CONTROL OF MULTI-SPECIES TRAPPED IONS IN A SCALABLE QUANTUM COMPUTER

MARTIN W. VAN MOURIK

A dissertation submitted to the Faculty of Mathematics, Computer Science, and Physics in the University of Innsbruck

Carried out under the supervision of Prof. Rainer Blatt at the Quantum Optics and Spectroscopy Group

Institute for Experimental Physics University of Innsbruck



Computers have revolutionized modern society and have provided us with immense calculating power. In recent decades, a new revolution in computing technology has begun to develop, in which the classically inexplicable properties of quantum physics are utilized to speed up certain computations. A prominent example of such a computational speed-up is Shor's algorithm for factorization, in which factorizing large prime numbers could take several minutes on a quantum computer, compared with thousands of years on a classical one.

The fundamental unit of information in quantum computation is a qubit, the quantum analog of a classical computer's bit. Ions stored in radio-frequency traps are a promising platform for storing and manipulating qubit information. In order for quantum computation to outperform its classical counterpart in certain algorithms, thousands of qubits will be required. As the field of trapped-ion quantum computation is still in its infancy, reaching such numbers of qubits is still a technological challenge.

This work documents the design and experimental use of a trapped ion quantum computer based on a cryogenic segmented surface ion trap, capable of trapping ⁴⁰Ca⁺ and ⁸⁸Sr⁺ ions. The apparatus is designed to demonstrate the prospects of scaling up the number of usable qubits. Three distinct projects are presented in this thesis.

Firstly, we investigate a property of radio-frequency traps known as RF heating, an adverse coupling between ions and the RF field required to operate the trap. This coupling perturbs an ion's motion and renders them useless for quantum computation. We present a conceptual and theoretical background of RF heating, perform ion-motion simulations and experimentally confirm our findings. Secondly, we investigate transport operations of trapped ions. As ion-based quantum computers are envisioned to have dedicated regions in the trap for various tasks (such as qubit manipulation, ion-ion coupling, storage, cooling, and loading), traps are designed to be able to move ions from region to region. We describe and experimentally realize techniques to perform and optimize several types of transport operations. Lastly, we extend the current toolbox of quantum computation by introducing a method of qubit control that combines elements of classical and quantum computation: we apply irreversible boolean operations to a quantum system. Such techniques could simplify quantum error correction protocols.

The results presented in this thesis have led to the following publications:

- M. W. van Mourik, E. A. Martinez, L. Gerster, P. Hrmo, T. Monz, P. Schindler, and R. Blatt. "Coherent rotations of qubits within a surface ion-trap quantum computer." In: *Physical Review A* 102.2 (2020), p. 022611.
- M. W. van Mourik, P. Hrmo, L. Gerster, B. Wilhelm, R. Blatt, P. Schindler, and T. Monz. "Rf-induced heating dynamics of noncrystallized trapped ions." In: *Physical Review A* 105.3 (2022) p. 033101.
- M. W. van Mourik, E. Zapusek, P. Hrmo, L. Gerster, R. Blatt, T. Monz, P. Schindler, and F. Reiter. "Experimental realization of nonunitary multi-qubit operations." In: Arxiv:2303.06098 (2023). Accepted in *Phys. Rev. Lett.* (Nov 2023).

During my Ph.D., I contributed to the following works:

- M. F. Brandl, M. W. van Mourik, L. Postler, A. Nolf, K. Lakhmanskiy, R. R. Paiva, S. Möller, N. Daniilidis, H. Häffner, V. Kaushal, T. Ruster, C. Warschburger, H. Kaufmann, U. G. Poschinger, F. Schmidt-Kaler, P. Schindler, T. Monz, R. Blatt. "Cryogenic setup for trapped ion quantum computing." In: *Rev. Sci. Instrum.* 87, 113103 (2016).
- L. Gerster, F. Martínez-García, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz. "Experimental Bayesian Calibration of Trapped-Ion Entangling Operations." In: *PRX Quantum* 3.2 (2022), p. 020350.
- P. Hrmo, B. Wilhelm, L. Gerster, M. W. van Mourik, M. Huber, R. Blatt, P. Schindler, T. Monz, M. Ringbauer. "Native qudit entanglement in a trapped ion quantum processor." In: *Nature Communications* 14.1 (2023), p. 2242.

It's dangerous to go alone

— Hylian proverb

My decision to leave my hometown of Eindhoven to come to Innsbruck for a Ph.D. was in part motivated by a desire to be put out of my comfort zone. That plan backfired, because the people that surrounded me during my years here made my time both in and out of the lab unexpectedly comfortable; for that I am extremely grateful.

First and foremost, my gratitude goes out to Rainer Blatt. From the day that I applied he has been supportive and motivating and has made me feel welcome within his group. He has a clear passion for physics, and attempts to cultivate that feeling in the rest of the group. Importantly, Rainer believes that this passion should stem from curiosity and an eagerness to discover. I'd like to think that that mentality has rubbed off on me a bit.

I extend that same gratitude to my supervisors Thomas Monz and Philipp Schindler. Tomi has done a tremendous amount of work to streamline my Ph.D., ranging from making sure I still get paid, to brainstorming and debugging physics issues, to randomly sneaking by to poke me. Philipp has always kept his door open for enjoyable and fruitful physics discussions and to many many iterations of paper corrections. As annoying as it is that he is almost always right, his help has been invaluable to the success of our experiments.

The SQIP setup has had its fair share of people tinkering with it. To begin with, Matthias Brandl, Rafael Rothganger, Lukas Postler and Anton Nolf welcomed me with open arms and eased the process of diving head first into the deep end of the world of cryogenic ion trapping. Matthias in particular has patiently taught me a great deal about quantum physics, mechanical engineering, electrical engineering, and lab-based innuendos. Lukas Gerster, like me a long-time child of the SQIP family, has been an integral part of many of the improvements that the experiment has seen in years, and has been a faithful, friendly and diligent lab-mate. Likewise, Pavel Hrmo and Benjamin Wilhelm were instrumental in some of SQIP's great successes, and were great company in the lab, in the office, and at the bar.

My work would certainly not have been possible without the aid of all the administrative and technical staff working behind the scenes, doing a lot of bureaucratic and technical tasks for me. In particular, Patricia Moser somehow finds the time to be of assistance to an increasingly large collection of disoriented group members, and deserves great credit and thanks. I would also like to thank Renate from accounting, Armin and Christoph from the mechanical workshop, and Wolfgang and Killian from the electronics workshop for their assistance and friendliness.

As thankful as I am for the fantastic work conditions that the group has provided, I firmly believe that a person's mental well-being should, for the majority, be established by life outside of work. As such, I have managed to maintain sanity due to the long list of people that have made living in Innsbruck unforgettable. These include (but are not limited to): Gerard, Cristine, Gabriel, Marina, Esteban, Nadine, Barbara, Bojana, Guilia, Veronica, Cristina, Daniel, Ben, Philip, Sasha, Josef, John, Mary, Fi, Jonas, Christine, Christian, Leni, Dominik, Matthias, Tuomas, Manoj, Lata, Elisa, Lorenzo, and Marco. I also got to share my Innsbruck experience with many of my old friends from the Netherlands who have visited me here over the years and have treated me as if I've never left whenever I visit them. These include Damien, Geert, Ulyana, Gillis, Lidewij, Gar, Marjolein, Manon, and Kees. Thank you all.

Of course I would not have become the well-adjusted young man I am today if it were not for my family. Mom and Dad, I am immensely thankful for your support, interest, compassion, and your restraint in asking if I'm ever going to be finishing my Ph.D. Reinier, Ed, you guys are like brothers to me.

Finally, my admiration, thanks, and love goes out to the one who has managed to keep me young: Milena, I look forward to our next adventures.

CONTENTS

```
1 INTRODUCTION
  TRAPPED IONS FOR QUANTUM COMPUTATION
                                                       5
   2.1 Ion trapping
        2.1.1
              RF trap overview
                                    5
              Linear Paul trap
        2.1.2
              Segmented surface traps
        2.1.4
              Ions in a linear Paul trap - positions and modes
                                                                10
   2.2 Trapped ions as qubits
                                  14
              Qubits
                         14
        2.2.1
              Ion - light interaction
        2.2.2
              Rabi oscillations
        2.2.3
              AC Stark shift
        2.2.4
        2.2.5
              Decoherence
              Light-atom interaction of a harmonically confined ion
        2.2.6
                                                                      19
              Qubit operations of Group II ions
        2.2.7
              Experimental sequences
        2.2.8
   2.3 Surface traps: potentials and characterization
              Trapping potentials - Spherical harmonics expansion
        2.3.1
                                                                     27
              Potential calibration
        2.3.2
              Micromotion compensation
        2.3.3
                                             31
        2.3.4
              Heating rates
                               33
              Coherence time
        2.3.5
              Entanglement
        2.3.6
                                37
  EXPERIMENTAL SETUP
   3.1 Setup overview
              Cryogenic apparatus
              Laser setup
        3.1.2
              Experimental control
        3.1.3
       'Golden Gate' surface ion trap
        3.2.1
              Previous traps
                                 56
              Golden Gate trap overview
        3.2.2
                                             59
        3.2.3
              Manufacturing
              Potential calibration
        3.2.4
              Micromotion compensation
        3.2.5
                                             60
              Heating rates
        3.2.6
              Coherence times
        3.2.7
        3.2.8
              Gate performance
                                    64
  CLASSICAL ION DYNAMICS IN AN RF PAUL TRAP
                                                         67
   4.1 Motivation
                      67
                                                      68
   4.2 Beyond the pseudopotential approximation
       RF heating of non-crystallized trapped ions
                                                      69
              RF heating: an overview
       4.3.2
              Ion clouds
                             70
              Heating in time-dependent potentials
       4.3.3
                                                       72
   4.4 Simulation overview
       4.4.1 Contributing forces
                                     74
       4.4.2 Analytic description of surface trap potentials
   4.5 Full ion dynamics simulations
       4.5.1
              Simulation overview
              Extracting secular energy
```

```
Simulation results
       4.5.3
                                   79
       Collision model
       4.6.1
              Model parameters
       4.6.2
              Simplified RF heating model
                                             86
              Recooling an ion cloud
       4.6.3
       Experimental validation
              Experiment Overview
       4.7.1
                                       96
              Experimental RF heating results
       4.7.2
              Experimental recrystallization results
   4.8 Outlook: Further considerations of RF heating
   4.9 Conclusions
                      105
  TRANSPORT OPERATIONS OF TRAPPED-ION QUBITS
   5.1 Introduction
                      107
   5.2 Shuttling
              Potential ramp design
       5.2.1
                                      110
              Voltage solution for shuttling potentials
       5.2.3
              Experimental implementation of shuttling 114
   5.3 Splitting
                   114
              Splitting potentials
       5.3.1
                                   115
              Splitting voltages
       5.3.2
                                  116
              Splitting field sensitivity
       5.3.3
                                        119
              Experimental demonstration of ion chain splitting
       Ion crystal rotations
       5.4.1 Rotation overview
                                   121
       5.4.2
              Rotation potentials
                                   123
              Experimental implementation
       5.4.3
              Experimental results
       5.4.4
              Summary and outlook
                                       140
       5.4.5
6 CONTROL OF QUBITS WITH DISSIPATION THROUGH ENGINEERED RESONANCE
                                                                                    143
   6.1 Dressed state splitting for dissipation through engineered resonance
              Dressed state splitting of one ion
                                               145
       6.1.2
              Two ion dressed state splitting
       6.1.3
              Dissipation
                            147
   6.2 Dissipation through engineered resonance for quantum error correction
              Minimal instance dissipative error correction
              Autonomous error correction with engineered resonance dissipation 152
   6.3 Boolean logic gates with engineered resonance dissipation
              Experimental overview
       6.3.1
       6.3.2
              Results
                       158
              Sources of error
       6.3.3
       6.3.4
              Conclusions
  SUMMARY AND OUTLOOK
   7.1 Recap
                165
   7.2 Next steps for the experimental apparatus
   7.3 Conclusion
A 3D COLLISION RATE
   A.1 Collision rate derivation
   A.2 Collision rate verification
B TWO-ION HEATING MEASUREMENTS
C DRESSED STATE DYNAMICS
   c.1 Dressed state dynamics of one ion
                                           175
   c.2 Dressed state dynamics of two ions
                                            176
   c.3 Dressed state dynamics of minimal instance error correction
D CLASSICAL QUBIT GATE ERRORS
```

E NOTES 187

LIST OF FIGURES

T.	
Figure 2.1	Schematic overview of linear Paul trap geometry 8
Figure 2.2	Layout of a linear surface trap 10
Figure 2.3	Various configurations of axial modes of ⁴⁰ Ca ⁺ and ⁸⁸ Sr ⁺ ion chains 12
Figure 2.4	Bloch sphere representation of a qubit state 14
Figure 2.5	Bloch sphere representation and excitation plots of typical qubit operations 16
Figure 2.6	Rabi excitation spectrum 17
Figure 2.7	Overview of qubit transitions of group II ions 22
Figure 2.8	Basic experimental operations of group II ions 24
Figure 2.9	Photon correlation method for micromotion compensation 33
Figure 2.10	Examples of motional mode occupation measurements 34
Figure 2.11	Bloch sphere representation of a Ramsey sequence for a coherence time mea-
Eigene a ca	surement 36
Figure 2.12	Schematic of MS gate levels and transitions 37
Figure 3.1	Experimental setup version history 39
Figure 3.2	Cross-sectional view of experimental apparatus 42
Figure 3.3	Overview of the inner shield 45
Figure 3.4	Gold mesh grid for electric field shielding 47
Figure 3.5	Schematic of DC electrode voltage wiring and filtering 48
Figure 3.6	Level schemes of Calcium and Strontium 50
Figure 3.7	Schematic overview of laser setup 51
Figure 3.8	Optical setup on the experiment table 54
Figure 3.9	Schematic overview of the experimental control hardware 55
Figure 3.10	Microscope image and electrode layout of the Berkeley trap 56
Figure 3.11	Microscope image and electrode layout of the Sandia HOA trap 57
Figure 3.12	Image and electrode layout of the Golden Gate trap 58
Figure 3.13	Pseudopotential of the Golden Gate trap 59
Figure 3.14	Results of micromotion compensation on the Golden Gate trap 61
Figure 3.15	Golden Gate trap heating rates 62
Figure 3.16	Golden Gate trap coherence times 63
Figure 3.17	Mølmer-Sørensen gate experimental results 64
Figure 4.1	Examples of ion dynamics where the pseudopotential approximation does not
	accurately portray ion dynamics 69
Figure 4.2	Schematic representation of secular energy change due to Coulomb interaction
	in an RF field 73
Figure 4.3	Removal of RF motion from ion trajectory data 79
Figure 4.4	Secular energy of simulated ion clouds 80
Figure 4.5	Statistics of ion-ion collisions and their secular energy changes 81
Figure 4.6	Analysis of energy change integrals of ion-ion collisions 83
Figure 4.7	Schematic overview of variables used for simulating collision energy 88
Figure 4.8	Comparison of the full ion dynamics simulation to the simplified simula-
	tion 90
Figure 4.9	Two-ion cloud energy dynamics, obtained with the simplified simulation 91
Figure 4.10	Simulated energy dynamics with Doppler cooling 96
Figure 4.11	Overview of fluorescence measurements in ion melting experiments 99
Figure 4.12	Experimental results of ion cloud energy change due to RF heating 100
Figure 4.13	Experimental results of ion cloud recooling 100
Figure 4.14	Simulated ion energy gain under various conditions 102
Figure 5.1	Overview of ion crystal transport operations 108
Figure 5.2	Overview of shuttling strategies 111

Figure 5.3	Voltages for shuttling an ion to a neighboring segment 114
Figure 5.4	Experimental demonstration of ion shuttling sequences 115
Figure 5.5	Schematic of ion chain splitting potentials 116
Figure 5.6	Axial mode frequency as a function of quadratic and quartic potential, for use in two-ion splitting 117
Figure 5.7	Ion chain splitting electrode voltages, mode frequency, and axial potential 118
Figure 5.8	Unsuccessful splitting due to a constant axial field 119
Figure 5.9	Experimental demonstration of ion chain splitting 120
Figure 5.10	Schematic overview of potentials used for ion crystal rotations 122
Figure 5.11	Properties of simulated ion crystal rotations 126
Figure 5.12	Electronic filter schematic and transfer function 128
Figure 5.13	Comparison of set potentials and simulated resulting potentials due to electronic voltage filtering 129
Figure 5.14	Schematic overview of rotation sequence optimization 130
Figure 5.15	Schematic of a neural network for machine learning 131
Figure 5.16	Automated rotation calibration performance 133
Figure 5.17	Simulated potentials and mode frequencies of the optimized rotation sequence 134
Figure 5.18	Rotation induced heating 136
Figure 5.19	Ramsey contrast after a set of ion crystal rotations 138
Figure 5.20	
Figure 5.21	Experimental success rate of Ca - Sr rotations under the influence of an offset
	field 139
Figure 5.22	A SWAP gate requires three entangling gates (MS _x ($\pi/4$)) and a set of local
	Pauli operations. 142
Figure 6.1	Conceptual diagrams of quantum error correction with classical feedback and dissipative state transfer 144
Figure 6.2	Excitation in a three-level system 145
Figure 6.3	Two-ion dressed state splitting 146
Figure 6.4	Level scheme for minimal instance dissipative error correction 149
Figure 6.5	Level schemes and excitation diagrams for minimal instance dissipative error correction 150
Figure 6.6	Autonomous error correction protocol 153
Figure 6.7	Truth tables for the NOR and OR operations 154
Figure 6.8	Level diagram for ⁴⁰ Ca ⁺ and ⁸⁸ Sr ⁺ states and transitions used in for classical gates with dissipation through engineered resonance 155
Figure 6.9	Excitation spectra for initial states $ 00\rangle$ and $ 01\rangle$, showing dressed state splitting.
E' (The solid lines are simulated data. 156
Figure 6.10	Pulse sequence diagram for the OR and NOR operations 157
Figure 6.11	Engineered resonance population transfer in the NOR gate 158
Figure 6.12	Measured dissipative population transfer using sympathetic cooling
Figure 6.13	Measured population truth tables of the OR and NOR gates 160
Figure a.1	Graphical representation of the collision condition, for deriving the 3D collision frequency 170
Figure a.2	Numerical analysis of collision time model 172
Figure d.1	Simulated engineered resonance OR gates with various sources of error 181
Figure d.2	Simulated engineered resonance OR gates with various sources of error (cont.) 182
Figure d.3	Simulated engineered resonance NOR gates with various sources of error 183
Figure d.4	Simulated engineered resonance NOR gates with various sources of error (cont.) 184
Figure d.5	Overview of engineered resonance gate errors 185

LIST OF TABLES

Table 3.1	Overview of lasers 50	
Table 4.1	Fit results for RF-induced energy diffusion model 92	
Table 5.1	Rotation success rate truth table 135	
Table d.1	Default parameters of the engineered resonance gate simulations 179	

ACRONYMS

RF radio-frequency

IXI	radio frequency
DC	direct current
ODE	ordinary differential equation
DE	differential equation
AOM	acousto-optic modulator
PMT	photomultiplier tube
CCD	charge-coupled device
MOS	metal oxide semiconductor
QCCD	Quantum Charge-Coupled Device
MS	Mølmer-Sørensen
AWG	arbitrary waveform generator
DAC	digital-to-analog converter
QIP	quantum information processing
QC	quantum computation
UCB	University of California, Berkeley
HOA	High Optical Access
UHV	ultra-high vacuum
OFHC	oxygen-free high conductivity
NA	numerical aperture
ITO	indium-tin-oxide
TTL	transistor-transistor logic
PGA	pin grid array
PI	photoionization
PBS	polarizing beam splitter
PDH	Pound-Drever-Hall
ULE	ultra-low-expansion
FNC	fiber-noise cancellation
VCO	voltage-controlled oscillator
YBCO	yttrium barium copper oxide
HTS	high-temperature superconductor

- TA tapered amplifier
- TrICS Trapped Ion Conrol Software
- DDS Direct Digital Synthesiser
- FPGA field programmable gate array
- SEM scanning electron microscope
- RWA rotating wave approximation
- MS Mølmer-Sørensen
- EIT electromagnetically induced transparancy
- PGC polarization gradient cooling
- QEC quantum error correction
- DEC dissipative error correction
- AOD acousto-optic deflector
- PCB printed cicuit board
- PID proportional-integral-derivative

INTRODUCTION

The development of the field of quantum physics over the past decades has been so astounding that it has become one of the greatest clichés for Ph.D. theses with the word 'quantum' in the title to start the introduction with some permutation of that statement. This thesis is no exception: The past decades have seen an unbelievable growth in the study and application of quantum properties of well-controlled systems [1], which less than a century ago was deemed to be impossible to experimentally reach and manipulate (Erwin Schrödinger compared working with individual quanta to raising "Ichthyosauria in a zoo" [2]). Quantum mechanics, the study of light and matter at the atomic scale, provides theories to help us understand *microscopic* systems, but has applications in the *macroscopic* 'classical' world [3], and therefore attracts great interest from the scientific community.

One of the many envisioned applications of quantum mechanics is its inclusion in the field of computation [4]. It is hard to overstate the power and capabilities of computers, both for research and public use. Computers open up a world of possibilities and applications, ranging from solving complex non-linear equations, to encrypting and decrypting data, to transmitting cat videos. Such applications are only limited by computational resources such as the available memory, processing speed, and energy supply. These finite resources make certain tasks a near impossibility with conventional classical computational techniques [5].

Fortunately, there are motivated minds that look for alternate routes around those conventional techniques. Richard Feynman is one example of such a mind: he discussed the prospect of using computers to simulate quantum dynamics of collections of interacting particles [6]. The limitation lies in the fact that a collection of N quantum particles can exist in any combination of 2^N different states. Simulating the full dynamics of such a system, including possible sources of external noise, could require solving a system of up to 2^N coupled equations. Storing information about a system containing N > 130 particles would already require more bits than there are particles in the universe—and we'd have to find whole other universes worth of particles to do any kind of calculation with all of that information. Feynman therefore proposed that for a device to realistically handle such a simulation, "let the computer itself be built of quantum mechanical elements which obey quantum mechanical laws" [6]. In such a simulation, known as *analog* quantum computation (QC), a system of particles with controllable quantum states and interactions are used to mimic, thus simulate, those properties of another less-controllable system [7, 8]. Feynman's ideas were seminal, but at the time (1981), purely theoretical. It would take a few more decades before such simulators made their way from textbooks to laboratories.

In the meantime, another branch of quantum computers was founding its roots: digital QC [9]. This type of computation, compared with analog QC, is fundamentally more similar to classical computation in that it operates through logical operations applied sequentially to binary encoded information carriers. In contrast, however, quantum physics opens up a new paradigm of information processing, by allowing information carriers to make use of properties of quantum mechanics that do not have a classical analog: superposition and entanglement. Harnessing the quantum analogue of bits, *qubits*, allows us to perform specific computation tasks on timescales notably shorter than required of classical computers.

The go-to example of a task whose computational cost is greatly reduced by the inclusion of quantum mechanics is Shor's algorithm for prime factorization [10]. Prime factorization of large numbers is notoriously time-consuming for classical computers, which is one of its integral features in the scope of data encryption: Decryption of such data requires the near-impossible task of factorizing large numbers. Running Shor's algorithm on a quantum computer, (specifically a subroutine known as period-finding), however, has the potential to lower the computational time required for this task, benefiting from polynomial scaling of number of required operations per number of digits of the number to be factorized, compared to an exponential scaling for the best known classical algorithm [11].

1

In the past decade, tremendous progress has been made in demonstrating Shor's algorithm on quantum computers [12–14]. However, it is worth keeping perspective: such implementations have been able to determine the prime factors of numbers as large as 21 [13, 14], a value that classical computers have little difficulty working with¹. That statement is not intended to downplay these achievements, but to emphasize the immense complexity in controlling a quantum system well enough to perform calculations that aim to one day compete with classical computers. We are now in the infancy of a new era of information technology, and are overcoming challenges and obstacles to work towards a regime where digital quantum computers undeniably outperform their classical counterparts. To reach this goal, there exist many promising platforms of quantum computers, such as solid-state superconductors [15], photonics [16], neutral atoms [17], implanted defects in solid-state crystals [18], and trapped ions. This thesis explores the prospects of using trapped ions as qubits for quantum computation [19].

Trapped ions have several properties that make them promising candidates to host qubits: they have long-lived electronic levels that can be used to encode the qubit, which can be manipulated using optical or microwave fields [20]. Other electronic levels have short-lived transitions, ideal for rapid readout of the qubit state. Ions, being charged particles, can be isolated from the surrounding environment through the use of dynamic electric fields [21]. This separation from the environment is crucial, because it is a source of uncontrolled perturbation of the qubit state. Another advantage of ions being charged, is that their mutual Coulomb interaction couples their motion, which provides a mechanism to exchange information and enable qubit entanglement.

One example of a device that produces the required dynamic electric fields to confine ions is known as a *Paul trap* [22], and comes in many shapes and sizes [23]. For the purpose of trapped-ion based quantum computation, Paul traps are generally a three-dimensional configuration of electrodes designed to stably confine ions along one dimension. Ions, separated by Coulomb repulsion, form a chain, referred to as an ion string. Laser beams can be focused down to manipulate the quantum state of each ion individually, all ions collectively, or some combination.

Using that concept of trapped-ion qubit manipulation, decrypting RSA-encrypted data, which involves factorizing numbers that are hundreds of digits, requires stably trapping some thousands—if not tens of thousands—of ions along a chain, probing them with a series of well-timed, well-aimed laser pulses, and reading out their final states. This is, of course, a drastic understatement of the complexity involved, as attested by the latest attempts to factorize the number 21 on a quantum computer. Many challenges stand in the way of building a quantum computer that competes with classical ones, most of which are related to reducing sources of noise that perturb the qubits' state. As the number of qubits and qubit operations in a quantum computer increases, so too does the error induced by these sources of noise.

One of the major envisioned obstacles in scaling up the number of usable qubits is that the single trapping potential that is produced by a traditional three-dimensional Paul trap is going to be overcrowded with ions. It becomes more challenging to keep ions in their crystalline structure along a single chain, and the increased density of motional modes in frequency space limit the success and speed of entangling operations. A proposed path towards scaling up the number of qubits for trapped-ion quantum computation is to use ion traps that contain multiple separated trapping regions, each containing a manageable amount of trapped ions. In 2002, Kielpinski et al. [24] proposed an array-like ion trap architecture, in which electrodes that carry the voltages required to produce ion-trapping potentials are confined to a two-dimensional surface. Such a Paul trap, which resembles the array-like layout of a charge-coupled device (CCD) chip, hosts numerous trapping regions with dedicated tasks, such as memory storage, qubit interaction, and state readout. This type of *segmented surface trap* chip, known as the Quantum Charge-Coupled Device (QCCD) architecture², remains one of the prominent visions of what scalable trapped-ion quantum information processing (QIP) will resemble [25–27].

