+0.602 802
c	-3	-2	3.495 825 733	1.940 390	+1.195 577
d	-2	-3	3.495 742 031	1.939 817	-0.204 417
e	-2	-2	3.448 567 039	1.643 113	-0.810 866
f	-2	-1	3.400 062 158	1.330 090	-1.040 756
g	-1	-2	3.399 978 455	1.329 517	-0.363 393
h	-1	-1	3.351 473 573	1.016 493	-0.932 839
i	-1	0	3.301 519 669	0.684 932	-0.876 791
j	0	-1	3.301 435 966	0.684 359	-0.528 271
k	0	0	3.251 482 062	0.352 797	-0.993 060
l	**0**	**+1**	**3.199 941 077**	**0 [2.416 kHz/G^2]**	**-0.702 127**
m	+1	+0	3.199 857 374	-0.000 573	+0.702 254
n	+1	+1	3.148 316 389	-0.353 370	+0.993 034
o	+1	+2	3.095 026 842	-0.730 728	-0.514 410
p	+2	+1	3.094 943 139	-0.731 301	+0.887 366
q	+2	+2	3.041 653 592	-1.108 659	-0.919 344
r	+2	+3	2.986 424 605	-1.514 736	+0.308 504
s	+3	+2	2.986 340 903	-1.515 309	-1.085 679
t	+3	+3	2.931 111 915	-1.921 385	+0.728 641
u	+4	+3	2.873 631 427	-2.362 039	+1.299 654

Table A.3: Properties of all magnetic dipole transitions in the $4S_{\frac{1}{2}}$ ground-level of $^{43}Ca^+$ at 146.094 2 G, labelled according to Fig. A.1. The clock transition is shown in bold.

	M_m	M_n	Frequency (GHz)	$\mathrm{d}f/\mathrm{d}B_{\mathrm{stat}}$ [$\mathrm{d}^2f/\mathrm{d}B_{\mathrm{stat}}^2$] (MHz/G)	$\langle M_n\|\mu_Q\|M_m\rangle$ (μ_B)
a	-4	-3	3.949 771 857	2.569 912	+1.355 692
b	-3	-3	3.867 368 475	2.337 131	+0.552 243
c	-3	-2	3.781 465 945	2.086 061	+1.233 104
d	-2	-3	3.781 301 064	2.085 488	-0.169 444
e	-2	-2	3.695 398 534	1.834 419	-0.756 705
f	-2	-1	3.605 294 276	1.560 347	-1.099 756
g	-1	-2	3.605 129 395	1.559 774	+0.305 995
h	-1	-1	3.515 025 138	1.285 701	+0.889 435
i	-1	0	3.420 035 468	0.982 944	+0.953 011
j	0	-1	3.419 870 588	0.982 371	+0.453 176
k	0	0	3.324 880 918	0.679 614	+0.971 138
l	0	+1	3.224 113 796	0.340 094	+0.789 034
m	+1	0	3.223 948 915	0.339 521	+0.616 028
n	**+1**	**+1**	**3.123 181 794**	**0 [2.514 kHz/G^2]**	**+1.001 026**
o	+1	+2	3.015 435 346	-0.388 298	+0.601 736
p	+2	+1	3.015 270 465	-0.388 871	-0.799 970
q	+2	+2	2.907 524 018	-0.777 169	+0.961 755
r	+2	+3	2.791 111 000	-1.233 169	-0.379 132
s	+3	+2	2.790 946 119	-1.233 742	-1.012 843
t	+3	+3	2.674 533 102	-1.689 743	+0.798 542
u	+4	+3	2.546 787 218	-2.246 218	-1.267 363

Table A.4: Properties of all magnetic dipole transitions in the $4S_{\frac{1}{2}}$ ground-level of $^{43}Ca^{+}$ at 287.782 7 G, labelled according to Fig. A.1. The clock transition is shown in bold.