Within the vision of the QCCD architecture, the aim is for an ion trap to produce a potential landscape in which separated trapping locations each form a quantum register and can be independently manipulated through varying electrode voltages. Independent control of ion registers requires that

¹ An exhaustive brute-force search on a cheap classical computer takes less than a millisecond to find $21 = 7 \times 3$.

² Note 1 in Appendix e

the separation between a register and an electrode is at most on the same order of magnitude as the separation between registers. This requires ions to be placed relatively close to trap electrodes, compared to most three-dimensional ion traps that have only one trapping region. Herein lies one of the prominent difficulties in surface trap-based quantum computation: electric field noise emanating from nearby surfaces adds thermal energy to ions, which negatively correlates with the success of many quantum operations [28]. The rate at which ions gain motional energy, the heating rate, due to surface electric field noise scales very unfavorably with the distance between ions and surfaces, more than an inverse-cubic scaling [29]. The design and fabrication of suitable surface trap architectures that minimize the influence of surface-induced electric field noise (along with a host of other obstacles) is an ongoing field of research in the ion-trap community [30–38].

An increase in ions' motional energy can be combated by laser-cooling techniques. Such techniques cannot be directly applied *during* computation sequences, since the access to the ions' electronic levels required for cooling disturbs the qubit state. One can bypass this restriction by introducing a second ion species into the register, which does not take part in computation and can be freely laser-cooled without affecting the qubits. Coulomb interaction couples the motion of qubits to that of the coolant species, thus remain cooled, a process known as *sympathetic cooling* [39–42].

Another method of maintaining a low motional energy of trapped ions is by cooling down the trap surface, which reduces the electric field noise emanating from it [43]. For this reason, surface traps are often operated at cryogenic temperatures. Cryogenic environments come with the additional benefit that they improve the vacuum pressure by orders of magnitude compared to room temperature setups [44], reducing the chance that particles in the background gas interfere with experiments by colliding with ions. Improved vacuum due to cryo-pumping reduces the amount of time required to bring a vacuum chamber down to acceptable pressures after installing a trap. It can take months after installing a trap into a room-temperature setup before it is ready for ion-trapping, whereas cryogenic setups could potentially be ready for loading ions within a day. Since the performance of surface traps, especially those with novel designs, is to-date often unpredictable, having a short turn-around time for switching out traps is beneficial.

When used for quantum computation, cryostats, ion traps, and experimental hardware are designed to minimize errors that undesirably influence a qubit's state. Despite the incredible technological achievements in such designs, qubits can never be perfectly protected from sources of noise. To ensure QC can still function in the presence of sources of error, quantum error correction (QEC) is a requirement [11]. QEC is, as the name suggests, the quantum variant of error correction schemes used in information theory to protect data from errors. The current technological status of trapped-ion quantum computation is at the point that errors induced by proposed correction algorithms [45–48] outweigh the ability of those algorithms to fix errors. Fault-tolerant QC requires both technological improvements to reduce sources of error, and improved error correction algorithms to combat those that remain [49–51].

Research in the field of *scalable quantum information processing* is tasked with implementing technological improvements and QEC protocols, to enable an increase of the number of trapped ion qubits from tens to hundreds to thousands, while avoiding and/or correcting errors. Many challenges must be dealt with, ranging from operating in cryogenic environments, to trap design and fabrication, to scaling of experimental hardware, to suppressing influence from noisy environments. This thesis documents a contribution to a vast field aiming to scale up the current paradigm of digital quantum computation with trapped ion qubits.

The work in this thesis details the design, development, and use of a cryogenic surface ion trap, designed for mixed-species operation, with the ultimate goal of advancing the field of scalable quantum information processing.

The thesis is structured as follows³: Chapter 1 recursively introduces the layout of the rest of this thesis. Chapter 2 gives a generalized overview of how ions are trapped and used as qubits for quantum computation. The core concepts of ion traps and their confining potentials are introduced, and a description of the motion of ions within such potentials is given. We describe how laser-ion interaction is used to manipulate an ion's electronic state, how this type of interaction is practically used to control

qubits, and how they are used within experimental sequences. Finally, experimental techniques that characterize qubits and the properties of surface traps are discussed.

Chapter 3 gives an overview of the experimental setup, introducing the cryogenic apparatus, the laser setup, the control software, and the surface traps designed and used throughout the course of this thesis. Experimental characterization of the performance of the most recent iteration of surface traps in our setup are presented.

Chapter 4 discusses numerical simulations and experimental results of a phenomenon unique to Paul traps that causes trapped ions to gain energy in specific regimes of ionic motion. This phenomenon, known as RF heating, can make ion motion chaotic and unpredictable, and regularly disrupts trapped ion experiments. Such disruptions are increasingly detrimental as the number of qubit registers increases. A solid understanding and effective strategies to manage the adverse effects of RF heating are therefore imperative for scalable quantum computation. The chapter investigates the physical origin of RF heating, and makes sense of the underlying chaos to produce accessible and generalized models that predict energy change rates of ions. Suggestions on recooling strategies are made, based on numerical simulations, which are experimentally verified.

Chapter 5 focuses on one of the most distinguishing features that makes segmented surface traps a good candidate for scalable QC: the ability to physically transport trapped ions during computional sequences. The main strategies for shuttling ions and splitting ion chains are discussed. Furthermore, experimental results of an ion reordering operation, ion crystal rotations, are shown, including the calibration routines required to realize them.

In Chapter 6, we explore a novel approach to qubit control, which extends the the current toolbox of quantum operations used within QC. The methods operate in a regime where quantum and classical control of qubits is combined, allowing for controlled non-unitary qubit operations. We discuss the engineering of a gate applied to trapped ions that generates state transfer only to specifically chosen initial states. In stark contrast with traditional QC gates, this state transfer is uni-directional, enabled by the dissipative process of sympathetic cooling. The primary envisioned application for this process, for which we have coined the term *dissipation through engineered resonance*, is in QEC, where it aims to circumvent the complications involved in non-destructive readout of error syndromes with a single non-unitary correction operation. Although a full implementation of this tool for QEC is technically challenging, we demonstrate the viability of the methods, through a proof-of-principle investigation using another less demanding type of non-unitary operation: we apply the tools of dissipation through engineered resonance to produce classical boolean logic gates on a quantum system.

TRAPPED IONS FOR QUANTUM COMPUTATION

It has been written for the technologist, and is not addressed in any sense to the pure mathematician, for whom I am not qualified to write. [...] Although, by virtue of necessity, the technologist may occasionally deviate from the narrow path followed rigorously by the pure mathematician, it must not be forgotten that the consequences of such deviation may be practical results of considerable benefit to the community at large.

- N.W. McLachlan, Theory and Application of Mathieu Functions

This section gives an overview of how trapped ions are used as qubits for quantum computation. In Section 2.1, we discuss how Paul traps confine ions in stable positions, while isolating them from the surrounding environment. Section 2.2 covers how the electronic states of trapped ions encode qubits, and how qubit states are manipulated and read out using electromagnetic fields. Finally, in Section 2.3, we discuss typical methods of characterizing performance specifications of trapped ions for use in quantum computation.

2.1 ION TRAPPING

Ions are charged particles, and therefore experience a force in the presence of an electric field. Paul traps [22], also known as radio-frequency (RF) traps, use a combination of static and oscillating electric fields to control and confine the position of ions, outlined below.

In this section, we start with a general overview of how electric fields can generate confining potentials for charged particles. Then, a common trap design, the macroscopic linear Paul trap, is introduced. Such a trap can store a linear chain of ions, an ideal configuration when used as individually addressable qubits. However, a macroscopic linear trap is not considered to be a scalable solution for ion-based quantum computation. We thus introduce another trap design as a promising candidate for scalable quantum computation, the planar segmented trap.

2.1.1 RF trap overview

In order to confine a charged particle at some point in space, one would desire to generate a static (DC) field that produces confinement in all three dimensions around that point. One of the most straightforward smooth confining potentials is a static harmonic potential, which corresponds to form

$$V_{\rm DC} = \frac{1}{2} \left(\alpha_x r_x^2 + \alpha_y r_y^2 + \alpha_z r_z^2 \right) \tag{2.1}$$

for positions r_k in an arbitrary orthogonal basis $k \in \{x, y, z\}$. Confinement in all dimensions is achieved when the field curvatures α_k are positive for all k. Unfortunately, Maxwell's equations prohibit this, since $\nabla^2 V_{\text{DC}} = 0$ requires that $\sum_k \alpha_k = 0$. Therefore, at least one of the values of α_k must be negative, and thus at least one axis will be anti-confining.

A Paul trap gets around this limitation by introducing, in addition to V_{DC} , a radio-frequency (RF) field, with potential V_{RF} . If this potential is harmonic in the same orthogonal basis as V_{DC} , and centered around the same position, this RF field can be expressed as

$$V_{\rm RF} = \frac{1}{2} \left(\beta_x r_x^2 + \beta_y r_y^2 + \beta_z r_z^2 \right) \cos(\Omega_{\rm RF} t) \tag{2.2}$$

with time t, and field curvatures β_k oscillating at a frequency Ω_{RF} .

A single charged particle, with charge q, mass m, and at position $\vec{r} = \{r_x, r_y, r_z\}$, experiences a force \vec{F} due to the combined potential $V_{\text{Tot}} = V_{\text{DC}} + V_{\text{RF}}$, given by $\vec{F} = -q \nabla V_{\text{Tot}}$. The equations of motion, $\vec{F} = m \frac{d^2 \vec{r}}{dt^2}$, are separable in each dimension, and are given by

$$\frac{d^2r_k}{dt^2} = \frac{q}{m} \left(\alpha_k + \beta_k \cos(\Omega_{RF}t) \right) r_k. \tag{2.3}$$

This equation can be written as

$$\frac{d^2r_k}{d\chi^2} + [a_k - 2q\cos(2\chi)] r_k = 0, \tag{2.4}$$

using the substitutions $t = 2\chi/\Omega_{\rm RF}$, and

$$a_k = \frac{4q\alpha_k}{m\Omega_{RF}^2}, \quad q_k = \frac{2q\beta_k}{m\Omega_{RF}^2}.$$
 (2.5)

Equation 2.4 is known as the Mathieu equation, with a_k and q_k aptly named Mathieu a and q parameters¹. Specific combinations of the a and q parameters lead to solutions of the Mathieu equations where the particle's motion is bounded [52] and are considered stable. Stable solutions can be analytically expressed as an infinite sum of trigonometric functions, whose coefficients are determined by infinitely continued fractions [53–55]. Ion traps are generally operated in a regime where $a_k \ll q_k \ll 1$, where luckily the infinite sums of infinite fractions can, by approximation, be truncated to a more intuitive equation of motion²:

$$r_k \approx r_k^{(0)} \cos(\omega_k t) \left(1 + \frac{q_k}{2} \cos(\Omega_{\text{RF}} t) \right).$$
 (2.6)

Here, $r_k^{(0)}$ is an amplitude that depends on initial conditions, and ω_k is a frequency whose exact value also requires a continued fraction to calculate exactly³. The same a_k and q_k restriction as before allows us to approximate ω_k as [55]

$$\omega_k pprox rac{\Omega_{
m RF}}{2} \sqrt{a_k + rac{q_k^2}{2}}.$$
 (2.7)

For sufficiently small a and q parameters, we have that $\omega_k \ll \Omega_{\rm RF}$. It is thus apparent that the motion described by Eq. 2.6 has two distinct timescales: The motion corresponding to the frequency ω_k is known as *secular motion*, and the motion at $\Omega_{\rm RF}$ is known as *micromotion*. This spectral separation (and the fact that in typical ion-trapping operation the micromotion amplitude is smaller than that of the secular motion), one can simplify an ion's motion by neglecting the higher-frequency micromotion. This is a useful simplification, because then the motion of an ion trapped in an RF trap can be seen as harmonic oscillator, with oscillation frequency ω_k . We will see in Section 2.2.6 that this harmonic motion plays a fundamental role in qubit manipulation. On the other hand, we will see in Chapter 4 that the simplification of neglecting micromotion is not always justified.

Starting from the assumption that a trapped ion has only harmonic motion, $r_k = r_k^{(0)} \cos(\omega_k t)$, we can derive an effective *static* potential \tilde{V}_{Tot} that would describe that motion⁴. Once again using the equations of motion, $-q\nabla \tilde{V}_{\text{Tot}} = m\ddot{r}$, with $\ddot{r}_k = -\omega_k^2 r_k$, we find

$$\tilde{V}_{\text{Tot}} = \sum_{k} \frac{1}{2} \frac{m\omega_k^2}{q} r_k^2. \tag{2.8}$$

¹ Note 3 in Appendix e

² Note 4 in Appendix e

³ Note 5 in Appendix e

⁴ Note 6 in Appendix e

 \tilde{V}_{Tot} can be written as the sum of the DC potential V_{DC} and an *effective*, time-independent potential \tilde{V}_{RF} :

$$V_{\rm DC} + \tilde{V}_{\rm RF} = \sum_{k} \frac{1}{2} \frac{m\omega_k^2}{q} r_k^2. \tag{2.9}$$

Substituting W_k with Eq. 2.7 and Eq. 2.5, and substituting V_{DC} with Eq. 2.1, yields

$$\tilde{V}_{\rm RF} = \frac{q}{4m\Omega_{\rm RF}^2} \sum_k \beta_k^2 r_k^2. \tag{2.10}$$

For a final substitution, we define the amplitude of the RF potential to be $V_{0,RF} \equiv 1/2 \sum_k \beta_k r_k^2$, which has the property

$$|\nabla V_{0,\text{RF}}| = \sqrt{\sum_{k} \beta_k^2 r_k^2},\tag{2.11}$$

thus giving

$$\tilde{V}_{RF} = \frac{q}{4m\Omega_{RF}^2} \left| \nabla V_{0,RF} \right|^2 \tag{2.12}$$

which is known as the pseudopotential approximation [56]. This approximation allows us to view the total trapping field as if it were a static harmonic potential, capable of confining trapped particles in all three dimension. The pseudopotential approximation simplifies the analysis of ion positions and their motional modes.

2.1.2 Linear Paul trap

Now that we have introduced the potentials that are used to confine ions, we can get creative with how to arrange electrodes to produce those potentials. There exist an abundance of variations of electrode configurations that produce confining potentials for charged particles [23], many of which have been realized in practice. However, when the goal is to utilize multiple trapped ions as qubits for quantum computation, a prevailing trap geometry is a *linear* trap. The design of a linear trap is such that there exists one axis that is free of any RF potential gradient. Referring to Eq. 2.2, this corresponds to one of the values of β_k being zero. We will choose $\beta_z = 0$, without loss of generality, which then sets the restriction that $\beta_x = -\beta_y$. The z-axis is free of RF, and is known as the RF-null. It is desirable to keep ions close to the RF-null, since along this region the RF field is at a minimum, and thus is the place where ions experience the least amount of micromotion. A collection of ions, to be used as qubits, can be stored in a chain along the RF-null, remaining separated from each other due to their mutual Coulomb interaction.

Producing an RF field with the potential curvatures $\beta_x = -\beta_y$ and $\beta_z = 0$ can be achieved by placing four hyperbolic electrodes symmetrically around $r_x = r_y = 0$ as shown in Figure 2.1(a), and extending those electrodes infinitely along the z-axis. Applying an RF voltage to two opposing electrodes produces the desired potential. In practice, designing perfectly hyperbolic electrodes that extend to infinity in all dimensions is impractical. Instead, it is only required to produce an approximate harmonic potential locally around the RF-null, where ions are stored. Simplified electrode structures, such as rods or blades, are therefore commonly used, depicted by the dashed lines in Figure 2.1(a). Since the deviation in an ion's position from the RF-null is, in regular operation, much smaller than the ion-electrode separation, higher order terms of the potential are negligible, and the field around the ions can be assumed to be harmonic. The electrodes extend far enough along the z-axis to ensure that β_z is negligible for the extent of the chain of ions. Specific design of the RF electrodes is based on machinability and required optical access to the trap center.

We have so far covered confinement of ions in two out of three dimensions, the *radial* directions. To confine ions along the *z*-axis, the *axial* direction, electrodes with a static voltage are placed at opposing




Figure 2.1: Overview of linear Paul trap geometry (a) An applied voltage on one of two pairs of hyperbolic electrodes produces a harmonic saddle-potential, indicated by the contour lines. This potential becomes effectively confining when an oscillating voltage (RF) is applied. In reality, electrodes are not hyperbolic in shape, but more blade-like, as indicated by the dashed lines. (b) A positive voltage on endcaps placed on opposing ends of the *z*-axis provides axial confinement.

ends of the trap, symmetrically around the *z*-axis. These electrodes, shown in Figure 2.1(b), are known as end-caps. This type of trap configuration is known as a 3D linear Paul trap. A positive end-cap voltage produces a positive axial curvature α_z . Noting that the sum of the potential curvatures is zero, $\sum_k \alpha_k$, and that by symmetry the radial curvatures are equal, we find that the radial curvatures due to the DC voltage on the endcaps are negative, given by $\alpha_x = \alpha_y = -\alpha_z/2$. One must therefore ensure that the radial confinement produced by the RF electrodes outweighs the anti-confinement of the DC electrodes. This requirement is often met by default, since the axial confinement is intended to be the weakest, such that multiple ions preferentially align themselves along this direction.

We denote the distance between the trap center and one of the end-caps as Z_0 and the distance between the trap center and one of the RF electrodes as R_0 . We can solve for the field curvatures α_k when a voltage $v_{\rm DC}$ is applied to the trap's end-caps, by combining the equation for the static harmonic potential (Eq. 2.1), the restriction $\alpha_x = \alpha_y = -\alpha_z/2$, and the boundary conditions that $V_{\rm DC} = v_{\rm DC}$ at $\vec{r} = (0,0,Z_0)$, and $V_{\rm DC} = 0$ at $\vec{r} = (R_0,0,0)$. The curvatures are then calculated⁵ to be

$$\alpha_z = -2\alpha_x = -2\alpha_y = \frac{4v_{\rm DC}}{R_0^2 + 2Z_0^2}.$$
 (2.13)

Similarly, applying an RF voltage with amplitude v_{RF} to the RF electrodes, with the conditions $V_{RF} = v_{RF} \cos(\Omega_{RF} t)$ at $\vec{r} = (R_0, 0, 0)$ and $V_{RF} = 0$ at $\vec{r} = (0, R_0, 0)$ and $\beta_x = -\beta_y$, and $\beta_z = 0$, we find

$$\beta_x = -\beta_y = \frac{v_{\rm RF}}{R_0^2}.$$
 (2.14)

We express the RF pseudopotential as a sum of effective potential curvatures, $\hat{V}_{RF} = 1/2 \sum_k \hat{\beta_k} r_k^2$. Using the definition of the pseudopotential \hat{V}_{RF} in Eqs. 2.11 and 2.12, the effective curvatures are derived to be

$$\hat{\beta}_z = 0, \quad \hat{\beta}_x = \hat{\beta}_y = \frac{2qv_{\rm RF}^2}{m\Omega_{\rm RF}^2 R_0^4}.$$
 (2.15)

The choice of electrode placement and voltages described above results in potentials that produce identical radial secular frequencies, $\omega_x = \omega_y$. For reliable operation of ions as qubits, this type of degeneracy is undesirable⁶, and can be lifted with a bias potential V_b , introduced by applying a bias

⁵ Often an additional scaling factor, a so-called geometric factor, is introduced to account for the non-hyberbolic shape of the electrodes.

⁶ In particular, Doppler cooling (see Section 2.2.8) becomes inefficient, and state manipulation through motional modes (see Section 2.2.6) becomes unreliable.

DC voltage v_b to the two opposing blades that do not carry an RF voltage. Similar analysis as above reveals that this bias voltage produces potential curvatures of $\pm 2v_b/R_0^2$ in the two radial directions. The full potential, using the pseudopotential approximation, is given by:

$$V_{\text{Tot}} = V_{\text{DC}} + V_{\text{RF}} + V_{\text{b}} \tag{2.16}$$

with

$$V_{DC} = \frac{v_{DC}}{R_0^2 + 2Z_0^2} (2r_z^2 - r_x^2 - r_y^2)$$

$$V_{RF} = \frac{qv_{RF}^2}{m\Omega_{RF}^2 R_0^4} (r_x^2 + r_y^2)$$

$$V_b = \frac{v_b}{R_0^2} (r_y^2 - r_x^2)$$
(2.17)

2.1.3 Segmented surface traps

The previous section has introduced a general trap architecture which produces potentials that confine ions in all three dimensions. For use in quantum computation, potentials are set such that multiple co-trapped ions align themselves in a chain along the RF-null, separated by Coulomb repulsion. We continue the discussion of ion positions and their motional modes in Section 2.1.4, and then how trapped ions are used as qubits in Section 2.2. However, while on the topic of trap geometry, and having introduced the basics of linear trap potentials, we take a small detour to discuss why and how one would migrate from a blade trap architecture, such as the one discussed in the previous section, towards what is considered a more scalable trapping solution, the linear *segmented surface* trap.

The problem of scalability boils down to increasing the number of trapped ions to be used as qubits. The macroscopic linear trap described in the previous section creates a single trapping region, in which ions are stored in a chain along the RF-null. Tens of ions can be confined in such a region, but such a trap is hard-pressed to introduce more ions while maintaining a stable linear chain [57], in which a one-dimensional chain buckles into a second axis. Keeping a linear chain at higher ion numbers requires either a higher radial confinement or a lower axial confinement, and both options come with technical complications. Motional modes (discussed in the following section) grow with ion number leading to so-called spectral mode crowding. Since shared modes form the basis for ion-ion communication, it is important that they remain spectrally well-separated.

One solution is to split up the trap geometry such that multiple separate trapping regions are formed [58, 59]. A small number of ions, typically single digits, can then occupy each trapping site. Such segmentation involves subdividing the non-RF-carrying blades into electrically isolated sections. Each section is an electrode whose voltage can be controlled independently. Applying appropriate voltages to a set of such electrodes produces multiple minima in the potential along the RF-null, with each minimum a region where ions are trapped. Electrode voltages can be independently adjusted to precisely control the trapping potential around each ion chain. They can be changed over time, allowing the physical positions of ion chains to be manipulated, giving rise to ion transport operations, the topic of Section 5.

Another step towards scalability lies in the ability to utilize not just a single one-dimensional RF-null as a channel to store ions, but to create an interconnected network of such channels [60–62]. Implementing such a network into a three-dimensional blade trap is technically challenging and poses limits in optical access. Instead, a two-dimensional trap, known as a surface trap, can be used, which hosts an interconnected set of rails that serve as channels over which ions can be transported. This concept follows the Quantum Charge-Coupled Device (QCCD) proposal [24].

Just as with macroscopic blade traps, a multitude of surface trap architectures exist. A commonly used design is shown schematically in Figure 2.2(a). It's functionality can be intuitively understood when viewing its electrodes as projections (with a bit of imagination) from the 3D design discussed in the previous section. Two rails carry RF voltage, and produce the radial pseudopotential, shown in Figure 2.2(b). The rest of the surface is segmented such that a supply of DC voltages can be applied

to produce an axial 'end-cap-like' confinement. In the example in Figure 2.2(a), this corresponds to applying a positive voltage on the corner electrodes.




Figure 2.2: (a) Example layout of a linear surface trap, which can be seen as a projection from the linear blade trap. Similar to the 3D blade trap, a surface trap produces a saddle potential, marked by a cross in the contour plot in (b).

The larger number of DC electrodes in a surface trap compared the non-segmented 3D trap introduces more degrees of freedom in producing a potential $V_{\rm DC}$ around the trapping region(s). It is therefore not as straightforward to produce trapping potentials similar to those as discussed earlier for the 3D trap. Our method of determining electrode voltages that produce the desired potentials is covered in Section 2.3, which also discusses characterization of surface traps. First, however, we discuss what "desired potentials" means, in terms of the motion of ions.

2.1.4 *Ions in a linear Paul trap - positions and modes*

In this section, we give an overview of ions positions and their secular motion in a linear trap. For this analysis, the pseudopotential approximation is used, and it is assumed that ions have low enough kinetic energy that their average displacement is much smaller than the mean distance between ions, such that they form a chain.

The total potential energy E_{pot} of a trapped ion crystal is given by the sum of the trap's potential energy E_{trap} and the interaction energy of the ions E_{Coul} ,

$$E_{\rm pot} = E_{\rm trap} + E_{\rm Coul} \tag{2.18}$$

The Coulomb energy for N ions of charge q_i and position $\vec{r}^{(i)}$ is

$$E_{\text{Coul}} = \frac{1}{4\pi\epsilon_0} \sum_{i}^{N} \sum_{j>i} \frac{q_i q_j}{\left| \vec{r}^{(j)} - \vec{r}^{(i)} \right|}.$$
 (2.19)

with ϵ_0 the vacuum permittivity, and i and j indices of the N ions.

The total potential energy in the trap is the sum of the potential energies of each ion, $E_{\text{trap}} = \sum_i q_i V_{\text{Tot}}(\vec{r}^{(i)})$. The potential is given by the sum of the DC potential V_{DC} and the RF pseudopotential V_{RF} , where we from here on omit the tilde notation that previously distinguished the pseudopotential from the time-dependent RF potential. Before it was assumed that both the RF and DC potentials had only quadratic contributions, both with principal axes in the same $\{x,y,z\}$ basis. Now, for completeness, we include the possibility that principal axes don't necessarily follow this basis, and that the potential minimum is not centered around zero (i.e., there exists a linear component to the field potential).

A notationally convenient way to express harmonic potentials is with the 3×3 curvature matrices, $\Phi_{RF}^{(i)}$ and Φ_{DC} , which are the Hessian of the potentials V_{RF} and V_{DC} . The Hessian of the RF pseudopo-

tential is indexed (*i*), since its value is dependent on the mass m_i and charge q_i of ion *i*. The potential V_E , originating from an undesired homogeneous electric field, can be expressed by its electric field vector, $\vec{E} = -\nabla V_E$.

The trap potential, considered to be harmonic (i.e. no third or higher order polynomial terms) is given by

$$E_{\text{trap}} = \frac{1}{2} \sum_{i}^{N} q_{i} \vec{r}^{(i)T} \Phi_{RF}^{(i)} \vec{r}^{(i)} + \frac{1}{2} \sum_{i}^{N} q_{i} \vec{r}^{(i)T} \Phi_{DC} \vec{r}^{(i)} - \sum_{i}^{N} q_{i} \vec{E} \cdot \vec{r}^{(i)},$$
 (2.20)

with the superscript T indicating the vector transpose. In order to satisfy Maxwell's equation $\nabla^2 V_{DC} = 0$ for static fields, it is required that Φ_{DC} satisfies $tr\{\Phi_{DC}\}=0$. The RF pseudopotential does not obey $\nabla^2\Phi_{RF}=0$. In a linear trap, it is assumed that there is one dimension along which there is no RF curvature, the trap's axial axis. The coordinate basis can be chosen such that the pseudopotential's curvature matrix is given by

$$\Phi_{\rm RF}^{(i)} = \frac{q_i}{4m^{(i)}\Omega_{\rm RF}^2} \left| \nabla V_{\rm RF} \right|_{min}^2 \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{2.21}$$

We denote the factor in front of the matrix as $v_{RF}^{(i)}$.

Grouping all individual ion positions $\vec{r}^{(i)}$ into a single vector \vec{r} of length 3N, the equilibrium positions \vec{r}_0 of all ions can be determined by solving the set of differential equations

$$\frac{\partial E_{\text{pot}}}{\partial \vec{r}} = \vec{0}. \tag{2.22}$$

Motional modes and their frequencies $\omega_{\{x,y,z\}}^{(i)}$ can be obtained by finding the eigenvectors and eigenvalues of the mass-corrected Hessian matrix of E_{pot} , with masses m_i , given by the terms

$$H_{ij} = \frac{1}{\sqrt{m_i m_j}} \left. \frac{\partial^2 E_{\text{pot}}}{\partial r_i \partial r_j} \right|_{\vec{r}_0}, \tag{2.23}$$

linearized around \vec{r}_0 . Stable trapping potentials require real positive eigenvalues e_k , and have motional mode frequencies given by $\sqrt{e_kq}$.

For a single trapped ion, the position and motional modes are analytically obtainable. If one can assume that the principal axes of the DC potential are the same as those of the RF pseudopotential, then Φ_{DC} does not contain off-diagonal elements⁷. Under this assumption, solving Eq. 2.22, the ion's position is given by the terms

$$r_k = E_k / \left(\Phi_{\mathrm{DC}}^{(kk)} + \Phi_{\mathrm{RF}}^{(kk)} \right) \tag{2.24}$$

and, solving Eq. 2.23, motional mode frequencies are

$$\omega_k = \sqrt{q(\Phi_{\rm DC}^{(kk)} + \Phi_{\rm RF}^{(kk)})/m},$$
(2.25)

where the superscripts denote the element of the respective matrices. In the coordinate system used here, ω_z is an axial frequency, and ω_x and ω_y are radial frequencies. In the absence of a homogeneous electric field \vec{E} , a single ion is located at the 'center' of the trap, $\vec{r} = \vec{0}$. The motional frequencies are unaffected by \vec{E} .

⁷ For many 'standard' trap designs, such as the ones shown in Figure 2.1 and 2.2, this is a reasonable assumption

The positions and motional frequencies of two trapped ions are also analytically attainable. The mean position of both ions $r_k^{(m)}$ is the same as the single-ion case. Under the assumption that the confinement along the *z*-axis is the weakest, ions align along this axis, and are placed at [63]

$$r_z^{(i)} = r_z^{(m)} + \frac{1}{2} \sqrt[3]{\frac{q}{2\pi\epsilon_0 \Phi_{DC}^{(zz)}}} (-1)^i.$$
 (2.26)

For ions i = 1, 2. This solution is also applicable if the masses of the two ions differ, as long as there is no homogeneous electric field in the radial directions, $E_x = E_y = 0$.

Two ions have a total of 6 motional modes. In each dimension, they can oscillate in phase with each other, in which case their motion is known as the *common mode* or *center-of-mass mode*, or out of phase with each other, known as the *stretch mode*. If ion masses are identical, common mode frequencies in each dimension $\omega_{k,\text{com}}$ are identical to the previously defined single-ion frequencies, ω_k . Stretch mode frequencies are given by [64, 65].

$$\omega_{x,\text{str}} = \sqrt{\omega_{x,\text{com}}^2 - \omega_{z,\text{com}}^2}$$
 (2.27)

$$\omega_{y,\text{str}} = \sqrt{\omega_{y,\text{com}}^2 - \omega_{z,\text{com}}^2}$$
 (2.28)

$$\omega_{z,\text{str}} = \sqrt{3}\omega_{z,\text{com}}.$$
 (2.29)

Analytic solutions for common and stretch mode frequencies also exist if the masses of the two ions are not identical [65]. Three same-mass ions also have analytically obtainable positions and motional frequencies [64], including that of a third type of mode, *Egyptian*, in which the central ion moves out of phase with the outer ions. The analytic solutions of two- and three-ion positions and frequencies require that off-diagonal terms in Φ_{DC} are zero, and that no higher-order potentials [66] are present. For all other cases, including larger numbers of ions, positions and mode frequencies can only be obtained through numerical calculation or from analytical approximations [64].

The axial mode structures and frequencies of trapped 40 Ca $^+$ and 88 Sr $^+$ ions in various configurations are shown in Figure 2.3, which reflects the relevant ion constitutions of works presented in this thesis. The values in the top right of each window is the motional frequency, expressed relative to the motional mode of a single 40 Ca $^+$ ion, where the same confining potentials are used for all examples. The values below the arrows, denoted with $\tilde{\eta}_z$ indicate a dimensionless coupling strength between an applied light field and the motional mode, through each ion, and is expressed relative to that of a single 40 Ca $^+$ ion. The significance and interpretation of this value is discussed in Section 2.2.6, where it is referred to as the Lamb-Dicke parameter, after having introduced light-atom interaction. For now, it can be loosely interpreted as the relative amplitude of each ion in a particular oscillator mode.

	Ca	Sr	Ca - Ca	Ca - Sr	Ca - Sr - Ca
Common	$\tilde{\omega}_z = 1$ $\tilde{\eta}_z = 1$	0.67 0.88	0.71 0.71	0.77 0.49 0.75	0.83 0.48 0.69 0.48
			1.73	1.53	1.56
Stretch			←⊕ ⊕ ≻	← →	←
			-0.54 0.54	-0.73 0.25	-0.53 0 0.53
					1.72
Egyptian					
					0.40 -0.33 0.40

Figure 2.3: Various configurations of axial modes of 40 Ca $^{+}$ and 88 Sr $^{+}$ ion chains. The upper values in each box are mode frequencies, scaled in reference to that of a single 40 Ca $^{+}$ ion. The lower value is the scaled Lamb-Dicke parameter, $\tilde{\eta}_z$, described in Section 2.2.6

As is described in the following section, ions' motional modes play an important role in quantum computation, since they enable interaction between qubits. For this purpose, ions should have a low enough kinetic energy that their motion is considered in terms of quanta of motion, *phonons* [67]. An

ion's motional state is expressed in terms of the occupation number n of each mode l, $|n\rangle_l$ and can be found in a superposition or mixed state of occupation numbers.

Manipulation of these modes is typically expressed in terms of the creation and annihilation operators \hat{a}_{l}^{\dagger} and \hat{a}_{l} for each mode l, with the properties

$$\hat{a}_l^{\dagger} |n\rangle_l = \sqrt{n+1} |n+1\rangle_l \tag{2.30}$$

$$\hat{a}_l |n\rangle_l = \sqrt{n} |n-1\rangle_l \tag{2.31}$$

$$\hat{a}_{1}^{\dagger}\hat{a}_{1}\left|n\right\rangle_{1}=n\left|n\right\rangle_{1}\tag{2.32}$$

Using these operators, the Hamiltonian that describes the energy of an ions motional mode l is given by

$$H_{\text{osc}}^{(l)} = \hbar \omega_l (\hat{a}_l^{\dagger} \hat{a}_l + \frac{1}{2})$$
 (2.33)

with \hbar the reduced Planck constant.




Figure 2.4: Bloch sphere representation of a qubit state

2.2 TRAPPED IONS AS QUBITS

2.2.1 Qubits

Trapped ions serve as the carriers of quantum information, *qubits*. A qubit is stored in two of an ion's electronic states. These states have the requirement that one should have the means to manipulate and measure their occupation in a controlled way[68]. Correspondingly, the qubit must have lifetimes much longer than the timescales associated with state control and detection. The internal level structure of the ionic species used in our experiment are discussed in Section 2.2.7, but for now we assume an ion with two electronic levels that meets these requirements. We denote these states of the qubit with $|0\rangle$ and $|1\rangle$, being the ground and excited state, respectively. An arbitrary pure quantum state $|\psi\rangle$ can exist in a superposition of these basis states

$$|\psi\rangle = c_0 |0\rangle + c_1 |1\rangle \tag{2.34}$$

where c_0 and c_1 are complex numbers that obey $|c_0|^2 + |c_1|^2 = 1$. We assume that spontaneous decay from the excited state $|1\rangle$ to the ground state $|0\rangle$ can be neglected. The Hamiltonian of this two-level system with the basis $\{|0\rangle, |1\rangle\}$ is given by

$$H_{\rm a} = -\frac{\hbar\omega_0}{2}\sigma_z \tag{2.35}$$

where $\hbar\omega_0$ is the energy difference between the states, and σ_z one of the Pauli matrices, which are given by

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$
 (2.36)

A convenient method of representing a qubit state is by means of the so called Bloch-sphere [11], depicted in Figure 2.4. Any possible combinations of the complex values c_0 and c_1 (adhering to the normalization criterion) can be uniquely represented by a vector whose endpoint lies on the surface of a sphere with radius 1. We take the Bloch-vector pointing to the southern-most point of this sphere to represent $|c_0|=1$, and to the northern-most point to be $|c_1|=1$ 8. Any intermediate point represents a superposition of $|0\rangle$ and $|1\rangle$. Notably, points on the 'equator' are an *equal* superposition, $1/\sqrt{2}(|0\rangle+e^{i\phi}|1\rangle)$. Here, ϕ is a phase difference between the states, which sets the angle along the equatorial circle on which that state is mapped.

In QC experiments, operations are applied to the qubit to induce controlled changes of the state $|\psi\rangle$. A basic unitary operation acting on the qubit state, $U = \exp(-i\sigma_k\theta/2)$, is represented by a rotation of the Bloch-vector around the axis $k \in \{x, y, z\}$, by an angle θ . The z-axis is the line that passes through the north and south pole, and the x and y-axes lie in the equatorial plane.

2.2.2 Ion - light interaction

We now discuss the controlled manipulation of a qubit state, which is achieved by coupling the two-level ion with an electromagnetic light field. We take this field to be a plane wave whose electric field at a given position, at time *t*, is given by

$$E(t) = E_0 \cos\left(w_f t + \phi_f\right) \tag{2.37}$$

with E_0 the field amplitude, ω_f the frequency, and ϕ_f the field's phase.

The two-level system, with a bare-state Hamiltonian H_a , is perturbed with an interaction Hamiltonian H_l , given by

$$H_l = \hbar \Omega \sigma_x \cos(\omega_f t + \phi_f) \tag{2.38}$$

with Ω the so-called Rabi frequency, whose value is dependent on the coupling strength between the ion's dipole moment and the applied field, and is proportional to the field amplitude E_0 .

The total system Hamiltonian, $H = H_a + H_l$, can be transformed to an interaction picture with respect to H_a , yielding

$$H_{\text{int}} = \exp\left(iH_{a}t/\hbar\right)H\exp\left(-iH_{a}t/\hbar\right)$$

$$= \frac{\hbar\Omega}{2}\left(\cos\left((\omega_{f} - \omega_{0})t + \phi_{f}\right)\sigma_{x} + \cos\left((\omega_{f} + \omega_{0})t + \phi_{f}\right)\sigma_{x} \dots\right)$$

$$+ \sin\left((\omega_{f} - \omega_{0})t + \phi_{f}\right)\sigma_{y} + \sin\left((\omega_{f} + \omega_{0})t + \phi_{f}\right)\sigma_{y}. \tag{2.39}$$

The rotating wave approximation (RWA) allows us to drop terms that contain fast oscillations, $\omega_f + \omega_0$. With the substitution $\Delta = \omega_f - \omega_0$, where Δ is the detuning between the light-field and the ion's transition frequency, the interaction Hamiltonian becomes

$$H_{\text{int}} = \frac{\hbar\Omega}{2} \left(e^{-i(\Delta t - \phi_f)} \sigma_+ + e^{i(\Delta t + \phi_f)} \sigma_- \right) \tag{2.40}$$

where we use the electronic raising and lowering operators $\sigma_{\pm} = (\sigma_x + i\sigma_y)/2$.

Another rotating frame transformation, with $U = \exp(-i\Delta t \sigma_z/2)$ allows us to to express H_{int} as a time-independent operator, which can be written in matrix form as

$$H_{\rm int} = \frac{1}{2}\hbar \begin{pmatrix} -\Delta & \Omega e^{-i\phi_f} \\ \Omega e^{i\phi_f} & \Delta \end{pmatrix}. \tag{2.41}$$

2.2.3 Rabi oscillations

The time-dependent dynamics of the quantum state $|\psi(t)\rangle$ can be calculated using the Schrödinger equation:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H_{\rm int} |\psi(t)\rangle.$$
 (2.42)




Figure 2.5: Bloch sphere representation and excitation plots for various light-atom interactions. (a) A light field on resonance with the atomic transition cyclically transfers population between $|0\rangle$ and $|1\rangle$. Applying the light field for a time $t=\pi/\Omega$ is known as a π -pulse, as depicted by the half rotation in the Bloch sphere, and the gray dashed line in the excitation plot. (b) Full population transfer is does not occur when the beam is detuned. (c) Electronic population does not change at very large detuning, but a phase shift between the states does occur.

Let us consider the case that the light field is on resonance with the atomic transition frequency, $\Delta=0$, and, for notational convenience set the light field phase to $\phi_f=0$. If the initial state is $|\psi(0)\rangle=|0\rangle$, solving Eq. 2.42 results in

$$|\psi(t)\rangle = c_0(t)|0\rangle + c_1(t)|1\rangle = \cos\left(\frac{\Omega t}{2}\right)|0\rangle - i\sin\left(\frac{\Omega t}{2}\right)|1\rangle.$$
 (2.43)

The probability of an ion being in the excited state $|1\rangle$ is given by

$$P_{|1\rangle} = |c_1(t)|^2 = \sin^2\left(\frac{\Omega t}{2}\right).$$
 (2.44)

This oscillatory behavior is known as a Rabi cycle, with Rabi frequency Ω . The Rabi frequency is proportional to the square root of the light field's intensity, and is additionally dependent on characteristics of the transition being probed (such as magnetic quantum numbers), and the polarization and orientation of the light field [69].

This type of interaction, resonant Rabi oscillations, is a crucial component of quantum computation, as it represents one of the basic mechanisms of control of a qubit's state. Applying a light field on resonance with an ion's electronic transition for a time $t = \pi/\Omega$, fully transfers population from the state $|0\rangle$ to $|1\rangle$, and vice versa. In the Bloch-sphere picture, this operation can be viewed as a rotation of the Bloch-vector around the *x*-axis, by an angle π , shown in Figure 2.5(a). The operation is therefore known as a π -pulse. Applying this field for half of that time, $t = \pi/(2\Omega)$, allows us to generate

superposition states, $|\phi\rangle = \frac{1}{\sqrt{2}}(|0\rangle - i|1\rangle)$. As one might be able to guess, this operation is referred to as a $\pi/2$ -pulse. Collectively, operations that rotate the Bloch-vector about a fixed axis are known as qubit rotations⁹.

Now let us consider what happens when the applied field is near, but not on, resonance with the atomic transition, $\Delta \neq 0$. The states now evolve as

$$c_0(t) = \cos\left(\frac{\Omega_{\rm eff}t}{2}\right) + i\frac{\Delta}{\Omega_{\rm eff}}\sin\left(\frac{\Omega_{\rm eff}t}{2}\right) \tag{2.45}$$

$$c_1(t) = -i\frac{\Omega}{\Omega_{\text{eff}}} \sin\left(\frac{\Omega_{\text{eff}}t}{2}\right) \tag{2.46}$$

with an effective Rabi frequency $\Omega_{eff}=\sqrt{\Omega^2+\Delta^2}.$ The population of the excited state is given by

$$P_{|1\rangle} = \frac{\Omega^2}{\Omega_{\rm eff}^2} \sin^2\left(\frac{\Omega_{\rm eff}t}{2}\right).$$
 (2.47)

This population transfer is shown in Figure 2.5(b), again for a pulse duration of $t = \pi/\Omega$, but with a detuning of $\Delta = -0.6\Omega$. Population transfer still exhibits oscillatory behavior, as in the resonant case, but Eq. 2.47 shows that the excited state will never be fully populated. Additionally, the effective cycling frequency is increased when the light-field is detuned from resonance.

In trapped-ion experiments with qubits, one typically finds the frequency at which a light field is resonant with an ion's transitions by applying light pulses of fixed duration for varying light frequencies. The resulting excitation spectrum, shown in Figure 2.6 for $t = \pi/\Omega$, is a simple and effective tool to help calibrate the laser frequency to be on resonance.




Figure 2.6: Rabi excitation spectrum as a function of detuning Δ , with a pulse duration $t = \pi/\Omega$

2.2.4 AC Stark shift

We can consider the case where the detuning is much larger than the on-resonance Rabi frequency, $\Delta \ll \Omega$. As can be seen from Eqs. 2.47, there is hardly any excitation from the ground state, as $\Omega^2/\Omega_{\rm eff}^2$ tends to zero. However, the energy levels still have a slight energy change with respect to the case where $\Omega=0$. This shift in energy, known as *AC Stark shift*, results in a relative phase change of the states $|0\rangle$ and $|1\rangle$ [70].

The change of phase over time, the AC Stark shift frequency, can be calculated by looking at the shift in eigenvalues, thus energies, of the Hamiltonian in Eq. 2.41, as done in [71]. Another instructive way to calculate this shift is by monitoring what happens to the states of $|0\rangle$ and $|1\rangle$, when initially placed in a superposition $c_0 = c_1 = 1/\sqrt{2}$ before switching on the detuned beam.

⁹ Not to be confused with ion crystal rotations in Chapter 5

In this case, the Schrödinger equation gives us

$$c_0(t) = \frac{1}{\sqrt{2}}\cos\left(\frac{\Omega_{\text{eff}}t}{2}\right) + i\frac{\Omega - \Delta}{\sqrt{2}\Omega_{\text{eff}}}\sin\left(\frac{\Omega_{\text{eff}}t}{2}\right)$$
(2.48)

$$c_1(t) = \frac{1}{\sqrt{2}}\cos\left(\frac{\Omega_{\rm eff}t}{2}\right) + i\frac{\Omega + \Delta}{\sqrt{2}\Omega_{\rm eff}}\sin\left(\frac{\Omega_{\rm eff}t}{2}\right). \tag{2.49}$$

For a small Rabi frequency, $\Omega \ll \Delta$, the populations in each state, $|c_0(t)|^2$ and $|c_1(t)|^2$, hardly change. However, the two states do acquire a non-negligible phase difference with respect to each other: The phase $\phi_{|0\rangle}$ and $\phi_{|1\rangle}$ of the states $|0\rangle$ and $|1\rangle$ are given by

$$\phi_{|i\rangle} = \tan^{-1}\left(\frac{\operatorname{Im}(c_i)}{\operatorname{Re}(c_i)}\right) = \tan^{-1}\left(\frac{\Omega \pm \Delta}{\sqrt{\Delta^2 + \Omega^2}} \tan\left(\frac{\sqrt{\Delta^2 + \Omega^2}t}{2}\right)\right) \tag{2.50}$$

with a minus sign for $|0\rangle$ and a plus sign for $|1\rangle$. The accumulated phase difference is $\partial \phi = \phi_{|1\rangle} - \phi_{|0\rangle}$. Since Ω is small, we can do a Taylor expansion of $\partial \phi$ around $\Omega = 0$. The leading term scales with Ω^2 , and is given by

$$\partial \phi pprox rac{2\Delta t - 4\sin(\Delta t) + \sin(2\Delta t)}{4\Delta^2}\Omega^2$$
 (2.51)

The mean AC Stark shift frequency between the states Δ_{AC} is the average phase change over time, and is thus

$$\Delta_{\rm AC} = \lim_{T \to \infty} \frac{\partial \phi(T) - \partial \phi(0)}{T} = \frac{\Omega^2}{2\Lambda}.$$
 (2.52)

The AC Stark shift is another fundamental tool (though occasionally also a hindrance) in quantum computation. Applying a light field far enough off-resonance from a transition allows us to change the phase between qubits, without altering their populations. In the Bloch-sphere picture, this operation corresponds to a rotation of the Bloch-vector around the *z*-axis, and is commonly referred to as a *z*-gate. In such an operation, a state in an equal superposition traverses the equator of the Bloch-sphere, as shown in Figure 2.5(c).

2.2.5 Decoherence

All previous descriptions of light-matter interaction have assumed perfect conditions: the laser frequency and power are perfectly stable, the ion's two-level transition frequency is constant and it's excited state does not spontaneously decay. A theorist might happily make such assumptions, but an experimentalist may not easily get away with that mentality. Imperfections in an experimental setup disrupt the intended qubit manipulation, leading to experimental errors. Imperfections can be categorized into coherent and incoherent sources of error.

Unintentional, but otherwise stable, offsets in experimental parameters lead to *coherent* errors. Qubit operations still consistently rotate any Bloch vector from one point to another on the surface of the Bloch sphere, but the intended mapping might not be achieved. The qubit remains representable as $c_0 |0\rangle + c_1 |1\rangle$ as in Eq. 2.34, though the values of c_0 and c_1 may deviate from the intended output.

Experimental parameters may unintentionally change during the execution of a quantum algorithm, leading to *incoherent* errors, resulting in qubit decoherence. In the presence of noisy experimental parameters, if one were to repeatedly initialize a qubit and execute a set of quantum operations, the final position of qubit's Bloch vector would be at a different point on the surface of the Bloch sphere on each run. Executing multiple experimental runs to determine the qubit state effectively averages out the different Bloch vectors. The resulting Bloch vector does not lie on the surface of the Bloch sphere, but somewhere within it. Similarly, interaction between the qubit and the environment leads to dissipative loss of qubit information, also represented by a reduction of the Bloch vector length. In such cases, the qubit cannot be expressed by the coefficients c_0 and c_1 alone, and is said to be in a

mixed state. It is instead expressed by a property that captures the mean of output qubit states in the ensemble of experimental runs, the density operator ρ . The density operator is defined in terms of the statistical probability p_i to produce the state $|\phi_i\rangle$ from an ensemble of possible states, and is given by

$$\rho = \sum_{j} p_{j} \left| \phi_{j} \right\rangle \left\langle \phi_{j} \right|. \tag{2.53}$$

To summarize, and introduce/reiterate some useful terms:

- A qubit state whose Bloch vector lies on the surface of the Bloch sphere is known as a *pure state*, which can be expressed as $|\psi\rangle = c_0 |0\rangle + c_1 |1\rangle$ with $|c_0|^2 + |c_1|^2 = 1$, and has a density operator given by $\rho = |\psi\rangle \langle \psi|$. More generally, if the system contains more than just one qubit, a pure quantum state is given by $|\psi\rangle = \sum_i c_i |i\rangle$, with i the index of the state, which lie in the set of all possible terms of the tensor product of the individual qubits. As before, the weight-terms are normalized, $\sum_i |c_i| = 1$. One could describe the Bloch vector as lying on the surface of a Bloch hypersphere, but maintaining an intuitive picture gets harder in more than three dimensions.
- If a qubit state is not pure, and is referred to as a *mixed state*, and is depicted by an average Bloch vector that lies within the volume of the Bloch sphere. This state is described by the *density* operator ρ , as in Eq. 2.53.
- The *purity* of a qubit state is a quantitative description of how close a Bloch vector is to the surface of a Bloch sphere, given by $\gamma = \text{tr}(\rho^2)$. A pure state has $\gamma = 1$. A maximally mixed state (i.e. a Bloch vector of zero length, $\vec{0}$) of a system with d states has a purity of $\gamma = 1/d$.
- *Coherence* is a qualitative property of a qubit that describes how well it can maintain a superposition. Decoherence is the process of coherence loss, and is paired with a reduction in purity.
- The Schrödinger equation (Eq. 2.42) cannot be used to describe the dynamics of a state that undergoes mixing. Instead, the dynamics of ρ can be obtained by using the Lindblad master equation:

$$\frac{d\rho}{dt} = -i\left[H,\rho\right] + \mathcal{L}(\rho) \tag{2.54}$$

 $\mathcal{L}(\rho)$ is known as the Lindblad decoherence term [72], a Liouvillian operator whose mathematical description is outside the scope of this thesis. This term describes the dynamics that lead to loss of purity, including dissipation and decoherence.

2.2.6 Light-atom interaction of a harmonically confined ion

In Section 2.1.4 we have discussed the motional modes of trapped ions, introduced the harmonic oscillator ladder operators \hat{a}_l^{\dagger} and \hat{a}_l of mode l, and introduced the oscillator's Hamiltonian, $H_{\rm osc}^{(l)}$. In this section, we extend the light-atom interaction analysis of the previous section, by additionally considering the ion's motion.

The interaction Hamiltonian was previously defined in Eq. 2.38. We can now include the ion's motion with respect to the laser field, by adapting our definition of the interaction Hamiltonian to include the ion's position:

$$H_{\rm int} = \hbar \Omega \sigma_x \cos(\hat{r}k + \omega_f t) \tag{2.55}$$

where $\hat{r} = \sqrt{\hbar/(2m\omega_l)}(\hat{a}_l^{\dagger} + \hat{a}_l)$ is the ion's position operator, and k is the projection of the laser field's wavevector \vec{k} onto the unit vector of the axis of ion motion \hat{e}_r (i.e. $k = \vec{k} \cdot \hat{e}_r$). The position term in the exponent $k\hat{r}$ can be written as $\eta(\hat{a}_l^{\dagger} + \hat{a}_l)$ using the substitution

$$\eta = k\sqrt{\frac{\hbar}{2m\omega_l}}.$$
 (2.56)

The value η is known as the Lamb-Dicke parameter. It relates the extent of the ion's ground state wavepacket to the wavelength of the light field.

The perturbed Hamiltonian is now adapted to

$$H_{l} = \frac{\hbar\Omega}{2}\sigma_{x}\left[\exp\left(i(\eta(\hat{a}_{l}^{\dagger} + \hat{a}_{l}) + \omega_{f}t)\right) + \exp\left(-i(\eta(\hat{a}_{l}^{\dagger} + \hat{a}_{l}) + \omega_{f}t)\right)\right]$$
(2.57)

This Hamiltonian's frame can be changed into an interaction picture, as $H_{\text{int}} = UHU^{\dagger}$ with the unitary operator $U = \exp(-i(H_a + H_{\text{osc}})t/\hbar)$. The interaction Hamiltonian then reads

$$H_{\text{int}} = \frac{\hbar\Omega}{2}\hat{\sigma}_{+} \left[\exp\left(i\eta(\hat{a}_{l}e^{-i\omega_{l}t} + \hat{a}_{l}^{\dagger}e^{i\omega_{l}t})\right) \right] \exp\left(-i\Delta t\right) + h.c.$$
 (2.58)

where, as in the previous section, we have made the RWA, and used the substitution $\Delta = \omega_f - \omega_0$. The interaction Hamiltonian can be further simplified under the assumption that the ion's wavepacket is much smaller than the wavelength of the light field, formally defined by the condition $\eta^2(2n+1) \ll 1$. This condition is the so-called Lamb-Dicke regime, and allows us to truncate the Taylor expansion,

$$\exp\left(i\eta(\hat{a}_{l}+\hat{a}_{l}^{\dagger})\right)) \approx 1 + i\eta(\hat{a}_{l}+\hat{a}_{l}^{\dagger}) - \frac{\eta^{2}}{2}(\hat{a}_{l}+\hat{a}_{l}^{\dagger})^{2}. \tag{2.59}$$

We apply this approximation to the interaction Hamiltonian, and expand it into three parts, $H_{int} = H_{car} + H_{rsb} + H_{bsb}$, with

$$H_{\text{car}} = \frac{\hbar\Omega}{2} (1 - \eta^2 n) \left(\hat{\sigma}_+ \exp(-i\Delta t) + \hat{\sigma}_- \exp(i\Delta t) \right)$$
 (2.60)

$$H_{\rm rsb} = \frac{\hbar \Omega \eta}{2} \left(\hat{\sigma}_{+} \hat{a}_{l}(0) \exp(-i(\Delta + \omega_{l})) + \hat{\sigma}_{-} \hat{a}_{l}^{\dagger}(0) \exp(i(\Delta + \omega_{l})) \right)$$
(2.61)

$$H_{\text{bsb}} = \frac{\hbar \Omega \eta}{2} \left(\hat{\sigma}_{+} \hat{a}_{l}^{\dagger}(0) \exp(-i(\Delta - \omega_{l})) + \hat{\sigma}_{-} \hat{a}_{l}(0) \exp(i(\Delta - \omega_{l})) \right)$$
(2.62)

where we have used $n = \hat{a}_{l}^{\dagger} \hat{a}_{l}$, and omit terms in H_{rsb} and H_{bsb} that scale with η^{2} .