B

Magnus expansion

Consider a system described by a Hamitonian

$$H(t) = H_0 + H_I(t), \tag{B.1}$$

which consists of a time independent part H_0 and a time dependent part $H_I(t)$. In order to study the time evolution of the system under the action of $H_I(t)$, it is convenient to first discard the effect of H_0 by stepping into the interaction picture by following the transformation rules

$$\begin{aligned} \left|\tilde{\psi}(t)\right\rangle &= U_0^\dagger \left|\psi(t)\right\rangle, \\ \tilde{H}_I(t) &= U_0^\dagger H_I(t) U_0, \end{aligned} \tag{B.2}$$

where $U_0 = \exp(-\frac{i}{\hbar}H_0 t)\left|\psi(t)\right\rangle$. The corresponding Schrödinger equation in the interaction picure becomes

$$i\hbar\frac{d}{dt}\left|\tilde{\psi}(t)\right\rangle = \tilde{H}_I(t)\left|\tilde{\psi}(t)\right\rangle, \tag{B.3}$$

and the time evolution of the state vector of the system

$$\left|\tilde{\psi}(t)\right\rangle = \tilde{U}(t,0)\left|\tilde{\psi}(0)\right\rangle \tag{B.4}$$

can be determined by finding the propagator $\tilde{U}(t,0)$. The differential equation describing the dynamics of the propagator is

$$\frac{d}{dt}\tilde{U}(t,0) = -\frac{i}{\hbar}\tilde{H}_I(t)\tilde{U}(t,0). \tag{B.5}$$

As a result of the time dependence of $\tilde{H}_I(t)$, one faces the problem that, in general, $\tilde{U}(t,0) \neq \exp(-\frac{i}{\hbar}\tilde{H}_I t)$. Instead, the formal solution to Eq. B.5 is a Dyson series

$$\begin{aligned} \tilde{U}(t,0) &= \mathcal{T}\{e^{-\frac{i}{\hbar}\int_0^t \tilde{H}_I(\tau)d\tau}\}\,\tilde{U}(0,0) \\ &= \mathbb{1} + \sum_{n=1}^{\infty}\left(-\frac{i}{\hbar}\right)^n \int_0^t dt_1 \ldots \int_0^{t_{n-1}} dt_n\, H_I(t_1)\ldots H_I(t_n), \qquad \forall\, t_n < t_{n+1} \end{aligned} \tag{B.6}$$

where $\tilde{U}(0,0) = \mathbb{1}$, and where $\mathcal{T}$ is the time ordering operator, which ensures that the product of operators in the series is ordered so that operators acting at earlier times are left of operators acting at later times. The complete series described in Eq. B.6 gives rise to a unitary propagator, but in practice the series has to be truncated and the resulting approximated propagator is not unitary any more.

An alternative series expansion which overcomes this problem is the Magnus expansion. When truncated, it gives an approximation to $\tilde{U}(0,t)$ that stays unitary. In some cases, but most importantly in the case of a forced harmonic oscillator, the series terminates, allowing for the calculation of the exact propagator.

The Magnus expansion relies on the assumption that there exists an exponential solution to Eq. B.5 of the form

$$\tilde{U}(t,0) = \exp\Big(-\frac{i}{\hbar}\sum_{k=1}^{\infty}\Omega_k(t)\Big)\,\tilde{U}(0,0), \tag{B.7}$$

where $\tilde{U}(0,0) = \mathbb{1}$,and where the $\Omega_k(t)$ are given by

$$\begin{aligned}
\Omega_1(t) &= \int_0^t dt_1\, \tilde{H}_I(t_1),\\
\Omega_2(t) &= \frac{1}{2}\int_0^t dt_1 \int_0^{t_1} dt_2\, [\tilde{H}_I(t_1), \tilde{H}_I(t_2)],\\
\Omega_3(t) &= \frac{1}{6}\int_0^t dt_1 \int_0^{t_2} dt_2 \int_0^{t_2} dt_3\, \big\{[\tilde{H}_I(t_1), [\tilde{H}_I(t_2), \tilde{H}_I(t_3)]] + [\tilde{H}_I(t_3), [\tilde{H}_I(t_2), \tilde{H}_I(t_1)]]\big\},\\
&\vdots
\end{aligned} \tag{B.8}$$

In the special case of the system being a forced harmonic oscillator, the exact propagator can be found using the Magnus expansion, as all Ω_k above $k = 2$ vanish.

C

Two-ion radial rocking mode heating rate

We present an attempt to measure the heating rate of the two-ion radial rocking mode. The three plots show the one-ion shelved signal (*) and the 2-ion shelved signal (o) for a small scan range of the Raman probe detuning over the rocking mode red sideband (red) and blue sideband (blue). The sidebands are centred at approximately zero detuning. The nearby radial COM mode is offset from the rocking mode by approximately 30kHz. The curves show an approximate numerical simulation.