We see three distinct cases of resonance $\Delta=0$, $\Delta=-\omega_l$, and $\Delta=+\omega_l$. In the first case, the dynamics of the ion's state are dominated by $H_{\rm car}$ and follow exactly the on-resonance Rabi cycles described in the previous section. There is no change of motional quanta at this detuning. This type of oscillation is known as a *carrier* excitation. The Rabi frequency depends on the mode occupation n, and is given by

$$\Omega_{n,n} = \Omega(1 - \eta^2 n) \tag{2.63}$$

If $\Delta=-\omega_l$, the light field is on resonance with a so-called *red sideband* transition. A transfer from the ions ground to excited state $\hat{\sigma}_+ \ket{0} = \ket{1}$ is paired with a removal of a phonon in mode l, \hat{a}_l . This exchange between electronic and motional excitation, with dynamics following from the Hamiltonian $H \propto (\hat{\sigma}_+ \hat{a} + \hat{\sigma}_- \hat{a}^\dagger)$, is known as the Jaynes-Cummings model [73]. Noting that $\hat{a}_l \ket{n}_l = \sqrt{n} \ket{n-1}_l$, we see from $H_{\rm rsb}$ that the red sideband transition from $\ket{0} \ket{n}_l \leftrightarrow \ket{1} \ket{n-1}_l$ occurs with a coupling strength of

$$\Omega_{n,n-1} = \Omega \eta \sqrt{n} \tag{2.64}$$

Similarly, if $\Delta = \omega_l$, the field is on resonance with a *blue sideband*, in which case an excitation of the electronic state is paired with the gain of a phonon¹⁰. The transition $|0\rangle |n\rangle_l \leftrightarrow |1\rangle |n+1\rangle_l$ has coupling strength

$$\Omega_{n,n+1} = \Omega \eta \sqrt{n+1} \tag{2.65}$$

The ability to control the motional state of an ion or multiple ions is a fundamental tool in quantum computation. Firstly, as described in Section 2.2.8, the process of *sideband cooling*, makes use of sideband transitions (as the name suggests) to bring ions close to the motional ground state, $|0\rangle_l$, where an ion's thermal energy does not negatively influence the outcome of qubit operations. Secondly, since the quantized states of motion are shared among multiple ions in a chain, these motional modes serve as a data bus to transfer information between ions, or to entangle them.

2.2.7 Qubit operations of Group II ions

The previous section introduced how light fields are used to manipulate a two-level system in a harmonic oscillator. In this section we look at the practical implementation of this type of matter-light interaction in our experiment. The experiments described in this thesis use 40 Ca⁺ and 88 Sr⁺ ions as information carriers, or qubits. Both species are group II ions (alkaline earth metals), and share similar level structures. In particular, for 40 Ca⁺(88 Sr⁺), the meta-stable $3D_{5/2}(4D_{5/2})$ level has an optical quadrupole transition to the $4S_{1/2}(5S_{1/2})$ ground state. The narrow linewidth, long lifetime, and accessible optical frequency of the transition makes it well-suited for coherent operations, and thus to be used as a qubit.

The following applies to both 40 Ca⁺and 88 Sr⁺, so we generalize the discussion for both species and omit the principle quantum number of the electronic levels. An external magnetic field splits the $S_{1/2}$ and $D_{5/2}$ levels into 2 and 6 Zeeman sublevels, respectively, denoted by their magnetic quantum numbers m, with $m = \pm 1/2$ for $S_{1/2}$ and $m = \pm \{1/2, 3/2, 5/2\}$ for $D_{5/2}$. The change in frequency of the individual Zeeman sublevels due an applied magnetic field B is given by

$$\Delta f = \frac{\mu_B B}{2\pi\hbar} (g_j(D_{5/2}) m_D - g_j(S_{1/2}) m_S)$$
(2.66)

where μ_B is the Bohr magneton, g_j the Lande g-factors of the respective levels¹¹, and m_D and m_S their magnetic quantum numbers. Selection rules for quadrupole transitions dictate that only transitions where the magnetic quantum number changes by $|\Delta m| = \{0,1,2\}$ are accessible. The allowed transitions between $S_{1/2}$ and $D_{5/2}$ are schematically shown in Figure 2.7(a).

In addition to the ion's internal electronic levels, its external modes of motion collectively represent ladders of possible states, depicted in Figure 2.7(b). For simplicity, we only consider one of the ion's motional modes, allowed since states of different modes ideally do not couple. An ion's full state is described by the tensor of its electronic and motional state.

The energy splitting due to an external magnetic field separates the levels such that a light field can selectively address specific Zeeman sublevel transitions by choosing an appropriate detuning. Figure 2.7(d) shows an illustrative example of an excitation spectrum similar to that in Figure 2.6, but now taking into account all possible excitations from the $S_{1/2}(m=-1/2)$ ground state, under the influence of an applied magnetic field. Next to the various $S_{1/2} \leftrightarrow D_{5/2}$ carrier transitions, we see blue and red sideband transitions, detuned from the carrier resonance by $\pm \omega_l$.

In principle any two of the electronic states in the $S_{1/2}$ and $D_{5/2}$ manifolds can be used to host a qubit. An *optical qubit* has an energy difference in an optical frequency, made by choosing one sublevel from $S_{1/2}$ and one from $D_{5/2}$. Here, a transition where $\Delta m = 0$ is a logical choice since its frequency is comparatively less sensitive to magnetic field noise (see Eq. 2.66), making it a more stable qubit. An alternative choice is to use the two ground state sublevels, $S_{1/2}(m = \pm 1/2)$, as the qubit, known

¹⁰ The omitted terms that scale with η^2 in H_{rsb} and H_{bsb} contain $\hat{a}_l\hat{a}_l$ and $\hat{a}_l^{\dagger}\hat{a}_l^{\dagger}$, known as second order sidebands, where two phonons are added or removed with one electronic excitation. Less truncation of the Taylor expansion of Eq. 2.59 reveals even higher order sidebands.

¹¹ $g_i(S_{1/2}) \approx 2.002$ [74] and $g_i(D_{5/2}) \approx 1.2$ [75, 76].




Figure 2.7: Overview of qubit transitions of group II ions. (a) Zeeman sublevels and allowed transitions of $S_{1/2}$ and $D_{5/2}$ in group II ions. (b) The motional mode occupation n of mode l represents a ladder of possible states. (c) Transitions between Zeeman sublevels of $S_{1/2}$ and $D_{5/2}$ have unique frequencies, which depend on the strength of the applied magnetic field (d) Theoretical excitation spectrum from $S_{1/2}(m=-1/2)$ at 3 G, including first-order motional sideband transitions.

as a *ground state qubit*, which has a few advantages: The upper of these two levels has an effectively unlimited lifetime and coherent operations on the transition, if accessed resonantly, are immensely less sensitive to frequency noise of the field that couples the two states¹². The downside is that the field that couples these states directly cannot be focused narrowly enough to address ions individually. Instead, the $D_{5/2}$ level can be used as an auxiliary level to couple the ground state qubit optically, for example using a two-photon Raman transition [77].

The qubit state of each of these levels can be manipulated using the light-matter interaction methods discussed in the previous section, many of which can be expressed as rotations on the Bloch sphere. In the following, we give an overview of the most common experimental tools and terminology for qubit manipulation.

- A light-field can be applied to multiple ions simultaneously, referred to as *global* operations, or to individual ions by focusing it down narrowly, referred to as *addressed* operations.
- Coherent population transfer is done with on-resonance Rabi oscillations, denoted by the unitary operation

$$R(\theta,0) = \exp(-i\theta\sigma_x/2) \tag{2.67}$$

The pulse duration and power of the light field determines the rotation angle $\theta = \Omega t$.

• A far-detuned light field produces a change in phase due to the AC Stark shift, which is a rotation around the *z*-axis, given by the operation

$$Z(\theta) = \exp(-i\theta\sigma_z/2) \tag{2.68}$$

where the rotation angle is determined by $\theta = \Omega^2 t/(2\Delta)$.

• Changing the phase of the light field results in a rotation around a different equitorial axis. In particular, if the light field is shifted by a phase of $\phi = \pi/2$, we have the operation

$$R(\theta, \pi/2) = \exp(-i\theta\sigma_y/2). \tag{2.69}$$

• In general, a light-field with phase ϕ_f is given by the rotation operation

$$R(\theta,\phi) = \exp\left(-i\theta(\sigma_x\cos\phi + \sigma_y\sin\phi)/2\right). \tag{2.70}$$

- The operation $R(\theta, \phi)$ can be realized with the concatenated operation $Z(-\phi)R(\theta, 0)Z(\phi)$. A laser phase-change is less error-prone than using AC Stark shifts.
- The phonon occupation of an ion (or multiple ions') motional modes can be manipulated through qubit rotations with the light field on resonance with red or blue sideband transitions.
- The shared motional modes are used for entanglement between ions. Several methods exist to achieve entanglement [67, 78]. The method most commonly used for optical qubits is the Mølmer-Sørensen (MS) gate [79], whose experimental implementation is described in Section 2.3.6. An MS gate's operation on two ions is given by:

$$MS(\theta) = \exp(-i\theta(\sigma_x^{(1)} + \sigma_x^{(2)})^2)$$
 (2.71)

with Pauli operators $\sigma_x^{(i)}$ acting on ions i = 1, 2.

The set of operations in this list is sufficient to form a universal set of gates for quantum computation [80].

¹² Compare how steady an archer needs to hold a bow to hit a target 1 meter away against how steady he'd have to hold it if the same target was tens of thousands of kilometers away

2.2.8 Experimental sequences

The light-atom interaction presented in the previous section provides a foundational framework of trapped-ion quantum computation. A set of such operations form the core of an experimental sequence, and are used as elements of a quantum circuit. In addition to executing a quantum circuit, an experimental sequence is preceded with steps of state preparation, and finalized with state readout. This section gives an overview of the components that make up a typical experimental sequence. The details of some of these steps are specific to which ionic species is used as a qubit. We limit the discussion to the ions used in this thesis, 40 Ca⁺and 88 Sr⁺. Generic level schemes, valid for both species, are shown in Figure 2.8(a), illustrating relevant transitions for various sequence steps.




Figure 2.8: (a) Experimental operations on group II ions, and (b) typical sequence for trapped ion quantum computation measurements. Times vary for different experiments, but the values indicate typical orders of magnitude.

-Doppler cooling-

The fidelity of quantum operations is dependent on having a well-defined motional state. Prior to executing a quantum circuit, ions are to be initialized as close to the motional ground state as possible. Whether due to collisions with the background gas, or motional excitation due to the relatively high-power beam used in state detection (described later in this section), one cannot assume that an ion is near the motional ground state at the start of an experimental sequence. Doppler cooling is a first step in reducing the ions' energy.

Doppler cooling is done by applying a light-field to the ions, red-detuned from a short-lived electronic transition. By now, Doppler cooling is well-enough documented [81, 82] and so commonplace in trapped atom and ion experiments, that it doesn't require a full treatment here. It is, however, worth noting that beam parameters, beam power and detuning from resonance, are generally chosen in such a way to minimize the final energy of the ions. This limit can be reached by choosing a detuning of $\delta = -\Gamma/2$ and a beam power (expressed in terms of the coupling strength, a.k.a. Rabi frequency — see previous section) $\Omega \ll \Gamma$, where Γ is the spontaneous decay rate of the probed transition¹³. Under these conditions, the ions energy can reach a mean phonon number of $\bar{n} \approx \Gamma/(2\omega_l)$, in each mode with frequency ω_l . For typical mode frequencies ($\sim 1\,\mathrm{MHz}$), this limit, the Doppler limit, is approximately ten phonons.

The Doppler cooling settings for reaching the Doppler limit ($\delta = -\Gamma/2$, $\Omega \ll \Gamma$) are not well-suited for efficiently cooling ions with energies orders of magnitude above the cooling limit, which can occur, for example, when a particle in the background gas collides with an ion. More efficient cooling is achieved by including a high power ($\Omega > \Gamma$), far-detuned ($\Delta \gg \Gamma$) Doppler cooling beam. This cooling

¹³ $\Gamma = 2\pi$ 22.4 MHz for $^{40}\text{Ca}^+(4S_{1/2} \leftrightarrow 4P_{1/2})$ and 2π 20.4 MHz for $^{88}\text{Sr}^+(5S_{1/2} \leftrightarrow 5P_{1/2})$

process, referred to as refreezing, spans multiple orders of magnitude of ion energy. The term 'refreeze' alludes to the phase transition that ions undergo while cooling, discussed in detail in Chapter 4.

The group II ions, undergoing Doppler cooling, suffer from occasional decay into the meta-stable $D_{3/2}$ level, where the cooling process is halted. To avoid population remaining trapped in this level, a 'repump' beam is on during Doppler cooling, returning the electronic state to the cooling cycle.

-Optical pumping-

Two of an ion's electronic sublevels are used to encode a qubit. At the start of a computation algorithm, the qubit is to be initialized in a particular state, which in our experiments is usually $S_{1/2}(m=-1/2)\equiv |0\rangle$. After Doppler cooling, the ion may be in either the m=+1/2 or m=-1/2 state. The transfer from this state towards a population of 100% in m=-1/2 is achieved by a process known as optical pumping [83].

There are two common methods of optical pumping. One involves using circularly polarized light on the Doppler cooling transition, where the light's polarization is tuned to favor $\Delta m=-1$ in the $S_{1/2}\to P_{1/2}$ transition: negative-circularly polarized light carries negative angular momentum, which must be conserved when an ion absorbs a photon. The $S_{1/2}(m=+1/2)$ sublevel takes part in this transition, whereas population transfer from $S_{1/2}(m=-1/2)$ is suppressed (since there is no $P_{1/2}(m=-3/2)$). Since decay to the $S_{1/2}(m=-1/2)$ level still occurs, population ends up trapped in that state. While this method is relatively quick (on the order of the spontaneous decay rate from $P_{1/2}$ to $S_{1/2}$, approximately 20 MHz), the probability to end up in the desired state is sensitive to the beam's polarization. A high-quality optical setup is required to maximize this probability. In our experiment, we elect to use another method, which is comparatively slower but has a higher state initialization fidelity.

The second method of optical pumping is, in contrast to the polarization selective method, frequency selective. The ions state is transferred by resonantly exciting the $S_{1/2}(m=+1/2)\leftrightarrow D_{5/2}(m=-3/2)$ quadrupole transition and the $D_{5/2}\leftrightarrow P_{3/2}$ transition simultaneously. Spontaneous decay transfers the state from $P_{3/2}$ back to $S_{1/2}$. This cycle continues, until the decay occurs to the sublevel $S_{1/2}(m=-1/2)$. While this method is conceptually similar to the previous one, the distinction is that the relatively narrow linewidth of the $S\leftrightarrow D$ transition ($\sim 1\,\mathrm{Hz}$) allows for a significantly higher resolved state transfer between the different sublevels. This comes at the cost that the $S\leftrightarrow D$ transition typically has a considerably lower Rabi frequency, making this process slower. In our experiment, typical optical pumping times are 400 µs, after which we reach > 99.9% population in the $S_{1/2}(m=-1/2)$ state.

-Sideband cooling-

Doppler cooling is limited to reducing an ion's energy to 7-10 quanta in each motional mode. Further reduction to the motional ground state is achieved with ground-state cooling. There exist multiple techniques to achieve near-ground-state cooling, such as EIT cooling [84] and polarization-gradient cooling [85]. However, by far the most commonly utilized technique is sideband cooling 14 .

Sideband cooling is conceptually similar to optical pumping, in the sense that decay due to spontaneous emission provides a dissipative transfer of population. In the case of sideband cooling, this dissipation process involves reducing the phonon occupation number of motional modes. A pulse is applied to the $S_{1/2}(m=-1/2) \leftrightarrow D_{5/2}(m=-5/2)$ transition, though the beam is detuned by $\Delta=-\omega_l$, where ω_l is the frequency of the mode that is to be cooled. In this red sideband interaction, described by Eq. 2.62, a phonon is removed whenever population is transferred from the *S*-level to the *D*-level. Simultaneously, a beam couples the $D_{5/2} \leftrightarrow P_{3/2}$ transition, the latter of which undergoes fast spontaneous decay back to the $S_{1/2}$ level. This transfer to a fast-decaying level is known as quenching. After the spontaneous decay, the process is repeated, and a phonon is removed with each cycle. The cycle continues until the mode occupation is zero, at which point a red sideband transition is no longer possible. The photon recoil of the last quenching and spontaneous decay processes poses a limit to the minimum attainable mean phonon number.

Sideband cooling is relatively time-consuming, taking up a considerable portion of an experimental sequence, possibly exceeding 10 ms if multiple motional modes require cooling. The frequencies and powers of the beams involved in sideband cooling should be tuned to optimize the cooling rate.

¹⁴ If one has the hardware in place for running QC algorithms, sideband cooling is doable by default. The same cannot be said for EIT and PG cooling, which require additional advanced setups.

Other than simply to reduce experimental time, optimizing the sideband cooling rate is important for minimizing the phonon number at the end of a cooling sequence, which finds an equilibrium value determined by the cooling rate and the heating rate (see Section 2.3.4). In practice, the beam coupling the $S_{1/2} \leftrightarrow D_{5/2}$ sideband transition is set to be as strong as possible (limited by the amount of laser power experimental hardware can provide), while the $D \leftrightarrow P$ "quenching" transition coupling strength is set relatively low with respect to the experimentally available power, since a lower coupling strength minimizes lowest attainable mean phonon number [69]. On the other hand, *too low* beam power results in a low cooling rate, and thus a higher final mean phonon number, as it competes with the trap's heating rate.

Though relatively unlikely, decay may occur from the $P_{3/2}$ to the $D_{3/2}$ level. While population trapped in $D_{3/2}$ can also be repumped, followed by decay to $S_{1/2}$, this process may result with state population in the $S_{1/2}(m=+1/2)$ sublevel, which terminates the cooling cycle. Sideband cooling is therefore interleaved with optical pumping pulses to reintroduce the state into the cooling cycle, if necessary. A final optical pumping pulse follows the sideband cooling step. Depending on cooling requirements, multiple modes may need to be cooled. In this case, each mode is cooled sequentially.

-Qubit operation-

The heart of an experimental sequence, and the most versatile and thus difficult to generalize, is the set of pulses that produce the desired qubit operations, for example implementing the set of gates of a QC circuit [86]. Generally, this part of the sequence entails the set of operations discussed in Section 2.2.7, but can include other operations, such as waiting times, or manipulation of an ion's position through changes in electrode voltages.

This step in the experimental sequence is specific to the experiment being carried out. Detailed descriptions are therefore deferred to the Chapters that describe the specific experiments, Chapters 4 through 6.

-State detection-

At the end of a sequence, the state of the ion(s) is determined using fluorescence detection [55]. The same set of beams used for Doppler cooling are used for state detection. If the qubit was in the $S_{1/2} \equiv |0\rangle$ state, the ion will fluoresce since this state is part of the Doppler cooling cycle. If it was in the $D_{5/2} \equiv |1\rangle$ state¹⁵, the electronic level is not part of the cooling cycle and the ion is 'dark'.

A photomultiplier tube (PMT) or a CCD camera is used to detect emitted fluorescence photons. If during the detection time the photon counts exceed a predetermined threshold, the ion is determined to have been in $S_{1/2}$, and otherwise in $D_{5/2}$. State detection through a CCD camera allows one to distinguish the state of individual spatially separated ions, whereas a PMT only detects the total number of bright and dark ions. Historically, the use of PMTs is based on their superior detection efficiency, high signal-to-noise ratio, and a relatively simple integration into experimental control software. With recent developments on each of these fronts for CCD cameras, using them for state detection is becoming more favorable. Our experimental setup has, as of yet, not implemented improved CCD camera detection. Therefore, most presented data makes use of PMT-based state detection, where the exceptions are cases that resolution of multiple-ion states is necessary.

2.3 SURFACE TRAPS: POTENTIALS AND CHARACTERIZATION

In Section 2.1, the foundation of ion traps and their trapping potentials was introduced. *Scalable* trapping architectures, to be used in the context of quantum computation, are envisioned to take on a QCCD design: a segmented surface trap. In contrast to the traditional 3D blade trap, producing a confining potential like in Eq. 2.17 is not trivial on a 2D segmented trap. For example, since trap electrodes are not placed symmetrically around the trap center, one cannot simply apply a high voltage on two opposing ends of the trap to create a harmonic confining potential, without introducing additional unwanted fields. On the other hand, the number of degrees of freedom in control over trapping potentials is comparable to the number of electrodes. The larger number of electrodes of a segmented trap provides more control over the trapping potentials compared with a 3D blade trap.

¹⁵ If the qubit is encoded in the Zeeman sublevels of $S_{1/2}$, (i.e., a ground-state qubit), an additional analysis pulse is used to first transfer one of the two states to a sublevel in $D_{5/2}$

Section 2.3.1 gives an overview of how we generate and control trapping potentials using voltages on a surface trap's electrodes. Once trapped, ions can be used as sensors to characterize properties of the experimental setup such as trap potentials, surface noise from the trap electrodes, magnetic field stability, and laser frequency stability. Such characterization makes use of the laser-ion interaction toolbox discussed in Section 2.2. Section 2.3.2 describes how potentials are calibrated in order to characterize the discrepancy between simulated and experimentally realized potentials. Section 2.3.4 discusses how an unintended gain in energy in the ion's motion, a notorious issue for ions in surface traps, is characterized. A method of determining how well quantum coherence in a qubit is maintained is discussed in Section 2.3.5. Finally, the ability to entangle qubits through a common entangling gate, the Mølmer-Sørensen (MS) gate, and characterization of the gate performance is discussed in Section 2.3.6.

2.3.1 Trapping potentials - Spherical harmonics expansion

Voltages on the trap's electrodes produce an electric field potential. One aims to apply voltages to the electrodes to place ions at a specific position within the trap, with a specific motional frequency. It is impractical to control these parameters by tweaking individual voltages, since this will likely incur undesired stray fields, leading to excess micromotion. Instead, we determine how a *set* of voltages applied to the electrodes produces a given potential. Multiple of such potentials form a basis from which a linear combination generates a total trap potential, which uniquely governs ion positions and motional frequencies.

For this basis of potentials, we chose spherical harmonic potentials [87]. We see later that this basis contains terms that are natural and intuitive for ion trapping. The terms in this basis are denoted by the spherical harmonic functions $Y_{l,n}$, characterized by a degree l and an order n. They are a function of position and can be analytically expressed in Cartesian coordinates. The terms make up an orthogonal set of solutions of the condition $\nabla^2 Y_{l,n} = 0$. Any potential in a charge-free volume can be expressed up to a chosen degree as a linear combination of these terms, known as a multipole expansion. The $Y_{l,n}$ terms are referred to as multipoles. Terms up to and including l = 2 are sufficient to precisely describe any *quadratic* potential.

The terms up to second degree¹⁶ are given by

$$Y_{0,0} = 1$$

$$Y_{1,\{-1,0,1\}} = \{y, z, x\}$$

$$Y_{2,-2} = xy$$

$$Y_{2,-1} = yz$$

$$Y_{2,0} = 2z^2 - x^2 - y^2$$

$$Y_{2,1} = xz$$

$$Y_{2,2} = x^2 - y^2$$
(2.72)

Referring back to Eq. 2.17, we can see why this basis is considered intuitive for ion trapping. The term $Y_{2,0}$ is the same as the axial confining potential brought on by the end-caps, V_{DC} . $Y_{2,2}$ corresponds to the radial bias introduced by applying a voltage to opposing blades, V_b . The $Y_{1,\{-1,0,1\}}$ terms represent homogeneous electric fields, typically used to compensate undesired stray fields that may be present around the trapping region. We will see in Chapter 5 that $Y_{2,1}$ plays a crucial role in manipulating the angular orientation of an ion chain.

The goal is to derive a set of electrode voltages that produces the potential corresponding to each individual multipole. The method of deriving these voltages is as follows (adapted from App. B. in Ref. [88]).

A model of our trap's electrode layout, including the gaps between the electrodes, and the slot in the center of our trap (see Section 3.2 for a description of surface trap layouts) is analyzed in the simulation

¹⁶ Often these terms are displayed with a normalization constant before them. We omit them because they are unnecessary for the purpose of generating ion trap potentials. In fact, it's easier to make an intuitive link to field potentials, which directly relates to motional frequencies and ion positions, without these scaling factors.

software Comsol. Comsol uses finite-element methods to calculate electrostatic potentials for a set of boundary conditions, such as voltages on electrodes. We calculate the potential $\Phi_{i,k}(x_i, y_i, z_i)$ that results from applying 1 V to electrode k, at discrete points $\{x_i, y_i, z_i\}$ indexed i, centered around the trapping region. All other electrodes are set to 0 V. The range of this grid of points is chosen such that higher-than-second order gradient terms are negligible, but large enough that the curvature (second-order gradient) is not underdetermined. We choose a range that is approximately 20% of the electrode-ion separation of our trap, 20 µm³. The potential is analyzed in a uniform grid within this range, with 21 grid points in each dimension, giving $n_i = 9261$ data points per electrode k, cast into the vector $\vec{\Phi}_k$.

The potential produced by an electrode is to be expanded into the basis of spherical harmonics. This is done by first analytically calculating the values of $Y_{l,n}(x_i, y_i, z_i)$ for all grid points, for the $n_m = 9$ terms of Y up to second degree, $l \le 2$. For notational convenience, all 9 possible combinations of l and n are denoted with a single index m. The center of the grid volume coincides with the desired trapping location, with x = y = z = 0.

The potential at each grid point can be expressed as a linear combination of spherical harmonic terms, $\sum_{m} Y_m(x_i, y_i, z_i) w_m^{(k)} = \Phi_k(x_i, y_i, z_i)$, where $w_m^{(k)}$ is a set of weights, cast into $\vec{w}^{(k)}$.