Figure C.1 : Zero heating delay. The simulation with initial mode temperatures of $\bar{n}_{\mathrm{COM}} = 0$ and $\bar{n}_{\mathrm{rock}} = 0.15$ matches the data reasonably well.

Figure C.2 : 80 ms heating delay. For the simulation, the COM mode was given its independently-measured heating rate of 60 quanta/s, and the rocking mode was given zero heating rate. The simulated red sideband curves match the data reasonably well.

Figure C.3 : Data as in Figure C.2. For the simulation, a heating rate of 60 quanta/s was again assumed for the COM mode, but here the rocking mode was given a heating rate of 5 quanta/s. The red sideband one-ion simulation is now significantly above the data, indicating that the heating rate is likely to be less than this value.

The data is unsatisfactory in several ways – the COM mode initial temperature was not measured, there is a non-negligible baseline due to photon scattering, and the rocking mode is so close to the COM mode that it is difficult to separate the effects of heating of the two modes. However, from examination of this and several similar data sets, we believe that 5 quanta/s is a reasonable upper limit for the rocking mode heating rate.

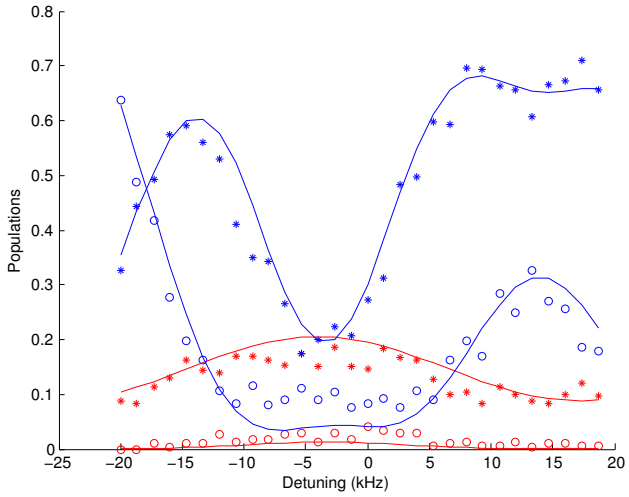

Figure C.1

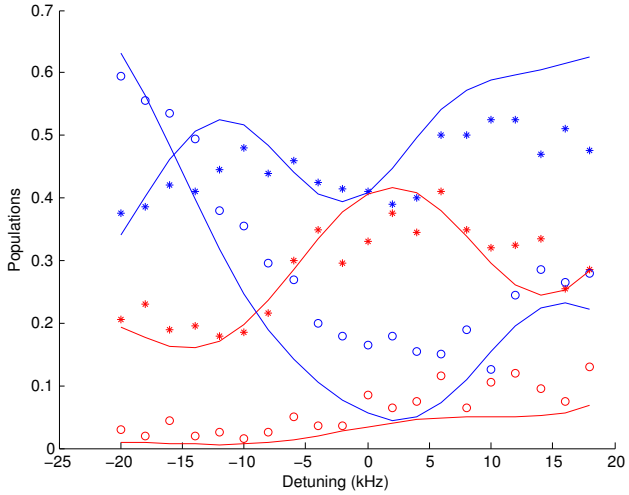

Figure C.2

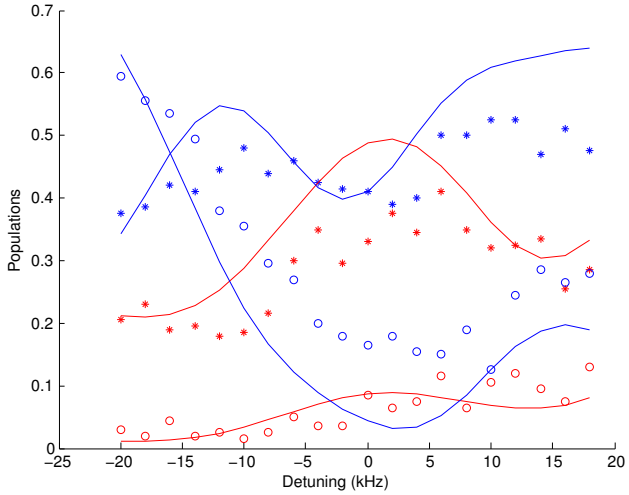

Figure C.3

Bibliography

[ACLH+14] D. P. L. Aude Craik, N. M. Linke, T. P. Harty, C. J. Ballance, D. M. Lucas, A. M. Steane, and D. T. C. Allcock. Microwave control electrodes for scalable, parallel, single-qubit operations in a surface-electrode ion trap. *Applied Physics B*, 114(1):3–10, 2014.

[ACLS+16] D. P. L. Aude Craik, N. M. Linke, M. A. Sepiol, T. P. Harty, C. J. Ballance, D. N. Stacey, A. M. Steane, D. M. Lucas, and D. T. C. Allcock. High-fidelity spatial addressing of Ca-43 qubits using near-field microwave control. *ArXiv e-prints*, Jan 2016.

[AGH+11] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Reduction of heating rate in a microfabricated ion trap by pulsed-laser cleaning. *New Journal of Physics*, 13(12):123023, 2011.

[AHB+13] 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(4), 2013.

[AHS+16] D. T. C. Allcock, T. P. Harty, M. A. Sepiol, H. A. Janacek, C. J. Ballance, A. M. Steane, D. M. Lucas, and D. N. Stacey. Dark-resonance Doppler cooling and high fluorescence in trapped Ca-43 ions at intermediate magnetic field. *New Journal of Physics*, 18(2):023043, 2016.

[All11] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. PhD thesis, University of Oxford, 2011.

[ASS+10] 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(5):053026, 2010.

[AUW+10] 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(3):033031, Mar 2010.

[Bal14] C. J. Ballance. *High-fidelity quantum logic in Ca^+*. PhD thesis, University of Oxford, 2014.

[BBC+95] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, Nov 1995.

[BDS+03] M. D. Barrett, B. DeMarco, T. Schaetz, V. Meyer, D. Leibfried, J. Britton, J. Chiaverini, W. M. Itano, B. Jelenković, J. D. Jost, C. Langer, T. Rosenband, and D. J. Wineland. Sympathetic cooling of $^{9}\mathrm{be}^{+}$ and $^{24}\mathrm{mg}^{+}$ for quantum logic. *Phys. Rev. A*, 68:042302, Oct 2003.

[Ber84] M. V. Berry. Quantal Phase Factors Accompanying Adiabatic Changes. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 392:45–57, 1984.

[BHL+05] K.-A. Brickman, P. C. Haljan, P. J. Lee, M. Acton, L. Deslauriers, and C. Monroe. Implementation of Grover's quantum search algorithm in a scalable system. *Phys. Rev. A*, 72:050306, Nov 2005.

[BHL+16] 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. *Phys. Rev. Lett.*, 117:060504, Aug 2016.

[BKRB08] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4:463–466, Jun 2008.

[BKRB15] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. Ion-trap measurements of electric-field noise near surfaces. *Rev. Mod. Phys.*, 87:1419–1482, Dec 2015.

[Bla01] E. D. Black. An introduction to PoundDreverHall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, 2001.

[BSPR12] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker. Robust trapped-ion quantum logic gates by continuous dynamical decoupling. *Phys. Rev. A*, 85:040302, Apr 2012.

[BSWL10] A. H. Burrell, D. J. Szwer, S. C. Webster, and D. M. Lucas. Scalable simultaneous multiqubit readout with 99.99% single-shot fidelity. *Phys. Rev. A*, 81:040302, Apr 2010.

[Bur10] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, University of Oxford, 2010.

[CBB+05] 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 Info. Comput.*, 5(6):419–439, September 2005.

[CGK98] I. L. Chuang, N. Gershenfeld, and M. Kubinec. Experimental implementation of fast quantum searching. *Phys. Rev. Lett.*, 80:3408–3411, Apr 1998.

[CKDO14] M. Carsjens, M. Kohnen, T. Dubielzig, and C. Ospelkaus. Surface-electrode Paul trap with optimized near-field microwave control. *Applied Physics B*, 114(1):243–250, 2014.

[CZ95] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Phys. Rev. Lett.*, 74:4091–4094, May 1995.