To solve for $w_m^{(k)}$, it is helpful to cast the values Y_m at each grid point i into a $n_i \times n_m$ matrix \mathbf{Y} . The

equation to be solved is

$$\mathbf{Y}\vec{w}^{(k)} = \vec{\Phi}_k \tag{2.73}$$

in which the least-squares solution can be obtained by using the singular-value decomposed pseudoinverse [89], Y^{-1} , a matrix-inversion technique for solving over-determined linear equations with a least-squares regression. The weights are given by

$$\vec{w}^{(k)} = \mathbf{Y}^{-1} \vec{\Phi}_k. \tag{2.74}$$

The coefficients in $\vec{w}^{(k)}$ describe the potential of electrode k, expressed in the basis of spherical harmonic potentials. Combining all $\vec{w}^{(k)}$ into the $n_m \times n_k$ matrix \mathbf{w} , we now need to determine what linear combination of voltages \vec{V}_m can produce each individual spherical harmonic potential. In matrix form, this is represented by the equation

$$\mathbf{wV} = 1 \tag{2.75}$$

where **V** is the $n_k \times n_m$ matrix containing \vec{V}_m for all m. Once again, a pseudoinverse is used to find the least-squares solution, given by

$$\mathbf{V} = \mathbf{w}^{-1}.\tag{2.76}$$

In calculating voltage solutions for spherical harmonic potentials, some harmonic terms can be omitted. This removes constraints on the remaining solutions, allowing solutions to be found at lower voltages, which may be useful for setups with limited voltage supply range. It is particularly useful to remove entries corresponding to $Y_{0,0}$, since this potential does not produce any electric field $(E = -\nabla Y_{0,0} = 0)$ and so does not affect an ion's position. Other rows can be omitted, for example, when the number of electrodes is low: as a rule-of-thumb, each electrode adds a degree of freedom to producing a unique potential. Therefore, if one wants to uniquely control 8 multipole potentials, at least 8 electrodes are required. For reasonable trap operation, at least the multipoles $Y_{1,-1}$, and $Y_{1,1}$, used for nudging the ions towards the RF-null, and $Y_{2,0}$, the primary DC trapping potential, should be included.

The significance of V is that it describes which set of voltages are required for each spherical harmonic potential. This means there is an independent set of voltages that controls the magnitude of each potential. We will describe later that a chosen position and set of motional frequencies can be uniquely described by a combination of spherical harmonic potentials. We will thus have a mapping between desired positions and frequencies, and electrode voltages.

2.3.2 Potential calibration

The solution for the required electrode voltages for a given static trapping potential may not produce the actual expected potentials, due to simulation imprecision, uncontrolled electric charges on the trap surface, and manufacturing tolerances. These discrepancies may lead to unexpected mode frequencies and ion positions. Deviations from intended mode frequencies are typically easily dealt with since they are straightforward to measure and correct. Deviations in an ion's intended position are caused by incorrect settings of the $Y_{1,\{-1,0,1\}}$ terms, often caused by the presence of stray fields that have not been accounted for. Crucially, offsets in $Y_{1,-1}$ and $Y_{1,1}$ force ions away from the RF-null, where they experience a higher micromotion amplitude, known as *excess micromotion*. There exist several commonly used methods of minimizing excess micromotion, as is discussed in Section 2.3.3.

For some applications, a higher degree of understanding and control of actual trap potentials is required. Notably, Chapter 5 discusses methods for physically manipulating ion positions through control of trap potentials, where it is clear that successful ion transport requires precisely tailored sequences of potentials. For such operations, the discrepancy between desired and actual potentials needs to be estimated. This section covers a measurement protocol to calibrate this discrepancy.

A list of voltages \vec{V} is generated for a chosen list of multipoles, $\vec{m}_{\rm set}$, using $\mathbf{V}\vec{m}_{\rm set} = \vec{V}$, with the multipole matrix \mathbf{V} found using Eq. 2.76, discussed in the previous section. $\vec{m}_{\rm set}$ contains the 8 multipole terms corresponding to $Y_{l,n}$ for l=1,2. Applying these voltages to the trap electrodes produces a potential that can be expressed in terms of multipole contributions, denoted by \vec{m} , which might deviate from $\vec{m}_{\rm set}$. We model this possible discrepancy as

$$\vec{m} = \mathbf{A}\vec{m}_{\text{set}} + \vec{b}. \tag{2.77}$$

A is an 8×8 correlation matrix between set and actual multipoles, and the vector \vec{b} is a constant offset. A method to estimate **A** and \vec{b} is outlined below.

The applied RF power and the eight terms of the vector \vec{m} uniquely define a potential with ellipsoidal equipotential surfaces. Such an ellipsoid can be characterized by nine parameters: the displacement of the center of the ellipsoid in three Cartesian directions, three tilt angles of the ellipsoid's principal axes with respect to the Cartesian axes, and three lengths of the principal axes of the ellipsoid. These parameters translate into measurable quantities of a single trapped ion. The calibration terms $\bf A$ and $\vec b$ can therefore be found by experimentally estimating the actual trapping potential for various multipole settings $\vec m_{\rm set}$, and comparing experimental results to calculated expected potentials. We will now describe how each of the potential parameters are experimentally determined.

The lengths of the principal axes of an equipotential ellipsoid define the curvature of the potential in three orthogonal directions $\alpha_{\{x,y,z\}}$, and are thus determined by the motional frequencies of the ion, as $\omega_k = \sqrt{\alpha_k q/m}$, with q and m the charge and mass of the ion. Sideband spectroscopy (see Section 2.2.7) is used to determine the ion's motional frequencies $\omega_k^{(j)}$ for various multipole settings $\vec{m}_{\rm set}^{(j)}$, indexed by j.

The tilt of the principal axes can be determined by measuring the relative coupling strengths of a beam with wavevector \vec{k} on resonance with the motional sidebands of an ion's electronic transition. As shown in Section 2.2.6, the Lamb-Dicke parameter, and therefore the coupling strength Ω_k of mode k, depends on the angle of incidence of the beam with respect to orientation of the motional mode, $\Omega_k \propto \vec{k} \cdot \hat{e}'_k$. Here \hat{e}'_k are the to-be-determined unit vectors of the potential's principal axes $k \in \{x', y', z'\}$, accented to denote that they are not necessarily the Cartesian axes of the lab system. The coupling strength also depends on the motional mode frequency, $\Omega_k \propto \omega^{-1/2}$. One can find the coupling strengths by measuring the effective Rabi frequency when exciting the blue sideband for each

motional mode of an ion prepared in the motional and electronic ground state. The tilt of the principal axes can be estimated by determining their unit vectors by solving the following set of equations:

$$\hat{e}_x' \times \hat{e}_y' = \hat{e}_z' \tag{2.78}$$

$$\frac{\Omega_x}{\Omega_y} = \frac{\vec{k} \cdot \hat{e}_x'}{\vec{k} \cdot \hat{e}_y'} \sqrt{\frac{\omega_y}{\omega_x}}$$
 (2.79)

$$\frac{\Omega_y}{\Omega_z} = \frac{\vec{k} \cdot \hat{e}_y'}{\vec{k} \cdot \hat{e}_z'} \sqrt{\frac{\omega_z}{\omega_y}}$$
 (2.80)

Using this method assumes that the wavevector \vec{k} is known. If there is uncertainty in the wavevector, one can also compare the sideband coupling strengths to that of a carrier excitation Ω_{car} , which gives the following additional equation

$$\frac{\Omega_k}{\Omega_{\rm car}} = \sqrt{\frac{\hbar}{2m\omega_k}} \vec{k} \cdot \hat{e}'_k. \tag{2.81}$$

 \hat{e}'_k is dependent on the configuration of multipoles, $\vec{m}_{\rm set}$. Similar to the measurements of $\omega_k^{(j)}$, we measure values of $\hat{e}_k^{(j)\prime}$ for various multipole configurations, $\vec{m}_{\rm set}^{(j)}$, where the index j refers to each distinct configuration.

Finally, a displacement of the ellipsoid, caused by a uniform field, results in an identical displacement of an ion, \vec{r} , since the center of the ellipsoid represents the potential minimum. Therefore, one can measure the position of the ion to characterize this component of the potential. In our experiment, a CCD camera monitors the position of a single ion with respect to the trap plane. We thus determine the displacement in the xz-plane, denoted by $r_x^{(j)}$ and $r_z^{(j)}$, for various multipole settings $\vec{m}_{\text{set}}^{(j)}$ by monitoring the position of an ion on the CCD image. The magnification is calibrated by using imaged trap electrodes as a scale reference. Alternatively, one can calculate the separation d of two trapped ions for a measured common mode frequency ω_z with $d^3 = q^2/(2\pi\epsilon_0 m\omega_z^2)$, with ϵ_0 the vacuum permittivity, and compare this to the separation on the CCD.

The displacement r_y is perpendicular to the image plane, and is therefore not readily detected by the CCD. An ion displaced in this direction experiences excess micromotion, as discussed in Section 2.3.3. The amplitude of this motion depends linearly on the displacement of the ion from the RF pseudopotential minimum. We can therefore infer ion displacement by measuring the micromotion amplitude. This amplitude is measured with micromotion sideband spectroscopy [90], in which we determine the ratio between the coupling strength of a micromotion sideband $\Omega_{\rm MM}$, and that of a carrier transition $\Omega_{\rm car}$. The micromotion amplitude, denoted by the micromotion modulation index β , is given by $\beta/2 \approx \Omega_{\rm MM}/\Omega_{\rm car}$. The modulation index β is related to ion displacement as $r_y = 2\beta/kq_y$ [91]. Here, k is the wavenumber of a beam propagating perpendicular to the trap surface, and q_y is the trap's stability parameter (see Eq. 2.5).

For any given multipole configuration \vec{m}_j , positions $\tilde{r}_k(\vec{m}^{(j)})$ can be calculated using equation 2.22, and frequencies $\tilde{\omega}_k(\vec{m}^{(j)})$ and tilts $\tilde{\tilde{e}}_k(\vec{m}^{(j)})$ are given by the eigenvalues and eigenvectors of equation 2.23. Calculated values of these parameters for various intended multipoles \vec{m}_j are compared with measured values, giving a cost function given by:

$$\sum_{k,j} \left[(\omega_k^{(j)} - \tilde{\omega}_k(\vec{m}_j))^2 + \left| \hat{e}_k^{(j)\prime} - \tilde{e}_k(\vec{m}_j) \right|^2 + (r_k^{(j)} - \tilde{r}_k(\vec{m}_j))^2 \right]$$
(2.82)

where the sum is taken over all measurements j, with measurement results in each of three axes k. We numerically minimize the cost function of Eq. 2.82. The correlation matrix \mathbf{A} and offset \vec{b} are used as optimization parameters, noting that the values of \vec{m}_j are dependent on these parameters, as given by Eq. 2.77. We thus find optimized values for \mathbf{A} and \vec{b} .

2.3.3 Micromotion compensation

The equations of motion of an RF trap, presented in Section 2.1.1, lead to solutions that include motion at two distinct timescales, secular motion and micromotion. The solution of Eq. 2.6 suggests that the amplitude ratio between secular motion and micromotion is $q_k/2$, with q_k the Mathieu q-parameter. Eq. 2.6, however, is derived with the assumption that the DC potential $V_{\rm DC}$ has no linear component around the trap center $\vec{r}=0$, such that $\nabla V_{\rm DC}|_{\vec{r}=0}=0$. In reality, charges on electrodes or miscalibrated electrode potentials can lead to undesired stray fields. These fields force an ion away from the trap center. If ions are radially further away from the trap center, they are subjected to a larger RF force (which grows quadratically with distance away from the RF-null). The ion is then said to experience excess micromotion.

While some excess micromotion may be tolerated, it has several adverse effects. For example, it reduces the nominal coupling strength between a light field and a carrier, it presents a fundamental limit to the lowest attainable phonon number with sideband cooling [90], and it may induce stronger coupling between motional modes [66], particularly in anharmonic trapping potentials. It is therefore a requirement for QC to compensate stray fields in order to minimize micromotion. There are multiple experimental techniques to minimize micromotion [91]. Two of these techniques have already been alluded to in the section on potential calibration, Section 2.3.2. There, we have introduced the vector \vec{b} , which describes the stray fields around the trapping region when all electrode voltages are zero. In particular, the first three terms of \vec{b} correspond to homogeneous electric fields, of which the x and y directions, the first and third entry of \vec{b} , require compensation. The values of \vec{m}_{set} must be chosen in such a way that these entries in \vec{m} become zero.

In practice, stray fields change on a day-to-day basis. Daily recalibration of all terms in **A** and \vec{b} is time-consuming, so a more direct method of micromotion compensation is desired, described below. The following methods are ones that are commonly used in our experiment, though is not an exhaustive list of micromotion compensation techniques. More methods can be found in Ref. [91].

-Camera detection-

The potential in one radial dimension $(r \in \{x, y\})$ is given by

$$V_r = \frac{1}{2}(\phi_{RF} + \phi_{DC})r^2 + (E_r + E_b)r$$
(2.83)

with the RF pseudopotential curvature ϕ_{RF} , the DC curvature ϕ_{DC} , and a DC electric field E_r that may be due to stray charges or potential miscalibration, and E_b an intentionally applied field to correct for it.

The ion's equilibrium position, found by solving $\delta V_r/\delta r = 0$, is given by

$$r_0 = -\frac{E_r + E_b}{\phi_{RF} + \phi_{DC}} \tag{2.84}$$

Micromotion is compensated when the ion is at the RF-null, $r_0=0$, valid when $E_b=-E_r$. Compensation through camera detection involves adjusting the RF power, thus changing the pseudopotential curvature $\phi_{\rm RF}$, while monitoring the ion's position by fluorescence detection with a CCD camera. If the ion is at the RF null, its position does not change. Otherwise, if $E_b\neq -E_r$ a change in RF power will result in a change in ion position. E_b is adjusted until the ion no longer visibly moves. This method is only applicable for compensation in the directions that are accessible in the image plane.

-Micromotion sideband-

It was shown in Section 2.2.6 that an ion's secular oscillatory motion produces motional sidebands in the laser excitation spectrum. A similar argument can be made for micromotion: excess micromotion produces sidebands, whose coupling strength Ω_{MM} to the light field depends on the amplitude of the motion. Simultaneously, the light field's carrier coupling strength Ω_{car} decreases when excess micromotion is present. Both coupling strengths are readily obtained by measuring Rabi oscillations

(see Section 2.2.7) on the carrier and micromotion sideband. The ratio of coupling strengths is given by

$$\frac{\Omega_{\text{MM}}}{\Omega_{\text{car}}} = \frac{J_1(\beta)}{J_0(\beta)} \tag{2.85}$$

with $J_n(\beta)$ the Bessel function of the first kind, with order n. The modulation index β relates the ion's micromotion oscillation amplitude \vec{r}_{MM} to the wavelength of the laser, as

$$\beta = \vec{r}_{\text{MM}} \cdot \vec{k} \tag{2.86}$$

with \vec{k} the light field's wavevector. The micromotion amplitude is related to the ion's displacement \vec{r} by

$$\vec{r}_{\text{MM}} \approx \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \vec{r} \frac{q}{2}$$
 (2.87)

with *q* the Mathieu *q*-parameter.

For $\beta \ll 1$, the coupling strength ratio is approximated by

$$\frac{\Omega_{\mathrm{MM}}}{\Omega_{\mathrm{car}}} pprox \frac{\beta}{2}$$
 (2.88)

Micromotion is thus minimized by measuring and minimizing the ratio of coupling strengths $\Omega_{\rm MM}$ and $\Omega_{\rm car}$ when varying the compensation fields E_b . This minimizes an ion's micromotion amplitude projected in the direction along the light field's wavevector. For compensation along both radial dimensions, at least two beams, each at a different angle in the radial plane, are required.

-Photon correlation-

A third method of micromotion compensation is through the photon correlation technique. The technique relies on the fact that the absorption and subsequent spontaneous emission rate of an ion undergoing Doppler cooling is dependent on the ion's velocity. An ion's velocity, in turn, is driven by the trap's RF field, and is dependent on how far the ion is from the RF null. Excess micromotion can be detected by temporally correlating the the phase of the RF field and the ion's emission rate. This process is illustrated in Figure 2.9, and described below.

An ion's velocity is correlated with the phase of the RF field that drives it, as shown in the twodimensional histograms in the upper panels of Figure 2.9, where lighter colors represent a higher likelihood for that combination of RF phase and ion velocity to occur. The plots are generated by simulating the motion of an ion in an RF trap (the simulations are discussed in detail in Section 4.4).

Velocity-phase correlation plots are shown for the case where there is no radial field (left), and where there is a $10 \,\mathrm{V}\,\mathrm{mm}^{-1}$ field (right), under otherwise typical trapping conditions. The interaction of the ion with the Doppler cooling beam is velocity dependent, as displayed in the plot on the right, calculated using typical Doppler cooling parameters¹⁷.

Since the ion's velocity depends on the RF drive phase, it is evident that the fluorescence rate is likewise dependent on that phase¹⁸. Combining the information from the upper plots in Figure 2.9, we display the Doppler cooling fluorescence as function of RF phase, where its phase dependence—or lack thereof—indicates how well micromotion is compensated.

In our setup, a portion of the trap's RF drive signal is sent to a signal counter. Simultaneously, photomultiplier tube (PMT)-pulses from the detection of an ion's fluorescence are sent to another port of the counter. The time delay between each RF drive signal and subsequent PMT counts are recorded, and displayed as a histogram. A flat histogram indicates that micromotion is compensated. As with the

¹⁷ $\Gamma/(2\pi)=22\,\mathrm{MHz}$, $\delta=-\Gamma/2$, $\Omega\ll\Gamma$, see Section 2.2.8

¹⁸ The natural decay time of the $P_{1/2} \rightarrow S_{1/2}$ is ~ 8 ns. Comparing that to the oscillation period of the RF drive, ~ 25 ns, we note that photon detection times are not perfectly correlated with the RF phase, but enough to be able to measure a correlation signal.

previous methods, the photon correlation technique is only sensitive to micromotion in the direction of the wavevector of the Doppler cooling beam.




Figure 2.9: Photon correlation for micromotion compensation: the upper plots are histograms of simulated Doppler cooled ion velocities, and the phase of the RF trap drive, shown for no stray field (left) and $10\,\mathrm{V}\,\mathrm{mm}^{-1}$ stray field (right). Light colors indicate a high occurrence. The lower plots show expected fluorescence as function of RF trap drive phase, noting that a higher fluorescence count is expected when an ion's velocity is negative (as shown in the right plot). Micromotion is compensated by flattening the correlation curve between trap drive phase and measured fluorescence.

In our experiment, we use a combination of techniques discussed above to detect and correct for micromotion. The micromotion sideband method for micromotion compensation is the most accurate of the three and provides a quantitative description of the amount of micromotion in terms of the modulation index β . However, the advantage of the photon correlation technique over the previous methods is that it non-invasively runs in the background, and does not require any calibration of micromotion sideband frequency or beam power. The necessary signals for photon correlation, PMT counts and the trap drive RF are continuously picked off without affecting the experiment. The photon-correlation technique is more sensitive when the Doppler cooling beam has low power ($\Omega \ll \Gamma$) and thus is less precise when high power is applied, for example during state detection.

2.3.4 Heating rates

The functionality of entangling operations relies on the ions' shared motional modes. High-fidelity quantum computation with trapped ions therefore requires that the motional state of an ion is well-controlled. In practice, this means that an ion's motional modes should be cooled to the ground state prior to any logical operations. While common entangling operations are designed to be resilient to imperfect ground state cooling [79], undesired changes in the motional state during a sequence, *motional heating*, disrupts the gate operation.

The predominant source of ion heating is usually electric field noise generated by surfaces near the ion, the trap electrodes [29, 32, 92–94]. Heating rates are especially a nuisance in *surface* ion traps, as the distance between the ion and the nearest electrode is typically an order of magnitude smaller than conventional macroscopic traps. Since motional heating presents a major source of error in QC sequences, it is useful to characterize a trap's heating rates. Two common methods [29] used are Rabi decay measurements and sideband thermometry¹⁹. The general sequence for both measurements is similar: ground state cool an ion, wait a fixed duration, and then perform Rabi cycles on either a carrier $S_{1/2} \leftrightarrow D_{5/2}$ transition, or on its sidebands. The ion's mode occupation can be inferred from measurements of the state population as a function of Rabi cycle time. The measurement is repeated

¹⁹ A third type of heating rate measurement is the Doppler recooling [95] technique, reserved for particularly high heating rates. If one needs to resort to using this kind of measurement on a trap intended for quantum computation, one would do well to consider replacing the trap altogether.

for various wait-times between cooling and probing the ion. The change in mean mode occupation \bar{n} over time is the heating rate $\dot{\bar{n}}$.

-Carrier excitation-

It was shown in Section 2.2.6 that the Rabi frequency of a resonant carrier transition is, to second order, given by $\Omega_{n,n} = \Omega(1 - \eta^2 n)$, with Ω the bare (n = 0) Rabi frequency, and η the Lamb-Dicke parameter. An ion whose motional state is in a thermal distribution²⁰ with mean phonon number \bar{n} has a phonon occupation probability given by

$$P_n(\bar{n}) = \frac{1}{\bar{n}+1} \left(\frac{\bar{n}}{\bar{n}+1}\right)^n \tag{2.89}$$

for phonon number n.

An ion with a phonon distribution $P_n(\bar{n})$ will exhibit Rabi oscillations with multiple frequency components, following Eq. 2.44. The excited state population develops as

$$P_{|1\rangle} = \sum_{n=0}^{\infty} P_n(\bar{n}) \sin^2\left(\frac{\Omega_{n,n}t}{2}\right) \tag{2.90}$$

The distribution of Rabi frequencies for $\bar{n}>0$ leads to a dephasing of the Rabi oscillations over time, which is seen as a decay in the excitation cycles of $P_{|1\rangle}$. A high rate of decay indicates a high mean phonon number. Examples of this decay for $\bar{n}=0$ and 10 are shown in 2.10(a), using a Lamb-Dicke parameter of $\eta=0.06$, that of a $^{40}\text{Ca}^+$ ion with 1 MHz motional frequency. Measured Rabi oscillations can be fit using Eq. 2.90 to infer the mean phonon number.




Figure 2.10: Examples of motional mode occupation (phonon number) measurements using $\eta = 0.06$, through (a) decay in Rabi cycles when exciting a carrier transition at mean phonon number $\bar{n} = 0$ and 10, and (b) development of blue and red sideband (bsb, rsb) excitation for $\bar{n} = 0$ and 1.

-Sideband thermometry-

Alternatively to the carrier oscillations, we can monitor the excitation development when applying a beam on resonance with the red and blue motional sideband. The Rabi frequencies of these interactions are $\Omega_{n,n-1} = \Omega \eta \sqrt{n}$ for the red sideband, and $\Omega_{n,n+1} = \Omega \eta \sqrt{n+1}$. For a thermal distribution, the excitation probabilities evolve as

$$P_{|1\rangle}^{(\text{rsb})} = \sum_{n=0}^{\infty} P_n(\bar{n}) \sin^2\left(\frac{\Omega\eta\sqrt{n}t}{2}\right)$$
 (2.91)

$$P_{|1\rangle}^{(\text{bsb})} = \sum_{n=0}^{\infty} P_n(\bar{n}) \sin^2\left(\frac{\Omega \eta \sqrt{n+1}t}{2}\right),\tag{2.92}$$

which are shown for $\bar{n} = 0$ and $\bar{n} = 1$ in Figure 2.10(b).

²⁰ A Doppler cooled ion, or one that has undergone heating from a ground state, is generally considered to have a thermal distribution.

Noting that the n = 0 term in the sum of $P_{|1\rangle}^{(\text{rsb})}$ is zero, we can rewrite it as

$$P_{|1\rangle}^{(\text{rsb})} = \sum_{n=1}^{\infty} \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}} \sin^2\left(\frac{\Omega\eta\sqrt{n}t}{2}\right)$$
 (2.93)

$$=\sum_{\tilde{n}=0}^{\infty} \frac{\tilde{n}^{\tilde{n}+1}}{(\tilde{n}+1)^{\tilde{n}+2}} \sin^2\left(\frac{\Omega \eta \sqrt{\tilde{n}+1}t}{2}\right) \tag{2.94}$$

$$=\frac{\bar{n}}{\bar{n}+1}P_{|1\rangle}^{(\text{bsb})}\tag{2.95}$$

where the substitution $\tilde{n}+1=n$ was used. The ratio between the red and blue sideband excitation is thus only dependent on the mean phonon number \bar{n} , given by $R=P_{|1\rangle}^{(rsb)}/P_{|1\rangle}^{(bsb)}=\bar{n}/(\bar{n}+1)$. Measuring this ratio allows us to estimate the mean phonon number, as $\bar{n}=R/(1-R)$. Note that for $\bar{n}\gg 1$, R approaches 1, and becomes increasingly less reliable as a measurement of phonon number. At higher phonon numbers, typically around ten or more, carrier excitation measurements are more accurate indicators of ion temperature. However, unlike with sideband thermometry, carrier excitations do not readily distinguish heating of multiple modes. Should this distinction be desired, more delicate curve fitting, and/or multiple beams with different wavevectors are required.

2.3.5 Coherence time

This section has so far only considered qubit operations involving *pure* states, which were formally defined in Section 2.2.5 in terms of the density operator $tr(\rho^2) = 1$, and informally defined as states whose Bloch vectors have unit length, and thus lies on the surface of the Bloch sphere. High fidelity quantum computation requires that qubit states remain pure, which implies that decoherence should be avoided [9, 96]. Qubit dephasing mechanisms, such as fluctuations in the magnetic field or laser phase, lead to a loss of coherence over time. It is therefore beneficial to characterize how long a qubit's quantum coherence is maintained. In a trapped ion experiment, the coherence time is obtained using a Ramsey sequence, which is outlined below and schematically shown in the Bloch sphere representation in Figure 2.11.

A qubit is initialized in $|0\rangle$. A Ramsey measurement consists of the sequence

$$R(\pi/2,0) \Rightarrow t_{\text{wait}} \Rightarrow R(\pi/2,\phi)$$
 (2.96)

where $t_{\rm wait}$ is a chosen waiting time between pulses. The first pulse prepares the state $1/\sqrt{2}(|0\rangle-i|1\rangle)$. If the second pulse is applied without a change in phase with respect to the first, $R(\pi/2,\phi=0)$, ideally the state $|1\rangle$ is fully populated. Similarly, a laser phase change of $\phi=\pi$, $R(\pi/2,\phi=\pi)$, ideally fully returns population to $|0\rangle$. Sources of noise (most commonly changes in magnetic field and laser frequency) during the waiting time cause the qubit phase to change on a shot-to-shot basis, after which the final pulse no longer fully transfers to $|1\rangle$ and $|0\rangle$. The loss of contrast between these two states indicates qubit dephasing, and can be obtained as a function of wait time. Figure 2.11 schematically shows this decay process in the Bloch sphere representation.