[CZ00] J. I. Cirac and P. Zoller. A scalable quantum computer with ions in an array of microtraps. *Nature*, 404:579–581, Apr 2000.

[DBE95] D. Deutsch, A. Barenco, and A. Ekert. Universality in Quantum Computation. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 449:669–677, 1995.

[DCG+09] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460:240–244, Jul 2009.

[Deh75] H. Dehmelt. Proposed $10^{14}\delta\nu > \nu$ laser fluorescence spectroscopy of Tl^+ mono-ion oscillator II. *Bulletin of the American Physical Society*, 20(1):60, 1975.

[Deu85] D. Deutsch. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 400(1818):97–117, 1985.

[DiV95] D.P. DiVincenzo. Two-bit gates are universal for quantum computation. *Phys. Rev. A*, 51:1015–1022, Feb 1995.

[DiV00] D. P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.

[EAZ05] J. Emerson, R. Alicki, and K. Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347, 2005.

[EF02] V. L. Ermakov and B. M. Fung. Experimental realization of a continuous version of the Grover algorithm. *Phys. Rev. A*, 66:042310, Oct 2002.

[Fey82] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.

[FMMC12] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, Sep 2012.

[Fox06] M. Fox. *Quantum Optics: An Introduction (Oxford Master Series in Physics, 6)*. Oxford University Press, USA, 2006.

[Gae16] Gaebler, J. P. and Tan, T. R. and Lin, Y. and Wan, Y. and Bowler, R. and Keith, A. C. and Glancy, S. and Coakley, K. and Knill, E. and Leibfried, D. and Wineland, D. J. High-fidelity universal gate set for $^9Be^+$ ion qubits. *Phys. Rev. Lett.*, 117:060505, Aug 2016.

[GAN14] I. M. Georgescu, S. Ashhab, and F. Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153–185, Mar 2014.

[GMT+12] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland. Randomized Benchmarking of Multiqubit Gates. *Phys. Rev. Lett.*, 108:260503, Jun 2012.

[GRB+01] 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.

[Gro96] L. K. Grover. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the* 28^{th} *Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, New York, NY, USA, 1996. ACM.

[Gro97] L. K. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Phys. Rev. Lett.*, 79:325–328, Jul 1997.

[GRZC03] J. J. García-Ripoll, P. Zoller, and J. I. Cirac. Speed Optimized Two-Qubit Gates with Laser Coherent Control Techniques for Ion Trap Quantum Computing. *Phys. Rev. Lett.*, 91:157901, Oct 2003.

[GWM+08] M. Gilowski, T. Wendrich, T. Müller, Ch. Jentsch, W. Ertmer, E. M. Rasel, and W. P. Schleich. Gauss Sum Factorization with Cold Atoms. *Phys. Rev. Lett.*, 100:030201, Jan 2008.

[HAB+14] 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. *Phys. Rev. Lett.*, 113:220501, Nov 2014.

[Har14] T. P. Harty. *High-fidelity microwave driven quantum logic in intermediate-field* $^{43}Ca^{+}$. PhD thesis, University of Oxford, 2014.

[HCD+12] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe. Coherent Error Suppression in Multiqubit Entangling Gates. *Phys. Rev. Lett.*, 109:020503, Jul 2012.

[HCW+12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-Fold Reduction of Electric-Field Noise in an Ion Trap Cleaned with *In Situ* Argon-Ion-Beam Bombardment. *Phys. Rev. Lett.*, 109:103001, Sep 2012.

[HFM+16] C. D. Herold, S. D. Fallek, J. T. Merrill, A. M. Meier, K. R. Brown, C. E. Volin, and J. M. Amini. Universal control of ion qubits in a scalable microfabricated planar trap. *New Journal of Physics*, 18(2):023048, 2016.

[HKVB11] D. Hayes, K. Khodjasteh, L. Viola, and M. J. Biercuk. Reducing sequencing complexity in dynamical quantum error suppression by Walsh modulation. *Phys. Rev. A*, 84:062323, Dec 2011.

[HSA+16] 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. *Phys. Rev. Lett.*, 117:140501, Sep 2016.

[JAG+11] M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose. Quantum annealing with manufactured spins. *Nature*, 473:194–198, May 2011.

[Jam98] D.F.V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B*, 66(2):181–190, 1998.

[Jan15] H. A. Janacek. *Optical Bloch equations for simulating trapped-ion qubits*. PhD thesis, University of Oxford, 2015.

[JBT+09] M. Johanning, A. Braun, N. Timoney, V. Elman, W. Neuhauser, and Chr. Wunderlich. Individual Addressing of Trapped Ions and Coupling of Motional and Spin States Using rf Radiation. *Phys. Rev. Lett.*, 102:073004, Feb 2009.

[JC63] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, Jan 1963.

[JMH98] J. A. Jones, M. Mosca, and R. H. Hansen. Implementation of a quantum search algorithm on a quantum computer. *Nature*, 393:344–346, May 1998.

[JPK00] D. Jonathan, M. B. Plenio, and P. L. Knight. Fast quantum gates for cold trapped ions. *Phys. Rev. A*, 62:042307, Sep 2000.

[KAN+14] S. Kotler, N. Akerman, N. Navon, Y. Glickman, and R. Ozeri. Measurement of the magnetic interaction between two bound electrons of two separate ions. *Nature*, 510:376–380, Jun 2014.

[Kei07] B. Keitch. *A Quantum Memory Qubit in Calcium-43*. PhD thesis, University of Oxford, 2007.

[KLR+08] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Phys. Rev. A*, 77:012307, Jan 2008.

[KMSW00] P. G. Kwiat, J. R. Mitchell, P. D. D. Schwindt, and A. G. White. Grover's search algorithm: An optical approach. *Journal of Modern Optics*, 47(2-3):257–266, 2000.

[KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, Jun 2002.

[Kni10] E. Knill. Physics: Quantum computing. *Nature*, 463:441–443, Jan 2010.

[KPS+12] A. Khromova, Ch. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and Ch. Wunderlich. Designer Spin Pseudomolecule Implemented with Trapped Ions in a Magnetic Gradient. *Phys. Rev. Lett.*, 108:220502, Jun 2012.

[LBB13] A. Lemmer, A. Bermudez, and Plenio M. B. Driven geometric phase gates with trapped ions. *New Journal of Physics*, 15(8):083001, 2013.

[LBC+12] E. Lucero, R. Barends, Y. Chen, J. Kelly, M. Mariantoni, A. Megrant, P. O'Malley, D. Sank, A. Vainsencher, J. Wenner, T. White, Y. Yin, A.N. Cleland, and J. M. Martinis. Computing prime factors with a Josephson phase qubit quantum processor. *Nature*, 8:719–723, Oct 2012.

[LBMW03] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003.

[LBYP07] C. Y. Lu, D. E. Browne, T. Yang, and J. W. Pan. Demonstration of a Compiled Version of Shor's Quantum Factoring Algorithm Using Photonic Qubits. *Phys. Rev. Lett.*, 99:250504, Dec 2007.

[LDH+03] D. M. Lucas, C. J. S. Donald, J. P. Home, M. J. McDonnell, A. Ramos, D. N. Stacey, J. P. Stacey, A. M. Steane, and S. C. Webster. Oxford ion-trap quantum computing project. *Phil. Trans. R. Soc. Lond. A*, 361(1808):1401–1408, 2003.

[LDM+03] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422:412–415, Mar 2003.

[LGA+08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibrandt, K. R. Brown, and I. L. Chuang. Suppression of Heating Rates in Cryogenic Surface-Electrode Ion Traps. *Phys. Rev. Lett.*, 100:013001, Jan 2008.

[Lin76] G. Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, 1976.

[LJL+10] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien. Quantum computers. *Nature*, 464:45–53, Mar 2010.

[LKH+07] D. M. Lucas, B. C. Keitch, J. P. Home, G. Imreh, M. J. McDonnell, D. N. Stacey, D. J. Szwer, and A. M. Steane. A long-lived memory qubit on a low-decoherence quantum bus. *ArXiv e-prints*, Oct 2007.

[LKOW07] D. Leibfried, E. Knill, C. Ospelkaus, and D. J. Wineland. Transport quantum logic gates for trapped ions. *Phys. Rev. A*, 76:032324, Sep 2007.

[LRH+04] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J.-P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photoionization for calcium ion trapping. *Phys. Rev. A*, 69:012711, Jan 2004.