The qubit phase could evolve coherently during the waiting time, which, in the presented scheme, would be detected as a loss of contrast. This could occur, for example, if the frequency of the laser pulses is marginally off-resonance from the transition it is probing, which results in a constant qubit phase accumulation during the waiting time. Scanning the laser phase of the second pulse over the full range $\phi \in [0,2\pi]$ instead of just at $\phi=0$ and $\phi=\pi$ ensures that maxima and minima in population transfer are still found and contrast due to decoherence can be established. The plot in Figure 2.11 displays an example of the output of such a scan at two wait times. The wait time required for contrast to reduce to $1/e\approx 0.37$ is known as the T_2 -time of the qubit.

When characterizing sources of decoherence, one can distinguish noise in two timescales. 'Fast' noise arbitrarily alters the qubit phase *during* an experimental cycle, and is difficult to actively correct for. 'Slow' noise is considered to be stable during a single cycle, but can result in a different qubit phase evolution in each subsequent experimental cycle. The influence of this type of noise can often




Figure 2.11: Bloch sphere representation of a Ramsey sequence for a coherence time measurement. A state along the equator dephases due to environmental noise, shown by the fanning out of Bloch vectors. This spread limits the population transfer to $|0\rangle$ and $|1\rangle$ with the final $\pi/2$ pulse. Scanning the phase of the final pulse results in the plot on the bottom left, where the contrast indicates qubit coherence.

be corrected for and can drastically increase the qubit coherence time. In Ramsey sequences for determining qubit coherence, this type of noise can be accounted for by including a so-called echo pulse, which is an additional π -pulse, $R(\pi,0)$, midway through the waiting time. Such a pulse maps the position of a Bloch vector on the equator to another point on the equator, mirrored on the yz-plane. As a consequence of this mirroring, a constant phase shift during the waiting time before the echo pulse is undone (canceled out) by an opposite phase shift after the echo pulse. The echo sequence is given by

$$R(\pi/2,0) \Rightarrow \frac{t_{\text{wait}}}{2} \Rightarrow R(\pi,0) \Rightarrow \frac{t_{\text{wait}}}{2} \Rightarrow R(\pi/2,\phi),$$
 (2.97)

which cancels out linear phase drifts. More complex phase drifts can be canceled with more elaborate echo schemes [97, 98].

Coherence time is inherent to the states chosen to encode the qubit. For example, the frequency of different transitions are unequally sensitive to fluctuations in magnetic fields and will exhibit different coherence time in the presence of magnetic field noise. One can therefore perform coherence time measurements on various transitions to infer the contribution of magnetic field noise to decoherence. Additionally, a ground state qubit is insensitive to laser field noise, and thus also serves as a good indicator of magnetic field noise. A coherence measurement of a ground state qubit (using only optical transitions) is done with the sequence

$$R(\pi/2,0)_{|0\rangle\leftrightarrow D} \Rightarrow R(\pi,0)_{|1\rangle\leftrightarrow D} \Rightarrow t_{\text{wait}} \Rightarrow R(\pi,\pi)_{|1\rangle\leftrightarrow D} \Rightarrow R(\pi/2,\phi)_{|0\rangle\leftrightarrow D}$$
(2.98)

where the subscript denotes the transition that the rotation operation applies to (see Section 2.2.7).




Figure 2.12: Schematic of MS gate levels and transitions

Measuring of coherence between states is not limited to electronic states of the ion: The coherence of motional states can be queried. This is particularly useful for estimating fidelity of gates that make use of the motional modes, such as entangling gates. Motional coherence is measured with the sequence

$$R(\pi/2,0) \Rightarrow R(\pi,0))_{\text{rsb}} \Rightarrow t_{\text{wait}} \Rightarrow R(\pi,\pi))_{\text{rsb}} \Rightarrow R(\pi/2,\phi)$$
 (2.99)

where the subscript rsb denotes a pulse on a red sideband. Such sequences also have variations that include spin echos.

2.3.6 Entanglement

Qubit entanglement plays a vital role in QC. A multitude of techniques to generate ion-ion entanglement exist [67, 79, 99–104], most of which make use of multiple ions' shared motion to mediate the gate action. One of the commonly used entangling gates for *optical* qubits is the Mølmer-Sørensen (MS)-gate [79], which produces an entangled Bell-state. The MS-gate is described thoroughly in Ref. [105] and is qualitatively summarized below. The desired gate operation is described in Section 2.2.7. The operation is realized by a bichromatic light-field which addresses two or more ions simultaneously. The bichromatic light field has a frequency component that is detuned by $-\delta$ from a red motional sideband and a second frequency component that is detuned by δ from a blue motional sideband. The levels scheme of two ions under such a light-field interaction is shown in Figure 2.12. The starting state $|00\rangle|n\rangle_n$, with mode occupation number n, is excited through $|10\rangle|n+1\rangle_n+|01\rangle|n+1\rangle_n$ to $|11\rangle|n\rangle_n$. Due to the detuning from the sidebands δ , the states containing $|n+1\rangle_n$ destructively interfere after time $t=2\pi/|\delta|$. The two-photon process that leads to $|11\rangle|n\rangle$ is effectively a Raman transition and does not destructively interfere. An appropriate choice of beam power creates a superposition of $|00\rangle|n\rangle_n$ and $|11\rangle|n\rangle_n$. The beam power is set to $\Omega=\delta/(4\eta)$, where Ω is the *carrier* Rabi frequency and η the Lamb-Dicke parameter (see Section 2.2.6).

At the gate time $t=2\pi/|\delta|$, the Bell state $1/\sqrt{2}(|00\rangle+i|11\rangle)$ is generated. To validate that this is a *pure* state, a global analysis pulse is applied to the output state of the MS-pulse. Using the notation introduced in Section 2.2.7, the analysis pulse operation is given by $R(\pi/2,\phi)$, with a variable phase ϕ . At phase $\phi=0,\pi/2,\pi...$, this pulse ideally produces an even distribution of populations of all possible states, $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. At phase $\phi=\pi/4$, odd terms ($|01\rangle$ and $|10\rangle$) destructively interfere, and population remains in $|00\rangle$ and $|11\rangle$. At $\phi=3\pi/4$, even terms destructively interfere, and population is placed in $|01\rangle$ and $|10\rangle$. The amount of occupation in either the even states or the odd states is quantified with the so-called *parity*. The state's parity is given by $P_{|00\rangle}+P_{|11\rangle}-P_{|01\rangle}-P_{|10\rangle}$, which has the value +1 when only even states are occupied and -1 when only odd states are occupied.

Scanning the phase of the analysis pulse and analyzing the parity of the resulting state reveals the purity of the entangled state, since phase coherence is required for destructive interference. The contrast in such a *parity scan*, together with the measured state population, are used to determine the fidelity of the MS-gate operation. Fidelity [11] is a measure of how close a measured state, characterized by the density operator ρ , is to the intended state, in the case of the MS-gate given by $|\psi\rangle=1/\sqrt{2}(|00\rangle+i\,|11\rangle)$. The fidelity F is given by $F=\sqrt{\langle\psi|\,\rho\,|\psi\rangle}$.

A common mistake that people make when trying to design something completely foolproof is to underestimate the ingenuity of complete fools.

— D. Adams, Mostly Harmless

This chapter covers the design of the experimental setup, which aims to demonstrate the prospect of scalability in QC. The setup is a cryogenic apparatus that houses a planar segmented surface trap, with the ability to trap and manipulate two different ion species, 40 Ca $^{+}$ and 88 Sr $^{+}$.

The cryogenic setup has been endearingly named 'Cryostina¹' by its initial designer, Matthias Brandl. The setup and many of the considerations that went into designing it are described in detail in the publication Ref. [106], and in greater detail in the thesis Ref. [107].

This chapter covers, in lesser detail, much of the setup, and highlights some of the notable upgrades that have been made since those publications. Throughout the course of my Ph.D. research, the setup has undergone many refurbishments and has seen multiple different ion traps, the majority of which have had the pleasure of trapping ions. The research discussed in further chapters of this thesis has been carried out at various stages of the setup's lifetime. An overview of the evolution of the cryogenic setup and ion traps is shown in Figure 3.1.




Figure 3.1: Timeline of various versions of the cryostat setup and traps. Trap images taken from [108-110]

¹ Cryostina follows our group's naming convention that cryostats should have human names with cold puns. Other examples include Coolien, Iceabelle, and Marcold.

-A bit of history-

The cryogenic vacuum chamber was first closed for operation in early-2014. At this point, a trap provided by Ferdinand Schmidt-Kaler's ion trapping group in Mainz was installed. Unfortunately, after several months of fruitless attempts to load ions, the trap was discarded. As it was becoming critical² to demonstrate that the experimental setup was at a stage in which ions could be reliably trapped, we sought out a quick alternative. A 'Yedikule' trap was provided by the cryogenic trap team in our group, courtesy of Michael Niedermayr [108]. This trap stably stored the setup's first calcium and strontium ions. While a good test-bed for learning the ropes of aligning optics, tuning trapping parameters, and trap characterization, the Yedikule trap ultimately would not be useful for QC experiments, as it lacked the necessary optical access to address single qubits without adverse laser scattering. Preparations were being made for a newer version of the cryogenic setup, and for a new trap.

After a solid year of designing, measuring, and otherwise preparing the upgrades, the new version of Cryostina was assembled, with a trap designed and provided by Hartmut Häffner's group at University of California, Berkeley (UCB)³. This version of the experimental setup was in use for several years. Coherent operations on the qubit transitions of both ⁴⁰Ca⁺ and ⁸⁸Sr⁺ were available, allowing us to characterize heating rates and coherence times, and set up gate sequences, even doing mixed-species gates. Furthermore, the experiment acquired an upgraded version of the arbitrary waveform generator (AWG) voltage supply, the 'Bertha', provided by Prof. Schmidt-Kaler's ion trapping group in Mainz. This allowed us to investigate ion transport, which led to an in-depth study of coherent ion crystal rotations and is the topic of Chapter 5.

In the meantime, preparations were made for another experimental adaptation, which was assembled in December, 2017. This version exhibited multiple upgrades, such as improved DC wiring and filtering. The main design goal, however, was to incorporate a new trap: Sandia National Laboratory's High Optical Access (HOA) trap [109]. The package holder, which uses electrical connectors known as fuzz buttons⁴ in a standardized land grid array, was designed with the idea that trap replacement would be reasonably simple. As it turned out, repeated trap replacement was a necessary operation (and was hardly "reasonably simple"). The first HOA trap never held ions, and the subsequent three HOA traps all exhibited high heating rates, and laser-induced trap charging, neither of which being desirable for QC. Also, we were unable to load calcium and strontium ions simultaneously, which very much defeated the purpose of a setup intended for mixed-species operation.

To bring the experiment back on the path of operating in the scope of quantum computation, a new trap was designed. The new trap, the 'Golden Gate' trap, was based on the already proven design of the Berkeley trap. It was bonded onto a chip carrier with the same footprint as the HOA trap, and was thus straightforward to plug in to the existing setup. Shortly after the installation, ions were trapped⁵, even mixed-species chains. This trap was and will be used for multiple projects, including an investigation of RF heating and recrystallization of ion clouds (Chapter 4), and dissipation through engineered resonance with mixed-species ion chains (Chapter 6). Additional projects, to be covered in the Ph.D. thesis of Lukas Gerster, include Bayesian gate optimization [111] and using phase gates for generating a higher-dimensional version of qubits, qudits. [112]. At the moment of writing, the Golden Gate trap is still in operation in our experimental setup.

The following section gives an overview of the experimental setup, specifically the latest version which has been used for most of the work presented in this thesis. The ion crystal rotation project (Chapter 5) was done on an earlier version of the setup. The description of the setup is, however, for the most part applicable to all experiments described in this thesis. One notable difference is the electronic filters of the direct current (DC) electrode voltage lines, which were redesigned based on experimental difficulties present in ion crystal rotations. The filters used during these experiments are therefore explicitly discussed in Chapter 5.

² Note 8 in Appendix e

³ Note 9 in Appendix e

⁴ Custom Interconnects

⁵ Even with accidentally scrambled DC electrode wiring

3.1 SETUP OVERVIEW

This section covers the experimental setup used for the experiments presented in this thesis. Most of what is described in this section is also covered in detail in [107]. Section 3.1.1 describes the cryogenic apparatus in which the surface trap is mounted. Section 3.1.2 describes the laser setup, and Section 3.1.3 covers the electronics and software that make up the experimental control.

3.1.1 Cryogenic apparatus

Figure 3.2 shows a schematic overview of the cryogenic apparatus, whose components are discussed below.

3.1.1.1 *Cryostat*

While not considered an absolute requirement, QC with ions trapped in a surface trap greatly benefits from a cryogenic setup. A primary reason is the known correlation between anomalous excitation of a trapped ion's kinetic energy and trap surface temperature [34, 43, 113]. Another benefit of a cryogenic setup is improved vacuum pressure due to cryopumping [44]. Gas particles condense on cold surfaces, and are thus removed from the vacuum volume. At temperatures of 4K and below, most molecules that are typically prominent in vacuum setups experience a significant drop in partial pressure. At this temperature, achievable with modern helium-based cooling techniques, vacuum pressures below 7×10^{-17} mbar are reachable [114]. Typical ion trapping setups used for QC aim to have ultra-high vacuum (UHV) environments below the 1×10^{-10} mbar range, to have acceptably low rates of collisions between trapped ions and background gas particles. The fact that such low pressures can be reached with cryopumping brings another advantage to the table: Cryogenic systems can benefit from a low turn-around time. After breaking vacuum in order to install a new trap or perform other setup upgrades, a room-temperature setup would ordinarily require several weeks of baking in order to reach a sufficiently low pressure. The ability to cryopump, however, eliminates the need for a bake-out. Turn-around time is then only limited by the time required to temperature-cycle the system, and the assembly/reassembly time. In our setup, we ideally do not wish to make use of this benefit: once a trap is installed, it should stay in. However, the performance of surface traps is difficult to predict before testing it with ions, and it may take several iterations of trap replacement until a suitable trap is installed.

There are many types of devices available that provide cryogenic environments. The main contenders for reaching the 4 K range or lower can be roughly subdivided into three types: 1) Dilution refrigerators, 2) closed-cycle cryocooling, and 3) wet cryostats. Dilution fridges are massive, complicated, offer very limited optical access [115, 116], and are therefore impractical for trapped-ion experiments. The achievable temperature is in the sub-Kelvin range, which for trapped ion purposes is considered 'overkill.' More commonly used in ion trapping experiments are closed-cycle cryostats, which operate through diabatic expansion and compression⁶ of helium gases. These types of cryostats benefit from a relative ease-of-use and do not require a supply of liquid cryogens. Unfortunately, their operation requires moving parts to drive the expansion and compression cycle, and thus introduces a source of acoustic, vibration, and magnetic field noise in the lab. Depending on the experimental requirements, these sources of noise may be tolerable. However, such machines rarely produce low enough vibrations and magnetic field noise to achieve the necessary qubit coherence associated with quantum information processing. Low-noise operation with closed-cycle systems for QC has been achieved, but requires extreme measures to decouple sources of noise from the ion trap [118] while maintaining high cooling power.

With wet cryostats, the ion trap is cooled down by thermally connecting it with a liquid coolant. For operation in a sufficiently low temperature (4 - 20 K, below the vapor pressure of most common background gas molecules [119]), the choice of liquid coolant boils down to liquid Helium. Wet cryostats can be further subdivided into bath cryostats and flow cryostats. In the former, a part of

⁶ Essentially the same principle of operation as a kitchen refrigerator, which would be considered a cryostat if it managed to cool food down below 120 K [117].




Figure 3.2: Cross-sectional view of experimental apparatus

the setup is simply submerged in a supply of helium (for example, the super-conducting magnets in MRI machines). Flow cryostats provide control over the flow of the cryogenic liquid, as the name suggests, thus allow a higher degree of control of operation temperature. Flow cryostats do not require moving mechanical components, thus do not produce magnetic field noise, or acoustic noise in the lab. While not immune to mechanical vibrations (a boiling liquid does, after all, vibrate), these are at a much more manageable level compared to closed-cycle systems. All advantages and disadvantages considered, given the finite lab space, and low-noise requirements of our setup, a flow cryostat has been chosen⁷. The cryostat uses liquid Helium as a coolant⁸.

As shown in Figure 3.2, the flow cryostat is installed in the top of the apparatus. The metal shaft provides an insert for a transfer line (not pictured) which mediates the flow from the liquid storage to the coldfinger, where it evaporates. The coldfinger provides the main source of cooling power for the apparatus, and is thermally connected to the trap. Two layers of radiation shielding limit radiative heat transfer from the surrounding room-temperature vacuum chamber to the trap. The trap sits within, and is thermally directly connected to, the inner heat shield (\sim 30 K), which in turn is surrounded by the outer heat shield (\sim 120 K).

The coldfinger is thermally connected to the inner shield with oxygen-free high conductivity (OFHC) copper wires. We have used a collection of thin wires⁹, crimped in cable lugs. Using a collection of thin wires allows us to mechanically decouple the coldfinger from the inner heat shield, and hence the trap, while maintaining good thermal contact between the two. The inevitable vibrations occurring during the evaporation of liquid coolant in the coldfinger are attenuated by the flexible copper wires. In total, 12 crimped collections of copper strands are used to connect the coldfinger and inner shield. The connecting rod between the main inner shield body and the connection to the coldfinger is used as thermal anchoring [120] for the wiring that provides DC voltages to the trap. The inner shield is discussed in more detail in Section 3.1.1.2.

The outer shield is divided into three sections, the upper, mid, and lower outer shield. The upper outer shield provides a secondary source of cooling power for the setup: A clamp above the exhaust enables heat transfer between evaporated coolant and the upper outer heat shield. Similar as before, crimped OFHC copper strands are used to thermally couple the upper and mid outer shield, while suppressing vibrations originating from the flow cryostat. The mid shield provides a large surface area for thermal anchoring of the DC electrode voltage wiring. Additionally, the mid shield houses the electrical filters for this wiring. The lower outer shield does not have any responsibilities other than being a radiation heat shield.

The flow cryostat has a CF flange connection to the vacuum chamber by means of a metal bellow, and a set of rubber spacers separating the two. The bellow, a ribbed deformable metal, attenuates vibrations between the cryostat and vacuum chamber.

The previously described copper strands and the bellow mechanically isolate the cryostat from the lower shields and vacuum chamber in which the trap is placed. The trap is thus mechanically loosely' connected to the cryostat, which ensures that vibrations caused by the cryostat are minimized at the trap. The trap must, however, be rigidly connected to the optical table, to ensure that trapped ions do not vibrate with respect to the optical beam paths. The inner and outer shields are therefore rigidly connected to the vacuum chamber, which is mounted to the optical table. An ideal trade-off between an as-high-as-possible rigid connection and an as-low-as-possible cross-section to minimize heat transfer is a stainless-steel hexapod structure. The hexapod, seen in the central column above the inner shield in Figure 3.2, holds the outer and inner shield in place with respect to the vacuum chamber, which is a lot of heavy lifting for just a few narrow steel tubes.

The ion trap is mounted inside the inner shield, at the bottom of Figure 3.2. The vacuum chamber and outer and inner shields provide optical access for various beams and fluorescence detection, from 8 different directions all within horizontal plane. The vacuum chamber is placed in a hole in the custom designed optical table such that this horizontal plane in which the trap sits coincides with the height of the optical components on the table.

⁷ Model ST-400-1, Janis

⁸ Note 10 in Appendix e

⁹ Stripped from a high-quality audio cable

3.1.1.2 Inner shield

The inner shield has the following purposes and requirements:

- Shield the trap from thermal radiation
- Shield ions from magnetic field fluctuations
- Maintain a high numerical aperture optical access to trapped ions to enable single ion addressing and fast state readout.

The design considerations of the inner shield are, as with much of this section, described in detail in Ref. [107]. Since that publication, the experimental setup has undergone a refurbishment, which will be (or is, depending on when this is being read) described in the Ph.D. thesis of Lukas Gerster [121]. A short overview of the latest version of the inner shield is provided here.

A schematic of the inner shield is shown in Figure 3.3. The shield is a copper cylinder with two-centimeter thick walls. Due to the skin effect, this thick copper wall attenuates magnetic field fluctuations. This attenuation is enhanced with the inner shield in the cryogenic environment, since the conductivity of the cold copper walls increases by orders of magnitude with respect to conductivity at room temperature. Measurements that determine the magnetic field attenuation are described in Ref. [106]. The inner shield is composed of two separate parts, as shown in Figure 3.3(a), to provide access inside of the shield during assembly. The parts have been chosen to separate in a vertical plane, as opposed to a previous version of the inner shield that opened horizontally, from the bottom. This decision ensures that Eddy currents around the axis corresponding to the externally applied magnetic field (the quantization axis) do not cross any seams in the shield parts. These currents therefore have the lowest resistance around the quantization axis, and the highest magnetic field attenuation.

Effective attenuation of magnetic field noise and shielding of radiation requires that the inner shield is as enclosed as possible, meaning that the surface area of 'holes' in the shield should be kept to a minimum. This requirement is in conflict with the necessity to have good optical access of the ion. Individual ion addressing requires a beam to be tightly focused, such that neighboring ions, typically spaced several micrometers apart, are not illuminated. This means the beam must have a high convergence angle, or in other words, a large numerical aperture (NA). For a 729 nm beam to achieve a waist diameter of 2 µm, an NA of about 0.22 is required.

Additionally, efficient state detection requires a large amount of fluorescence to be collected from ions, which also implies a large NA. This high optical access is achieved by placing lenses inside of the inner shield, shown in the cross-sectional cutout in Figure 3.3(b) (Note, this cross-section does not align with the cut between shield parts in (a)). A copper plate inside the inner shield serves as a lens mount. The same plate is a holder for the trap assembly, including the resonator, discussed later in this section. The lens mount holds lenses on both sides of the surface trap. The trap, also discussed later, is positioned such that the normal of the surface plane is horizontal. The lenses are placed 'in front of' and 'behind' the surface of trap, meaning that the lenses' principal axes align with the surface normal. The trap has a slot in the surface to allow imaging and addressing from the back side. The back lens has an NA of 0.25 and a focal length of 25 mm. The front lens is discussed below.

One of the design goals of Cryostina V₃ was to push state readout to its limits, by attempting to collect fluorescence from nearly a full semi-spherical solid angle of the ion source. A 25 mm o.83 NA lens¹⁰ was chosen for this purpose. The lens has a focal length of 15 mm. The short focal length makes it such that the lens must be placed in a very precise position with respect to the ion trap: Ideally, when the lens is placed 15 mm away from the ion, the ion's fluorescence is collected and collimated. Similarly, the intended use of the lens is to focus a *collimated* laser beam onto an ion. However, if the lens were 1 mm closer to the ion, the collected fluorescence diverges, and more than half of the fluorescence intensity does not escape through the apparatus's viewports. Lens placement is therefore critical for good ion imaging and addressing. A similar argument can be made for the radial placement of the lens, i.e. ensuring that the lens's principal axis coincides with the ion.

Unfortunately, it is difficult to perfectly predict the lens's position after assembly and cooldown of the cryostat. We have therefore opted to place the lens on a stack of two nanopositioners¹¹, which

¹⁰ Edmund Optics, #67-245

¹¹ Attocube ANPx51 and ANPz51




Figure 3.3: Overview of the inner shield

allows us to control the lens's height and distance from the trap. The inner shield is not big enough to include a third nanopositioner to control the third translational degree of freedom. Since this direction corresponds to the trap's axial direction, we can instead move the ion to align it with the lens's principal axis.

The lens is held in a metal mount with a set of plastic screws. The metal mount has extruded arms in which the lens is placed, designed to minimize mechanical stress when the mount contracts when cooled down. The lens is not centered above the nanopositioner stages. To reduce torque on the nanopositioners, a counterweight is made by extruding the lens mount towards the rear, away from the lens.

The lens has a back focal length (the distance between the flat side of the lens and the focal point, the ion) of 8.4 mm. From a trapped ion's point of view, this is a very close proximity for a large dielectric material to be [31]. Charging of the lens can alter the electric field that the ion experiences, which can alter the ion's position and motional frequency, and induce motional excitation. In a previous version of the setup, Cryostina V2, a metal clamp above the trap was used to shield these fields. The opening in the clamp limited the optical access to the trap to an NA of approximately 0.25. As this defeats the purpose of the 0.83 NA lens, another solution is required.

One option is to coat the lens with a conductive film, such as indium-tin-oxide (ITO). A typical ITO-coated substrate is has around 85% transmissivity, and about $10 \times 10^{-4} \,\Omega$ cm resistivity [122]. The 85% transmissivity is tolerable but not preferable. Instead, we place a grid of thin gold wires midway between the back of the lens and the trap. Such grids are commercially available, but offer up to about 95% transmission for our desired spacing. We decided we could do better, and handmade a gold wire grid, shown in Figure 3.4(a). Gold wires with a diameter of 20 μ m, normally used for wire-bonding, are spanned across the clamp that holds the trap carrier in place. The clamp has recessed grooves in which the gold wires can be placed, so that they are neatly equidistant, with a 1 mm spacing. The wires are glued in place at the grooves. The clamp itself is made of titanium and is electrically grounded.

To determine the effectiveness of the shielding grid, electrostatic simulations are done using Comsol, which models the electric field of the lens-grid-trap setup. A test voltage of 1 V is applied to the surface of the lens. This corresponds to a uniform surface charge of $0.05\,\mathrm{C}\,\mu\mathrm{m}^{-2}$, which is motivated by Ref. [123]. Figure 3.4(c) shows the electric field (in the direction normal to the trap surface) between the trap (kept at 0 V) and the lens, shown with and without the grounded mesh grid in place. The placement of the trap, grid, and lens are shown schematically in Figure 3.4(b), which also shows how the field potential is shielded by the mesh. If the mesh were not there, the potential transitions more or less linearly from 1 V to 0 V from the lens to the trap. The dashed line in Figure 3.4(c) and the blue dot in (b) show the height at which the ion is trapped, 110 μ m above the surface of the trap. We conclude that the trap region experiences a tenfold reduction in electric field noise coming from the lens when the grid is in place.

The lens-grid-trap setup has been similarly studied using the software Zemax to determine the optical properties of the grid. The grid has a 98% transmission. A beam passing through the lens, focused onto the ion, experiences negligible distortion due to the grid.

Ultimately, the high-NA lens has been utilized for state-readout of 40 Ca⁺ions. We have found that we can distinguish the qubit state with a certainty of more than 99.99%, with a 50 µs detection time. This is more than an order of magnitude faster than typical detection times in other ion trap setups.

3.1.1.3 DC voltage supply

Trapping ions in a surface trap requires a combination of RF and DC voltages to be applied to trap electrodes. This section describes the DC voltage supply. We use the term DC to contrast it with the RF supply, but note that these voltages are not necessarily static throughout an experimental sequence. In particular, in Chapter 5 ion positions are manipulated during a sequence by applying time-dependent voltages to trap electrodes.