[LWL+07] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White. Experimental Demonstration of a Compiled Version of Shor's Algorithm with Quantum Entanglement. *Phys. Rev. Lett.*, 99:250505, Dec 2007.

[Mau16] P. Maunz. High-fidelity two qubit quantum gates in a scalable surface ion trap. reported in MPQ seminar, May 2016.

[MK13] C. Monroe and J. Kim. Scaling the Ion Trap Quantum Processor. *Science*, 339(6124):1164–1169, 2013.

[MLLL+12] E. Martin-Lopez, a. Laing, T. Lawson, R. Alvarez, X. Q. Zhou, and J. L. O'Brien. Experimental realization of Shor's quantum factoring algorithm using qubit recycling. *Nature Photonics*, 6:773–776, Nov 2012.

[MSW+08] 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. *Phys. Rev. Lett.*, 100:200502, May 2008.

[MW01] F. Mintert and C. Wunderlich. Ion-Trap Quantum Logic Using Long-Wavelength Radiation. *Phys. Rev. Lett.*, 87:257904, Nov 2001.

[NC11] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.

[NFB14] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin. Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links. *Phys. Rev. X*, 4:041041, Dec 2014.

[OIB+07] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Phys. Rev. A*, 75:042329, Apr 2007.

[OLA+08] 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. *Phys. Rev. Lett.*, 101:090502, Aug 2008.

[OLK03] J. E. Ollerenshaw, D. A. Lidar, and L. E. Kay. Magnetic Resonance Realization of Decoherence-Free Quantum Computation. *Phys. Rev. Lett.*, 91:217904, Nov 2003.

[OWC+11] 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, Aug 2011.

[PMO09] A. Politi, J. C. F. Matthews, and J. L. O'Brien. Shor's Quantum Factoring Algorithm on a Photonic Chip. *Science*, 325(5945):1221–1221, 2009.

[Pre98] J. Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 454:385–410, Jan 1998.

[PSVW14] C. Piltz, T. Sriarunothai, A. F. Varón, and C. Wunderlich. A trapped-ion-based quantum byte with 10^{-5} next-neighbour cross-talk. *Nat Commun*, 5, Aug 2014.

[RKS+06] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt. Process Tomography of Ion Trap Quantum Gates. *Phys. Rev. Lett.*, 97:220407, Dec 2006.

[SDSM16] E. Strohmaier, J. Dongarra, H. Simon, and M. Meuer. TOP500 list. http://www.top500.org, Jun 2016.

[Sho97] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[SKGR+03] F. Schmidt-Kaler, S. Gulde, M. Riebe, T. Deuschle, A. Kreuter, G. Lancaster, C. Becher, J. Eschner, H. Häffner, and R. Blatt. The coherence of qubits based on single Ca^+ ions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 36(3):623, 2003.

[SKHR+03] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422:408–411, Mar 2003.

[SM99] A. Sørensen and K. Mølmer. Quantum Computation with Ions in Thermal Motion. *Phys. Rev. Lett.*, 82:1971–1974, Mar 1999.

[SNJ+11] E.W. Streed, B. G. Norton, A. Jechow, T. J. Weinhold, and D. Kielpinski. Imaging of Trapped Ions with a Microfabricated Optic for Quantum Information Processing. *Phys. Rev. Lett.*, 106:010502, Jan 2011.

[SSB+92] I. Siemers, M. Schubert, R. Blatt, W. Neuhauser, and P. E. Toschek. The "trapped state" of a trapped ion - line shifts and shape. *EPL (Europhysics Letters)*, 18(2):139, 1992.

[Szw09] D. Szwer. *High Fidelity Readout and Protection of a $^{43}Ca^+$ Trapped Ion Qubit*. PhD thesis, University of Oxford, 2009.

[Tar] J. E. Tarlton. PhD thesis, Imperial College London. In preparation.

[TBJ+11] N. Timoney, I. Baumgart, M. Johanning, A. F. Varón, M. B. Plenio, A. Retzker, and Ch. Wunderlich. Quantum gates and memory using microwave-dressed states. *Nature*, 476:185–188, Aug 2011.

[TGB+13] T. R. Tan, J. P. Gaebler, R. Bowler, Y. Lin, J. D. Jost, D. , and D. J. Wineland. Demonstration of a Dressed-State Phase Gate for Trapped Ions. *Phys. Rev. Lett.*, 110:263002, Jun 2013.