A schematic overview of the wiring and filtering for the DC voltage supply is shown in Figure 3.5. Voltages are supplied by a multi-channel AWG [124], provided to us by the group of Ferdinand Schmidt-Kaler at the Johannes Gutenburg University, Mainz¹². A centralized field programmable

¹² Where the device has acquired its name, Bertha




Figure 3.4: (a) Picture of the gold mesh grid for shielding the trap from stray electric fields. (b) Electrostatic simulations indicating how the grid (not to scale) shields the potential of a uniformly charged lens surface. (c) Electric field perpendicular to the trap plane, with and without the grid in place. The dashed line indicates the ion position.

gate array (FPGA) controls the digital voltage waveform of 80 digital-to-analog converter (DAC) output channels. The DACs have a $\pm 40\,\mathrm{V}$ output range, with a voltage resolution of 1.2 mV. We have not measured noise characteristics directly. See Ref. [124] for noise characteristics on a similar device.

Digital voltage waveforms, in this thesis usually referred to as *voltage sequences*, have a maximum sample rate of 2 MHz. Sequences can be initiated by a transistor-transistor logic (TTL) trigger signal. Predetermined voltage sequences can thus be used as sub-steps of experimental cycles. In our setup, however, the voltage sequences must be pre-computed and cannot be altered 'on-the-fly' between experimental cycles. Section 3.1.3 provides more information on how the AWG is implemented within the context of experimental control.

Voltages output by the AWG are redistributed, into four handmade 25-channel cables, designed to optimally shield the voltage-carrying wires from external electric field noise. The cables are optionally passed through a set of filters. These third-order RC filters ($R = 100 \,\mathrm{k}\Omega$, $C = 120 \,\mathrm{nF}$) have a cut-off frequency at 12 Hz, and are removed when experiments include dynamic voltage sequences. Otherwise, when only static voltages are required, the filters are kept in, since they reduce motional mode heating of the ion (see Section 2.3.4).

A flange on the vacuum chamber with four 25-pin D-Sub connectors connects the air side and vacuum side. On the vacuum side, it is desirable to use wires with a non-negligible ($\gtrsim 1\,\Omega$) electrical resistance, to reduce thermal coupling between the room-temperature vacuum chamber and cryogenic trap. Previously, individual strands of phosphor-bronze were used, which were impossible to keep untangled, similar to having 50 headphone cables in your pocket. To simplify the cable management, the in-vacuum DC lines are distributed over four loom cables 13. Outgassing rates of these cables are low enough for use in cryogenic operation, but are not recommended for room-temperature UHV experiments.

The mid outer shield has grooves on its outer surface to which the loom is clamped, before passing into the shield. This thermally anchors the loom to the outer shield temperature. Inside the mid outer shield, there are four sets of filter boards, each with a micro-D-Sub input and output. The filter boards are second-order RC filters for each DC line, with $R=500\,\Omega$ and $C=1.2\,\mathrm{nF}$. The cutoff frequency, $f_c=270\,\mathrm{kHz}$, is chosen to filter noise at the ions' motional frequencies (in the MHz range), while allowing voltage waveforms that transport ions on tens of μ s timescales to pass through. An RC

¹³ Custom made manganin Cryoloom, micro-D-Sub terminated, CMR-direct




Figure 3.5: Schematic of DC electrode voltage wiring and filtering

filter circuit has been chosen instead of the previous version's LC filter. The reason for this change is the behavior of the ground wires, which in reality are not perfectly grounded due to finite cable resistance. This non-negligible wire resistance creates an effective crosstalk between electrodes when time-dependent voltages are applied. This effect is discussed in more detail in Section 5.4.3.2, where the filtering is put in context of generating sequences of electrode voltages for manipulating the position of ions. Here, it is discussed how the effective crosstalk across electrodes of time dependent voltages affects the design of the sequence. The main lesson learned is that designing voltage sequences is much simpler when the effective crosstalk is minimized. Circuit simulations with the Python PySpice package have shown that RC filters outperform LC filters in this regard.

After the set of outer shield filters, the wires are thermally anchored to the outside of the inner shield, enter the inner shield, and are attached with micro-D-Sub connectors to a printed cicuit board (PCB) on which the trap carrier is mounted. Capacitors (4 nF) are placed on the PCB, which shunts parasitic RF noise picked up from the trap's RF drive. The PCB is mounted on the same mount that holds the in-vacuum lens assembly, described in the previous section. The DC lines on the PCB terminate at pads, laid out in using the standardized 100-pin PGA footprint. The PCB has a ceramic mount to hold a trap chip and function as an electrical interposer. The mount has integrated Fuzz buttons in the carrier mount electrically connect the trap carrier and the PCB. The design of the trap carrier, and the incorporation of the trap, are covered in Section 3.2.

3.1.1.4 RF voltage supply

The functionality of Paul traps, also known as RF traps, require an RF voltage to be applied to dedicated trap electrodes. In this section, the RF supply is discussed.

The voltage amplitude of RF electrodes of surface traps is typically on the order of hundreds of volts, supplied at tens of MHz. Such voltages are often reached with the help of RF resonators, which produce a voltage gain of an input RF signal at a specified frequency. At the highest level, an RF resonator is an LC-circuit. The capacitance *C* is for the most part given by the trap's RF electrodes. The inductance comes from a coil, where *L* is chosen such that the desired resonance frequency,

 $\Omega_{RF} \approx 1/\sqrt{LC}$ is achieved. An externally applied electric field can resonantly build up power in the LC circuit. Losses in this LC circuit, caused by resistance in the wiring and loss in the trap's dielectric layers, limit the power that can be stored in the resonator. The performance of resonator is characterized by its quality factor and voltage gain. The quality factor indicates how narrow the frequency response of the resonator is. It represents the relative amount of energy lost in an oscillation cycle, and is proportional to the square of the voltage gain. Aiming for a high quality factor, and thus a high voltage gain, minimizes the amount of input power needed to reach the desired voltage on the trap's RF electrodes. Less supplied RF power means less heating of the trap, which is desirable for stable cryogenic operation. The resonator's quality factor is optimized by minimizing the effective resistance in the LC circuit. It is therefore imperative to place the coil as close as possible to the trap, and to use low-resistance materials in the circuit.

If the coil were to be placed outside of vacuum, as is commonly done with room-temperature traps, the required low resistance electrical connection between the coil and the trap would present a tremendous heat load directly on the cryogenic trap. Operation with an out-of-vacuum coil is therefore not an option. It should instead be placed at the same temperature stage as the trap, namely in the inner shield. Space in the inner shield is limited, so implementing a bulky helical coil, a commonly used type of coil for traps operated at room temperature, is impractical.

The suitably-sized inductor that we have landed on is a PCB-style toroidal wire coil¹⁵, a brainchild of Matthias Brandl. The design considerations and characterization of the resonator coil are described in detail in Ref. [125], and will not be covered here. We have chosen a 63 loop toroidal coil, using 0.6 mm silver wire. The coil was hand-woven into a ceramic substrate with holes milled out. An additional coil loop acts as an inductive pickup wire, allowing us to monitor the power in the LC circuit. A PCB that holds the RF matching network [125] is attached directly to the coil holder. The coil and matching network assembly fit in the inner shield mount, directly below the trap PCB, see Figure 3.3. A small hole allows the coil's RF and ground wire to pass upwards to be soldered onto the trap PCB.

The matching network minimizes reflections in the RF signal in the transition from the supply cable to the resonator. We use an additional LC circuit for matching, using the self-inductance of the input line, and a set of capacitors to ground. The capacitance must be chosen such that impedance matching is optimized for the electrical properties of the resonator at cryogenic temperatures, which differ from those at room temperature. Therefore, before installing the resonator, several matching tests have been done with a dummy¹⁶ trap in room-temperature and in liquid nitrogen baths, at 77 K. We then extrapolate which matching capacitance will minimize RF reflections at the intended cryogenic temperature, about 30 K. We intentionally choose a capacitance for ideal matching at a higher temperature than the inner shield temperature, since the resonator is expected to heat up when supplied with RF power. In our experiment, we see that at the start of the day when first switching on the RF supply that the resonance frequency drifts over the course of about an hour by a few percent and that the on-resonance reflected RF signal decreases.

At stable cryogenic operation, the resonator operates at 40.38 MHz, has a voltage gain of about 80, a quality factor of about 150, and less than 5% reflected power.

3.1.2 Laser setup

In this section, an overview of the lasers and the optical setup used in our experiments is given. We have sets of lasers that allow us to trap two ion species, ${}^{40}\text{Ca}^{+}$ and ${}^{88}\text{Sr}^{+}$. The lasers used in our experiments are summarized in Table 3.1, and the relevant level diagrams are displayed in Figure 3.6. The lasers are placed on optical tables that are separated from the cryogenic ion trap apparatus.

3.1.2.1 Laser tables

The wavelength of most beams are monitored and stabilized by a wavemeter¹⁷ which includes an eight-channel fiber switch, and a proportional-integral-derivative (PID) module that feeds back onto

¹⁵ Note 11 in Appendix e

¹⁶ i.e. a trap whose RF electrode has a similar capacitance as the actual trap, but subjected to out-of-vacuum mistreatment that we wouldn't want to put the real trap through.

¹⁷ HF-WS/8-2, Toptica




Figure 3.6: Level schemes of Calcium and Strontium showing relevant levels for photoionization of atomic ${}^{40}\text{Ca}^{+}$ and ${}^{88}\text{Sr}^{+}$, and levels used in trapped ion experiments.

	λ (nm)	Description	Transition
⁴⁰ Ca	423	Photoionization	$4S_0 \leftrightarrow 4P_1$
	375	Thotolomzation	$4P_1 o \infty$
	397	Doppler cooling/ Detection	$4S_{1/2} \leftrightarrow 4P_{1/2}$
	866	Repumping	$3D_{3/2} \leftrightarrow 4P_{1/2}$
	854	Repullipling	$3D_{5/2} \leftrightarrow 4P_{3/2}$
	729	Qubit control	$4S_{1/2} \leftrightarrow 3D_{5/2}$
⁸⁸ Sr	461	Photoionization	$5S_0 \leftrightarrow 5P_1$
	405	Thotolomzation	$5P_1 \rightarrow 5D_2$
	422	Doppler cooling/ Detection	$5S_{1/2} \leftrightarrow 5P_{1/2}$
	1092	Repumping	$4D_{3/2} \leftrightarrow 5P_{1/2}$
	1033	Tepumping	$4D_{5/2} \leftrightarrow 5P_{3/2}$
	674	Qubit control	$5S_{1/2} \leftrightarrow 4D_{5/2}$

Table 3.1: Overview of lasers used in the experiment

some of the laser drivers, through which their wavelength is stabilized. This is an upgrade with respect to Ref. [107], in which lasers' wavelengths were locked to an optical cavity, by means of a Pound-Drever-Hall (PDH) lock. Now only the beams that are responsible for coherent qubit manipulation are locked to an optical cavity, as a narrower linewidth is required than what the wavemeter lock can provide. The wavemeter is specified to have a 2MHz absolute frequency accuracy.

Below we provide an overview of the laser setups and tasks of the ⁴⁰Ca⁺lasers. The overview for the ⁸⁸Sr⁺lasers is nearly identical, and have respective wavelengths given in parentheses is the section headings.

-Photoionization (PI): 423, 375 (461, 405) nm-

Ions are generated with two PI beams. A neutral beam of calcium atoms are emitted from an 'oven' (a tube of solid Calcium, heated above the temperature that Calcium vaporizes in vacuum, $> 500\,\mathrm{K}$ [126]), providing a flux of atoms in the trapping region. The PI beams are continuously on while loading ions into the trap. The 423 nm beam ¹⁸ resonantly excites the $4S_0 \leftrightarrow 4P_1$ transition ¹⁹ and the 375 nm beam has an energy that exceeds what is required to ionize the excited atom. The 375 nm is a free-running diode laser. Having the two beams on simultaneously at the trapping region produces

¹⁸ Some works describe this beam as 422 nm light. To avoid confusion with the strontium Doppler cooling beam at 422 nm, in this thesis the calcium PI beam is correctly rounded to 423 nm.

¹⁹ The fluorescence emanating from the ions in this transition can be used to detect if there is an atomic flux at the trap site, and is thus used to tune the temperature of the oven, or to decide that an oven needs replacing.




Figure 3.7: Simplified schematic of beam paths on the laser tables. Laser wavelengths are shown for ⁴⁰Ca⁺and, in parentheses, ⁸⁸Sr⁺, which have nearly identical setups.

Calcium ions in the trap. Ions with low enough energy do not immediately escape the trap, and can be cooled further by Doppler cooling.

The wavelength of the 423 nm beam is monitored by the wavemeter. No feedback is applied to stabilize the wavelength, other than occasional manual fine-tuning. A suitable wavelength is initially calibrated by maximizing atomic fluorescence from the flux of neutral atoms near the trapping region. While attempting to load ions, the wavelength is swept around this calibration point by several hundred MHz.

One slight difference between the ionization process of $^{40}\text{Ca}^+$ and $^{88}\text{Sr}^+$ is that after the initial $5S_0 \leftrightarrow 5P_1$ transition, strontium happens to have an excited level, $5D_2$ at a convenient wavelength of $405\,\text{nm}$, from which auto-ionization is possible. While not strictly necessary to make use of this transition, ionization is enhanced when coupling to it [127], which allows ion loading with lower atomic flux, and therefore less contamination of the trap surface. $405\,\text{nm}$ diodes are easily available, as they can be ripped out of any obsolete Blu-ray player.

The 423 and 375 nm beams were initially overlapped using a polarizing beam splitter (PBS) on the laser table, as depicted in Figure 3.7, before being sent into an optical fiber to the experiment table. In a revised version of the setup, the 375 nm diode laser was moved to the experiment table, and is overlapped with the 423 nm beam there. The Strontium setup is still as indicated in Figure 3.7.

-Doppler cooling and state detection: 397 (422) nm-

Doppler cooling and state detection is done using the $4S_{1/2} \leftrightarrow 4P_{1/2}$ dipole transition, at 397 nm. This transition has a short natural lifetime, 7.7 ns. The 397 nm light is provided directly by a diode laser (as opposed to the previous frequency-doubled 794 nm setup). The wavelength of the laser output is monitored and stabilized by the wavemeter. The beam is split into two arms, each passing through a separate double-pass acousto-optic modulator (AOM) setup. One arm, with a total frequency shift of 200 MHz, is used as the main Doppler cooling and detection beam. The second arm, the so-called 'refreeze' beam has a frequency shift of 120 MHz, and is thus red-detuned by 80 MHz with respect to the main Doppler cooling beam. The 'refreeze' beam is a high-power beam, far-red-detuned from cooling transition. The reason for including this beam is discussed in detail in Chapter 4, but to summarize here: The main Doppler cooling beam is often not sufficient for cooling high-energy ions to the point that they form an ion chain. The refreeze beam aids in this process. The main Doppler beam and the refreeze beam are recombined with a PBS, and sent through an optical fiber to the experiment table.

-Repumping: 866, 854 (1092, 1033) nm-

The cooling transition is not a closed cycle: there is a 6% chance that the spontaneous decay from $4P_{1/2}$ is to $3D_{3/2}$, which is a meta-stable state. Population is repumped to the $S \leftrightarrow P$ cycle with 866 nm light, which couples the states $3D_{3/2}$ and $4P_{1/2}$, thus returning population to the cooling cycle states.

Population can additionally be 'trapped' in the $3D_{5/2}$ state, for example at the end of an experimental sequence in which this level is used as one of the two qubit states. Repumping back to the cycling transition is done with 854 nm light, which couples $3D_{5/2}$ and $4P_{3/2}$. 854 nm light is also used for dissipative state transfer used for ground state preparation (optical pumping) and resolved sideband cooling (see Section 2.2.8).

Both repumping beams are sent through a double-pass AOM setup, and are combined with a PBS before being sent through an optical fiber to the experiment table. The AOMs function as switches, allowing us to turn on and off the light at the ions within experimental sequences (see Section 3.1.3).

-Qubit manipulation: 729 (674) nm-

Most of our experiments use the $4S_{1/2}$ and $3D_{5/2}$ levels as the two qubit states. Coherent transfer between these states is done with pulses of 729 nm light. Qubit coherence, and thus the fidelity of quantum operations, require the 729 nm light to be stable in phase, such that the laser linewidth is in or below the single-Hertz regime. The wavemeter lock, used to stabilize the frequency of most of the other lasers, has an absolute frequency accuracy of 2 MHz, and a feedback bandwidth of 1 kHz, at best. This is by no means accurate nor fast enough for locking the 729 nm light for coherent control of qubits. Instead, a high-finesse ($\mathcal{F} \approx 250,000$) optical cavity is used as a frequency reference. The 729 nm light is locked to this cavity using the Pound-Drever-Hall (PDH) method²⁰.

The cavity mirrors are separated by ultra-low-expansion (ULE)-glass, which is temperature controlled to the material's point of zero expansion, at -0.6 °C. From beat measurements with other stabilized 729 nm sources (from the University's 'big lab' and from the IQOQI), an approximate linewidth of our 729 nm light of 10 - 15 Hz is inferred. More information about the high-finesse cavity, and the locking scheme can be found in the Master's theses of Lukas Postler [128] and Roman Stricker [129].

The cavity has a free-spectral-range at 729 nm of 3 GHz. A 240 MHz AOM is used to help bridge the gap between the frequency of one of the cavity's TEM_{00} -modes, and the frequency of the ion's qubit transition. The AOM can be used to control the intensity of light going into the cavity. A photodiode before the cavity measures this intensity. Its signal is fed into a PID-controller²¹ which feeds back onto the AOM power to stabilize the light intensity. This further stabilizes the frequency of the locked laser, since the cavity frequency is dependent on the incident laser power [130].

The frequency-stabilized light is sent to the experiment table through an optical fiber. Acoustic noise in the 20 m-long fiber alters its optical path length, thus modulating the phase of the beam at the output, which results in a broadening of the laser linewidth, referred to as fiber noise. This broadening is counteracted by fiber-noise cancellation (FNC) [131]. In short: an AOM before the fiber shifts the beam's frequency by $f_{\rm AOM}$. The beam experiences fiber noise, which produces additional phase shifts $\psi_n(t)$. A small portion of the light, reflected at the experiment-side facet of the fiber, is sent back through the fiber and subsequently through the AOM. Here the beam is recombined with a bit of light picked off before it goes through all that frequency shifting. The combined beams have a beat at frequency $2(f_{\rm AOM}+d\psi_n/dt)$, which is detected by a photodiode. The beat signal is fed into a home-built FNC-module. The module compares the beat with a stable reference signal with a set frequency $f_{\rm ref}$, and feeds back the error signal onto a voltage-controlled oscillator (VCO) that drives the AOM. The AOM frequency $f_{\rm AOM}$ is thus actively modulated to ensure that $2(f_{\rm AOM}+d\psi_n/dt)=f_{\rm ref}$.

3.1.2.2 Experiment table

All beams are routed onto the experiment table through optical fibers. Figure 3.8 gives a schematic top-down overview of the optical setup on the experiment table. To make the following description of orientations of the components in the setup less ambiguous, we refer to the cardinal directions shown in the top left of the figure²². In the center (in gray) is the vacuum chamber. The ion trap is positioned in the center of the chamber with trap surface oriented vertically, out of the page in Figure 3.8. The normal of the trap surface points south-east from the trap, and the trap's axial direction is along the north-east direction. A pair of Helmholtz coils (not shown) are placed to the west and east of the chamber, producing a magnetic field along the figure's horizontal direction.

The Doppler cooling and repumping beams for both $^{40}\text{Ca}^+$ and $^{88}\text{Sr}^+$ are overlapped into a single beam line with dielectric mirrors and PBSs. This joint line comes in from the south, and is thus aligned at 45° with respect to the trap's axial direction, and at 90° with respect to the quantization axis. The PI beams of both species are overlapped and sent in from the south-west viewport, and graze across the trap surface.

The beams that enable qubit control are first amplified by a tapered amplifier (TA). The TA output is sent through a short fiber to clean the beam mode. A portion of the beam is split off and sent to a photodiode, whose signal is used to stabilize the intensity of the beam through feedback on the TA current. The beam is split into two arms, one of which is intended for global operations, and the other for single-ion addressing. Each beam is sent through a double-pass AOM setup. These AOMs provide the primary source of frequency control of the qubit beam within experimental sequences, further detailed in Section 3.1.3. Each beam is subsequently directed through an AOM in single-pass configuration, a short mode-cleaning fiber, and a set of lenses that expand the beam. The global 40 Ca⁺ and 88 Sr⁺ beams are overlapped and reach the trap from the north-east viewport, aligned with the trap's axial axis. The overlapped addressing beams approach through the south-east viewport, and pass through the high-NA lens, which ideally focuses the beam down to a sub-micrometer waist size.

Ionic fluorescence is collected by the in-vacuum lenses, and can be observed from both the front and back of the slotted trap. Front-side imaging with a CCD camera allows us to see the surface

²¹ SIM960, SRS

²² Which coincidentally matches real-world cardinal directions quite closely




Figure 3.8: Simplified schematic of the optical setup on the experiment table, with beam lines shown in red, and paths for fluorescence detection in blue.

of the trap and how various beams reflect off of it. Imaging from the front is therefore used for aiding in beam alignment²³. Fluorescence collected through the backside of the trap is used for state detection. Fluorescence from ⁴⁰Ca⁺and ⁸⁸Sr⁺is separated into two paths with a dichroic mirror. ⁴⁰Ca⁺fluorescence is either sent to a PMT or a CCD camera by means of a mirror on a flip-mount. A second PMT detects strontium fluorescence.

3.1.3 Experimental control

A schematic overview of the main hardware and software components that provide experimental control²⁴ is shown in Figure 3.9. The experiment is centrally controlled from a PC, running Trapped Ion Conrol Software (TrICS), a home-built program that provides connectivity to the various components of the experiment and a user-interface for controlling them. A thorough overview of TrICS can be found in the thesis of Daniel Heinrich [132].

²³ Note 12 in Appendix e

²⁴ Note 13 in Appendix e




Figure 3.9: Schematic overview of the experimental control hardware

Most experimental parameters (AOM settings, trap drive power, etc.) require control that is not sensitive to precise timing, as they do not change within experimental sequences. In contrast, a few dedicated pieces of hardware are responsible for controlling parameters that must be updated with sub-microsecond timing during experimental sequences. We will first discuss the *quasi-static* experiment control, that remains unchanged during an experimental cycle.

The control PC has an I/O PCI card²⁵, which is controlled by TrICS. The card has a set of digital outputs, which are used to communicate with devices in the lab. Digital signals include 8 address bits, to isolate which device is being communicated with, and 16 data bits that specify device settings. Connectivity to these devices is provided by a 50-wire ribbon cable²⁶, denoted as *bus system* in Figure 3.9. The ribbon cable is separated into sections by repeater and splitter stations, which galvanically isolate separate parts of the lab and refresh digital signals. Ribbon cables are attached to the backplane of 19" racks, which house the devices that need to be controlled. The bus system supports communication with Direct Digital Synthesisers (DDSs), DACs, and digital in and output cards. In practice, almost without exception²⁷, the only devices connected to the bus system are DDSs, which provide the RF signals for AOMs, and the trap drive. A lone DAC card drives a pressure controller, which controls the cryostat's liquid flow rate, and thus its temperature.

The aforementioned control of AOMs allows us to set frequencies and powers of lasers used in experimental sequences. However, the bus system does not allow for precisely timed switching of these laser parameters. Such switching within experimental sequences is achieved by a dedicated piece of hardware, the *pulse box* [133]. The pulse box has TTL outputs and RF outputs, the latter of which can be phase coherently switched in amplitude, frequency and phase. Pulses can be generated with a time resolution of 10 ns. Pulse timings are controlled by a built-in FPGA. Pulse sequences, designed within the TrICS framework, are communicated to the pulse box's FPGA through a dedicated sequence generation and communication software [134] on the main control PC.

The RF pulses emitted by the pulse box are externally amplified and subsequently drive 4 specific AOMs, namely the ones in double-pass configuration in the top right of Figure 3.8. These AOMs are responsible for coherent switching of the qubit interaction beam, at 729 nm for 40 Ca $^{+}$ and 674 nm for 88 Sr $^{+}$. In addition to simply switching these beams on and off, the frequency, phase, and amplitude of these fields are regulated by the pulse box.

The TTL pulses emitted by the pulse box are connected to various devices: TTL signals enable or interrupt the passage of RF signals going to the AOMs in double-pass configuration, thus switching the beams on or off. TTL pulses are also used as a start and stop signal for a counter card²⁸ which detects PMT pulses. This allows us to register fluorescence rate within a specified detection window. Count results are communicated back to the control PC through a dedicated PCI card.

²⁵ NI-DAQmx-PCI-6534

²⁶ Migration to an ethernet-based communication network has begun in our lab, but most of our devices still run on the ribbon-wire bus

²⁷ In the past, the bus system also controlled DACs, which provided a voltage to piezos attached to cavity mirrors, allowing us to tune the frequencies of lasers locked to these cavities. The cavity locks have since been replaced by a wavemeter lock.

²⁸ NI-DAQmx-PCI-6733




Figure 3.10: Microscope image and electrode layout of the Berkeley trap. Dimensions are in micrometers.

The control PC has stand-alone software to handle the voltage output of the AWG, which provides voltages to the trap's DC electrodes (see Section 3.1.1.3). This software allows us to set static voltages or, alternatively, supply the AWG with pre-computed voltage sequences. The AWG has a trigger-input, allowing us to start these voltage sequences on-demand within experimental cycles using the pulse box's TTL outputs. Similarly, separate software provides control of the CCD camera, but can be programmed to detect fluorescence within experimental cycles, initiated by TTL signals from the pulse box.

3.2 'GOLDEN GATE' SURFACE ION TRAP

This section provides an overview of the surface trap installed in the experimental setup. As shown in the beginning of this chapter, in Figure 3.1, the experiment has seen the comings and goings of multiple different traps. This section focuses on the most recent of these, the 'Golden Gate' surface ion trap. Before discussing that trap though, two previous traps used in the course of my Ph.D. work deserve some attention as they have sourced some of the results presented in this thesis. A brief description of those traps is given in the following sections.

3.2.1 Previous traps

3.2.1.1 'Berkeley' trap

The 'Berkeley' trap, so named because it was designed, and processed at UC Berkeley in the group of Hartmut Häffner, was installed in our setup as a platform to demonstrate the prospects of scalability in quantum computation (as opposed to the previously installed traps, which merely demonstrated that the apparatus could trap in cryogenic operation). The layout of the trap is shown in Figure 3.10, with a microscope image on the left, and a schematic layout on the right. The trap is based on a gold-coated, micro-structured silica substrate²⁹. Careful inspection of the trap image reveals many regions of dark spots, most pronounced on the mid-left. These are artifacts due to stress on the substrate during the microstructuring procedure. They are below the surface, and do not affect the surface quality. More about the trap's fabrication, coating, and wire-bonding, which is nearly identical to the 'Golden Gate' trap, is covered in Section 3.2.3³⁰.