[TKK+00] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Phys. Rev. A*, 61:063418, May 2000.

[TMK+00] Q. A. Turchette, C. J. Myatt, B. E. King, C. A. Sackett, D. Kielpinski, W. M. Itano, C. Monroe, and D. J. Wineland. Decoherence and decay of motional quantum states of a trapped atom coupled to engineered reservoirs. *Phys. Rev. A*, 62:053807, Oct 2000.

[TPR+03] G. Tommaseo, T. Pfeil, G. Revalde, G. Werth, P. Indelicato, and J. P. Desclaux. The g_J-factor in the ground state of Ca^+. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 25(2):113–121, 2003.

[TWK+98] Q. A. Turchette, C. S. Wood, B. E. King, C. J. Myatt, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. Deterministic Entanglement of Two Trapped Ions. *Phys. Rev. Lett.*, 81:3631–3634, Oct 1998.

[VCA+10] A. P. VanDevender, Y. Colombe, J. Amini, D. Leibfried, and D. J. Wineland. Efficient Fiber Optic Detection of Trapped Ion Fluorescence. *Phys. Rev. Lett.*, 105:023001, Jul 2010.

[vS03] M. Šašura and A. M. Steane. Fast quantum logic by selective displacement of hot trapped ions. *Phys. Rev. A*, 67:062318, Jun 2003.

[VSB+01] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang. Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414:883–887, Dec 2001.

[VSS+00] L. M. K. Vandersypen, M. Steffen, M. H. Sherwood, C. S. Yannoni, G. Breyta, and I. L. Chuang. Implementation of a three-quantum-bit search algorithm. *Applied Physics Letters*, 76(5):646–648, 2000.

[WAIQ13] M. Waseem, R. Ahmed, M. Irfan, and S. Qamar. Three-qubit Grover's algorithm using superconducting quantum interference devices in cavity-QED. *Quantum Information Processing*, 12:3649–3664, Dec 2013.

[WHZ+16] M. Wahnschaffe, H. Hahn, G. Zarantonello, T. Dubielzig, S. Grondkowski, A. Bautista-Salvador, M. Kohnen, and C. Ospelkaus. Exploring and understanding 2D microwave near-fields using a single ion. *ArXiv e-prints*, 2016.

[WI79] D. J. Wineland and Wayne M. Itano. Laser cooling of atoms. *Phys. Rev. A*, 20:1521–1540, Oct 1979.

[WMW+98] D.J. Wineland, C. Monroe, Itano W.M., D. Leibfried, King B.E., and Meekhof D.M. Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions. *J. Res. Natl. Inst. Stand. Technol.*, 103(3):259, 1998.

[WOC+13a] U. Warring, C. Ospelkaus, Y. Colombe, K. R. Brown, J. M. Amini, M. Carsjens, D. Leibfried, and D. J. Wineland. Techniques for microwave near-field quantum control of trapped ions. *Phys. Rev. A*, 87:013437, Jan 2013.

[WOC+13b] U. Warring, C. Ospelkaus, Y. Colombe, R. Jördens, D. Leibfried, and D. J. Wineland. Individual-ion addressing with microwave field gradients. *Phys. Rev. Lett.*, 110:173002, Apr 2013.

[Woo80] G. K. Woodgate. *Elementary atomic structure*. Oxford science publications. Clarendon Press, 1980.

[WRR+05] P. Walther, K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger. Experimental one-way quantum computing. *Nature*, 434:169–176, Mar 2005.

[WRW+15] S. Weidt, J. Randall, S. C. Webster, E. D. Standing, A. Rodriguez, A. E. Webb, B. Lekitsch, and W. K. Hensinger. Ground-State Cooling of a Trapped Ion Using Long-Wavelength Radiation. *Phys. Rev. Lett.*, 115:013002, Jun 2015.

[WRW+16] 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. *ArXiv e-prints*, Mar 2016.

[WSGS12] G. Wilpers, P. See, P. Gill, and A. G. Sinclair. A monolithic array of three-dimensional ion traps fabricated with conventional semiconductor technology. *Nature Nanotechnology*, 7:572–576, Sept 2012.