The trap consisted of 10 pairs of DC electrodes, each 200 μ m wide. The traces leading up to the electrode are narrower than the electrodes to minimize capacitance between them, thus minimizing electrical crosstalk. The dark area in the center of the trap in Figure 3.10 is the slot region, which is 100 μ m wide. Flanking the trap slot are the central DC electrodes, which are 30 μ m wide. These electrodes join up on the right side, thus share the same voltage. One strip further out are the RF electrodes, which provide the trap's RF confinement. The RF electrodes produce a field, whose potential

²⁹ Fabricated by *Translume*

³⁰ Note 14 in Appendix e




Figure 3.11: Microscope image of the Sandia HOA trap surface [109], and schematic layouts corresponding to the purple and blue dashed regions in the image. Dimensions are in micrometers.

saddle-point, which corresponds to the pseudopotential minimum, is 110 µm above the surface of the trap. All electrodes are separated by trenches, which are 10 µm wide and 30 µm deep.

The Berkeley trap was the first trap in our cryogenic apparatus that demonstrated coherent qubit operations, on both $^{40}\text{Ca}^+$ and $^{88}\text{Sr}^+$. The system was characterized in terms of coherence times and trap heating rates. In fact, the trap boasted a respectably low axial heating rate of 2 phonons per second at 1 MHz trap frequency. Entangling gates (MS gates, see Section 3.2.8) were performed on $^{40}\text{Ca}^+$, and manipulation of the ion position was demonstrated by precise control of the trap voltages. Specifically, we did an in-depth study of ion crystal rotations, in which a sequence of voltages causes two ionic qubits to switch positions, while maintaining their quantum coherence. This operation is the focus of Chapter 5.

3.2.1.2 Sandia High Optical Access (HOA) trap

The Sandia HOA trap was a successor to the Berkeley trap, and was installed along with a new-and-improved version of the cryostat. The trap was designed, manufactured, and packaged by Sandia National Laboratories, as part of IARPA's Multi-Qubit Coherent Operations (MQCO) collaboration. The design goals include [109]

- Providing high optical access (hence the name) for laser beams and detection optics, on various axes
- Having a high degree of control of local trapping potentials, such that multiple neighboring trapping sites can be independently controlled
- The ability to split ion chains at electrode voltage that do not exceed the voltage supply limit or electrode voltage breakdown limit
- Having a junction region to allow re-ordering of ion chains

An scanning electron microscope (SEM) image of the trap surface is shown in Figure 3.11. The electrode layout of the dashed areas are shown alongside. The 'quantum region' of the trap, denoted by the purple dashed box, is slotted to provide optical access for beams perpendicular to the trap plane, thus allowing single-ion addressing. The slot is $60\,\mu m$ wide. The electrode layout is conceptually similar to the Berkeley design, with the notable difference that the segmented DC electrodes flank the slotted region and are enclosed by the RF rails. The segmented DC electrodes are therefore relatively closer to the ions, and provide finer control over local potentials. The RF rails produce a potential whose saddle-point is $70\,\mu m$ above the surface of the trap.

Two junction regions at each end of the linear central region separate the trap into four arms. The junction regions do not have a slot. The jagged electrode edges in the slit-to-no-slit transition and in the junction minimize deviations in the curvature of the RF pseudopotential experienced by an ion traversing these regions.




Figure 3.12: (a) Image of the Golden Gate trap. (b) Portion of the electrode layout (c) Cross section, cutting vertically through (a). Dimensions are in micrometers

The trap is wire-bonded onto a chip carrier, with a standardized 100-pin PGA footprint (see Section 3.1.1.3). This makes it easy to swap out traps 31 . In our experiment we are no strangers to trap replacement: Several versions of the Sandia HOA trap have been tested, with varying results, though unfortunately mostly discouraging ones. Overall, the heating rates (see Section 2.3.4) were on the order of several thousand phonons per second at typical trap parameters (at an axial common mode frequency of $\sim 1\,\mathrm{MHz}$). This far-exceeds desirable heating rates for high-fidelity quantum operations. The mechanism that leads to high heating rates (which apparently is more prominent in cryogenic operation) is as of yet not fully understood but is believed to be related to incomplete metal coverage of the slit.

An additional technical difficulty was that the slotted region was too narrow for our 397 nm Doppler cooling beam to pass through without notable scattering from the slot's sidewalls. This made it difficult to have an acceptable signal-to-noise ratio for state detection. A final technical difficulty was that we never managed to co-trap two different ion species.

Our way forward was to return to the previously shown-to-work Berkeley trap, though with minor modifications, in line with the design goals of the Sandia HOA trap. This successor trap, known as the 'Golden Gate' trap, was used for most of the work described in this thesis and is discussed in the following section.




Figure 3.13: (a) Contour of the pseudopotential for 40 Ca $^{+}$ at an RF electrode voltage amplitude of 200 V, and a 40 MHz trap drive. The potential minimum is marked with a red cross. (b) Pseudopotential along the radial direction perpendicular to the trap surface. At 200 V, the curvature around the potential minimum, denoted by the dashed line, is 190 V mm $^{-2}$, corresponding to a 3.5 MHz radial motional frequency. The trap depth is 90 meV.

3.2.2 Golden Gate trap overview

The credit of naming the 'Golden Gate' trap goes to Pavel Hrmo, and was so named for the following reasons: 1) The traces leading up to the DC electrodes are reminiscent of the arches of the Golden Gate Bridge³², 2) the trap is coated in *gold*, 3) the trap is used to do quantum *gates*, and 4) it's an homage to the group in Berkeley where we borrowed the design from, with a view of the Golden Gate bridge from the lab. An image of the trap and its electrode layout are shown in Figure 3.12. The trap chip dimensions are $9 \times 4.5 \times 0.5$ mm³.

The trap has 27 pairs of DC electrodes. A central DC electrode is made up of two co-wired strips (30 μ m) that flank a central slot. The central DC electrodes are enclosed within a pair of RF rails (60 μ m). The DC electrodes are narrower than the Berkeley trap, (100 μ m pitch), providing finer control of trap potentials around the ion. Trenches between the electrodes are 10 μ m wide and are designed to be 50 μ m deep, though actual manufactured depth has some tolerance and might, according to the manufacturer³³, more realistically be around 40 μ m. The slotted region is 100 μ m wide, 70 μ m deep, and the opening angle behind the slot towards the back side of the trap is 80°. These values were chosen to reduce light scattering of the 397 nm beam, both on the front side and back side, while retaining structural integrity of the chip. The 397 nm beam is used for fluorescence detection, which requires a high signal-to-noise ratio for quick state detection.

Figure 3.13(a) shows a contour plot of the RF pseudopotential in the radial direction (see Eq. 2.12. The potential is shown for 40 Ca $^+$, with a 40 MHz trap drive frequency and an RF voltage amplitude of 200 V. These are typical parameters in our experiment, and result in radial motional frequencies of 3.5 MHz. The RF null, the location where ions are intended to be trapped, is 110 μ m above the surface of the trap. The pseudopotential in the radial direction perpendicular to the trap surface is shown in Figure 3.13(b), from which it can be seen that the trap depth, the energy required for an ion to escape, is 90 meV. With the same trap voltages, a 88 Sr $^+$ ion needs 41 meV to escape the trap.

3.2.3 Manufacturing

The trap chip is based on a fused-silica substrate. The electrode structure and slot are fabricated with laser micro-machining. The surface was gold-coated in our in-house cleanroom facility in the

³¹ The process of switching out one trap on the trap PCB with another takes a matter of minutes. Though that is ignoring the amount of time it takes to disassemble and the setup to reach the trap in the first place, and to reassemble it afterwards.

³² If you squint

³³ Translume

University of Innsbruck using evaporation deposition. During evaporation, the trap surface is tilted at a 45° angle with respect to the evaporation flux. First, a 10 nm layer of Titanium is applied as an adhesion promoter, and secondly a 150 nm layer of gold is applied. This process was repeated four times, twice on the top surface and twice on the bottom surface. Between the two top-side evaporation runs, the trap chip is rotated 180° around the surface normal, though keeping the 45° with respect to the evaporation flux. This ensures that all trench walls are coated, but keeps the bottom of the trenches uncoated, thus avoiding electrical contact between neighboring electrodes³⁴.

The coated trap is glued³⁵ to a trap carrier PCB. Electrode pads are wire-bonded to the traces on the carrier PCB. The RF electrode had multiple wire-bonds This PCB has the same footprint as the Sandia HOA trap package, and could therefore be installed in our setup without any further modifications to the apparatus.

Shortly after installation, ⁴⁰Ca⁺ and ⁸⁸Sr⁺ ions were stably trapped in the Golden Gate trap. In the following sections, trapped ions are used to characterize the Golden Gate trap.

3.2.4 Potential calibration

Following the methods described in Section 2.3.1, a set of voltages are derived that control trapping potentials, expanded in a basis of spherical harmonic multipoles. Tolerances in the fabrication of the trap, limits in accuracy of potential simulations, and undesired stray charges influence how well these sets of voltages produce the actual desired potentials. We use the methods of section 2.3.2 to calibrate any possible discrepancy between set potentials and measured potentials, as given by Eq. 2.77.

The measured motional frequencies behave as expected according to the set potential: the multipoles that intentionally influence motional frequencies lead to frequency changes that range within 5% of the predicted value. Multipoles that should leave motional frequencies unaffected for the most part do so. The exception is that the $Y_{1,1}$ term (see Eq. 2.72), which ideally should only produce a homogenous field in the radial direction parallel to the trap surface, additionally induces splitting of the two radial mode frequencies. The optimized calibration allows us to predict motional frequencies with less than 1% error.

3.2.5 Micromotion compensation

In this section, the techniques described in Section 2.3.3 are applied to minimize and characterize the motion of an ion due to the ion trap's RF field, micromotion. We use the methods of photon correlation, CCD detection, and micromotion sideband spectroscopy. *Photon correlation* is the detection of the arrival time of an ion's fluorescence photons with respect to the phase of the RF drive. It can detect a component of micromotion parallel to the Doppler cooling beam. *CCD detection* uses fluorescence imaging to detect a shift in position of the ion when the trap drive power is changed. It detects micromotion in the image plane. These two methods together are a quick method of compensation and are used on a daily basis. Since the Doppler cooling beam in our setup does not lie in the image plane, the two techniques are sensitive to micromotion along different axes, and are therefore suitable together for compensation in both radial directions.

Figure 3.14(a) shows results of micromotion compensation using photon correlation. Histograms of photon counts as function of detection time (in terms of the phase of the RF drive) are shown for two applied electric fields, at $0.75\,\mathrm{V\,mm^{-1}}$ and $0.85\,\mathrm{V\,mm^{-1}}$. These fields are produced perpendicular to the trap surface, by applying a set of voltages to the trap's electrodes, as described in Section 2.3.1. In the presence of excess micromotion, the fluorescence rate is measurably correlated with RF phase, as seen in the left plot, emphasized by a sinusoidal fit. Correcting micromotion flattens this curve, as seen in the right plot. We extract the amplitude from the sinusoidal fit for similar measurements over a range of applied fields, and show the results in the lower plot. The amplitude has a minimum around $0.85\,\mathrm{V\,mm^{-1}}$.

Complementary to this method, we monitor the position of the ion using a CCD camera, whose image plane is parallel to the trap surface. Applying a field parallel to the trap plane, in the radial

³⁴ Note 15 in Appendix 🤇

³⁵ EPO-TEK H20E




Figure 3.14: Results of micromotion compensation on the Golden Gate trap, using various techniques. (a) Photon correlation: upper plots are histograms of ion fluorescence counts as function of RF phase, shown for two field strengths perpendicular to the trap plane. The blue line is a sinusoidal fit. The lower plot shows the amplitude of the fitted lines, for various field strengths. (b) Ion position as function of field strength (parallel to the trap surface, in the radial direction) for two RF powers. (c) Excitation by applying a pulse on a carrier transition (upper) and micromotion sideband (lower). The ratio of their Rabi frequencies is used to calculate the micromotion index.

direction, displaces the ion. The amount of displacement depends on the power of the trap's RF drive. Figure 3.14(b) shows the ion displacement as function of the applied field, for two RF powers (where the 'high' power corresponds to a radial secular motional frequency of 3.5 MHz and the 'low' power to 1.8 MHz). The ion position linearly changes with applied field, as can be seen by the fit lines. The value of applied field where the two lines cross corresponds to a minimum in micromotion. Combining the methods shown in Figure 3.14(a) and (b) allows us to compensate micromotion in both radial directions to an accuracy of about $\pm 10 \,\mathrm{V}\,\mathrm{m}^{-1}$.

A more accurate fine-tuning of micromotion is done by probing a micromotion sideband of one of the ion's $4S_{1/2} \leftrightarrow 3D_{5/2}$ transitions. The ratio of coupling strengths of the micromotion sideband and the carrier transition is used to determine the micromotion modulation index (Eq. 2.85). Figure 3.14(c) shows excitation of a carrier transition (top) and its motional sideband (bottom). The lines are sinusoidal fits, of which the respective Rabi frequencies are determined to be $\Omega_{\rm car} = 51.3(3)$ kHz and $\Omega_{\rm MM} = 1.0(1)$ kHz. The modulation index is then $\beta = 0.037(6)$, corresponding to a micromotion oscillation amplitude of about $r_{\rm MM} \approx 4.3$ nm. This value indicates the amplitude of the component of micromotion in the direction parallel to the beam used to probe these transitions, which in our case is perpendicular to the trap plane.

In addition to the measurements shown in Figure 3.14(c) the micromotion modulation index is measured using a beam that is parallel to the trap's axial direction. Ideally, in a linear trap, ions should experience no micromotion in this direction, thus should not have a detectable micromotion sideband. In our setup, we measure a modulation index of 0.25 in the axial direction. This value is independent of which DC fields are applied. This suggests there is a source that produces a homogeneous RF field (as opposed to the harmonic RF field produced by the trap's RF electrodes). We suspect that the source of this field is a wire a few centimeters away from the trapping region that connects the resonator coil to the trap PCB. Electric-field simulations indicate that, even with shielding provided by the trap surface and the trap clamp, this wire can indeed produce a high enough field to account for the measured axial micromotion modulation index. Note also that this axial micromotion can offset the result of the photon correlation method for micromotion compensation. Since the Doppler cooling beam wavevector has overlap with the axial direction, a minimization of correlation amplitude is not necessarily reflected by the ion being at the RF null. Despite this, the photon correlation technique is still useful as a coarse optimization, after which probing the micromotion sideband can be used for fine-tuning.

For single species operation, this additional micromotion is most likely not a hindrance since the added motion is the same for all ions. In mixed-species operation, however, the RF field unequally




Figure 3.15: Heating rate measurements of the Golden Gate trap.

modulates the ions' axial position, which leads to anharmonic secular motion. This, in turn, can lead to coupling between motional modes, mode heating, and motional decoherence. Future versions of the setup will have this design flaw corrected.

3.2.6 Heating rates

Heating of an ion or group of ions' secular modes, as discussed in Section 2.3.4, lead to a reduction in fidelity of many quantum operations. In order to be able to characterize errors in ion trap experiments, the trap's heating rates must be measured. In this section, heating rate measurements of the Golden Gate trap are discussed.

Figure 3.15 displays heating rate results of the Golden Gate trap. 'Typical' trapping parameters in our setup correspond to an axial common mode frequency of 1.06 MHz and radial frequencies of 3.3 MHz and 3.6 MHz parallel and perpendicular to the trap plane. Most results in this thesis use similar motional frequencies, the exception being the results in Chapter 6, in which experiments are performed at an axial common mode frequency of 0.55 MHz. Figure 3.15 shows heating rates \dot{n} of the axial mode at various mode frequencies ω , measured using a single 40 Ca $^{+}$ ion, using the sideband thermometry technique described in Section 2.3.4. A heating rate of approximately 28 phonons per second is measured at 1.06 MHz, and around 100 phonons per second at 0.55 MHz.

A commonly used model for mode heating in surface traps poses that the heating rate is inversely proportional to the mode frequency ω and proportional to the spectral density of the electric field noise $S(\omega)$. The measured heating rate as function of mode frequency can therefore be used to glean information about the electric field noise. For example, many electrical systems have noise spectral densities that follow a $S \propto 1/\omega$ frequency scaling [93], which would be reflected by a ω^{-2} scaling in the heating rate. The dashed line shows a least-squares fit through the data using a power law model, $\dot{n} \propto \omega^{-\alpha}$, from which a frequency scaling factor of $\alpha = 2.0(3)$ is determined. This scaling is consistent with basic surface noise models [29].

The measurements shown in Figure 3.15 were taken in May 2022, at which point the trap was in operation for about three-and-a-half years. In early 2021 and earlier, heating rates were measured on the order of 20 phonons per second. The increase in heating rate over time may be attributed to an increased amount of contaminants on the surface of the trap [36, 135]. Nevertheless, the heating rate of the Golden Gate trap is in an acceptable range for many demonstrative quantum computation sequences. We have also measured heating rates of motional modes in the radial directions, at frequencies between 3 and 4 MHz, and find rates between 50 and 100 phonons per second. The works in this thesis make use of the axial motional modes, so radial heating is less consequential. The measured radial heating rates are therefore also within an acceptable range.

3.2.7 *Coherence times*

A defining property of a qubit is its coherence, which is effectively a measure of how well-defined the phase-relation between the two qubit states is. Imprecision in experimental properties like magnetic field strength and laser frequency cause this phase-relation, thus the coherence, to decay over time.




Figure 3.16: Coherence time measurements on the Golden Gate trap, inferred from Ramsey contrast. Coherence times are shown between various states, including (a)-(b) two optical transitions, (c) an RF transition, and (d) a motional mode transition. Light blue lines are exponential fits, from which indicated T_2 times are derived. The teal lines are a guide to the eye linking smoothed contrast data in the optical qubit.

In this section, we display results of coherence time measurements, using the Ramsey measurement technique covered in Section 2.3.5.

Firstly, a 'standard' Ramsey measurement, as in Eq. 2.96, is executed. Unfortunately, our system suffers from slow magnetic field drifts; individual measurements (i.e. 50 repetitions of an experimental cycle) can provide results with high contrast, suggesting that during this set of cycles the magnetic field is stable. However, a full Ramsey scan is typically done with a sequence of such measurements, over a range of Ramsey phases, to get a full Ramsey fringe as in Figure 2.11. The magnetic field drifts in our setup are large enough that regardless of which levels are used as qubits, a Ramsey time $t_{\rm wait}$ of about 4 ms results in a measured excitation as function of Ramsey phase with high contrast, but no clear sinusiodal behavior. Since such measurements are difficult to analyze unambiguously, we do not display these results. However, we can conclude that we maintain a near unity coherence at around 1 ms, but cannot draw definitive conclusions about coherence from these measurements after a few milliseconds.

Instead, we add an echo pulse to the Ramsey sequences, also described in Section 2.3.5. For these sequences, one does not need to scan the laser phase of the final Ramsey pulse, because linear phase drifts due to incorrect laser frequency settings are canceled out by the spin-echo. Coherence can therefore be extracted from the excitation contrast for Ramsey sequences with the laser phases $\phi=0$ and $\phi=\pi$. Figure 3.16 displays results of Ramsey measurements performed on $^{40}\text{Ca}^+$, with spin-echo, where we probe the coherence of various levels: From left to right, we probe two optical transitions, $(4S_{1/2}(-1/2) \leftrightarrow 3D_{5/2}(-5/2))$ and $4S_{1/2}(-1/2) \leftrightarrow 3D_{5/2}(-1/2)$, the ground state $(4S_{1/2}(-1/2) \leftrightarrow 4S_{1/2}(+1/2))$ and check the coherence of the ion's motion $(|n=0\rangle_n \leftrightarrow |n=1\rangle_n)$. In the plots in Figure 3.16, the blue dots are measured contrast and the blue lines are exponential fits through the data³⁶, from which the coherence time T_2 is extracted, using the assumption that the measured contrast follows $\exp(-t_{\text{wait}}/T_2)$. The large statistical spread in the measured data is mostly due to quantum projection noise, as each data point was obtained with 50 repetitions of an experimental cycle (which in hindsight should have been measured more thoroughly).

The T_2 times of both optical qubits and the ground state qubit all exceed 100 ms. The optical qubits both exhibit a feature of coherence recovery at around 50 ms, after an initial drop around 25 ms, which can be seen with the teal guide-to-the-eye that connects smoothed measured data in Figure 3.16. Since

³⁶ Different noise models lead to different decay functions in Ramsey contrast. Our data likely suffers from various types of noise, and cannot be described by a simple model. The exponential decay model provides a good indication of coherence time but is not intended as an exact representation of the decoherence mechanism.




Figure 3.17: Experimental realization of a two-ion Mølmer-Sørensen gate on the Golden Gate trap. (a) State population as function of the MS-gate pulse duration. The Bell state $\frac{1}{\sqrt{2}}(|00\rangle+i|11\rangle)$ is generated at a pulse duration indicated by the purple bar, confirmed by (b) a parity measurement, in which an analysis pulse with varying phase prepares even or odd states.

the ground state qubit does not contain this feature, this points to a source of coherent noise in the laser frequency, likely related to the 50 Hz electrical line cycle.

The optical qubit with upper state $3D_{5/2}(-5/2)$ is five times more sensitive to magnetic field fluctuations than $3D_{5/2}(-1/2)$ (see Eq. 2.66). If magnetic field noise was the dominant source of noise, then the T_2 times of the two optical qubits would differ by a factor of five. Since this is not the case (we measure a factor of about 1.2), the coherence time seems to be limited by laser phase noise. One would then expect a significant improvement when probing the coherence time of the ground state qubit, which shares the same magnetic field sensitivity as the $4S_{1/2}(-1/2) \leftrightarrow 3D_{5/2}(-5/2)$ transition, but is insensitive to laser phase noise. However, the measured T_2 time of the ground state qubit, 116(8) ms does not reflect this expectation. We suspect, instead, that motional heating affects the outcome of the ground state qubit coherence measurement: The Ramsey sequence for the ground state qubit, including the spin-echo, requires a total of seven pulses, compared to the three required for the optical qubit. Motional heating during the wait time reduces the success of full population transfer that these pulses are expected to make, which makes the outcome of sequences with more pulses more sensitive to mode heating.

Finally, the coherence time of an ion's motional mode (using the axial mode, at $1.06\,\text{MHz}$) is measured to be 17 ms. The two main sources that lead to motional decoherence are motional heating and noise in the motional frequency due to undesired changes in the trap's confining potential. An ion at an axial frequency of $1.06\,\text{MHz}$ has a heating rate of 28 phonons per second, which corresponds to a theoretical motional T_2 time of 36 ms. The measured T_2 time is thus dominated by instabilities in the confining potential, likely caused by voltage fluctuations on the trap surface.

The duration of many of the sequences in the works detailed in this thesis are on the order of a millisecond or less. Decoherence, even when no spin-echo sequences are used, represent only a minor source of error in such sequences.

3.2.8 Gate performance

We characterize the setup's ability to generate qubit entanglement by applying the Mølmer-Sørensen (MS) gate, described in Section 2.3.6. Figure 3.17(a) shows an experimental realization of an MS-gate pulse applied to two trapped ions in the Golden Gate trap. Individual data points are the measured population of the states $|00\rangle$, $|11\rangle$ and the combined population of the states $|01\rangle$ and $|10\rangle$. The lines are simulated results, showing the excitation development of an ideal gate pulse.

At a pulse duration of $t=145\,\mu s$, indicated in Figure 3.17(a) with the purple bar, the applied pulse produces the desired populations of an entangled state, $\frac{1}{\sqrt{2}}(|00\rangle+i|11\rangle)$. A gate fidelity of 97.3(5)% is obtained from gate population and parity scans. An example of a parity scan is shown in Figure 3.17(b). A portion of the error in gate fidelity comes from errors in state preparation and measurement

(SPAM). The gate error of the MS-gate alone, without SPAM, can be deduced by measuring gate fidelity for multiple repeated applications of the gate, and determining the overall per-gate drop in fidelity. Only fidelities of odd numbers of gates are accounted for, since even numbers undo the entangling operation. We have repeated the gate up to 9 times, and determine a gate fidelity of 98.6(1)%. The main source of error is resonant excitation of the carrier due to the 729 nm diode laser's non-negligible noise floor. Depending on the intended gate application, more robust techniques of analyzing gate performance may be desirable [136].

4

CLASSICAL ION DYNAMICS IN AN RF PAUL TRAP

Assume that a penguin is a circular cylinder...

- Fundamentals of Physics, Halliday et al.

In this chapter, we will discuss classical motion of ions trapped in a Paul trap. Here, the term "classical" is used to emphasize that we take into consideration the Newtonian mechanics of charged particles, whose energies are high enough such that quantum mechanics can be neglected. We investigate classical motion by executing numerical simulations of ions in time-dependent fields. This investigation leads us to a discussion of an often overlooked —but in some cases dominant— effect known as RF heating.

4.1 MOTIVATION

Our motivation for studying ion motion stems from experimental difficulties in utilizing surface traps for scalable quantum computation: Surface traps intrinsically have a low energy barrier for ions to escape, compared with macroscopic 3D traps and are therefore more susceptible to ion loss. Ion loss should be minimized, as such events lead to tedious ion reloading times, and constitute a decrease in computational speed.

On the surface¹, it would seem that keeping ions trapped should not pose any difficulty: Ions are Doppler cooled to energies in the μeV range, whereas trap depths of typical surface traps are typically tens or hundreds of meV. In the absence of Doppler cooling, an ion might gain energy at rates on the order of $1 \mu eV s^{-1}$, due to electric field noise emanating from the trap surface (see Section 3.2.6). At this rate, one would thus expect uncooled lifetimes of at least tens of minutes. During experimental sequences, Doppler cooling is never off for more than a few tens of milliseconds. Anomalous heating alone therefore would not likely lead to ion expulsion. A more dramatic change of energy can occur if a particle in the background gas collides with an ion. However, some back-of-the-envelope ballistic collision calculations suggest that it's nearly impossible² to have a collision that leads to an ion energy above 40 meV. From this we can naively conclude that as long as Doppler cooling isn't interrupted for too long of a period of time, ion loss should not be a problem, even for surface traps.

Unfortunately, ion loss is a persistent problem for many ion traps with low escape barriers, and the cause is not fully understood. Throughout the many iterations of traps in our setup, ion lifetimes have varied from days to minutes to seconds, without a clear indication of which experimental parameters lead to these changes in lifetime. Ion lifetimes (and loading times) in mixed-species operation was particularly perplexing: in one trap (see Section 3.2.1.2) we were able to stably trap a chain of 40 Ca $^+$ for several hours, stably trap a chain of 88 Sr $^+$ for several hours, but never managed to keep a mixture of Ca and Sr in the trap, let alone crystallize it. If a trap can hold multiple specimens of either species, what process is responsible for hindering loading of a mixed-species crystal?

There is much to be learned about the processes involved in ion loss and loading. This prompted us to study the dynamics of trapped ions in simulation. Such simulations enable the investigation of many aspects of trapped ion stability, for example the effects of asymmetry in the trapping potential, parametric resonances, varying ion numbers, and mixed species interaction. Our classical ion dynamics investigations have helped us determine, understand, and mitigate the main source of heating that limits mixed-species loading rates and trapping lifetimes.

The treatment of ion dynamics is usually massively simplified by assuming that an ion's motion ignores the trap's time-dependent RF field and instead is governed by an effective static potential. We

¹ Pun intended

² Granted, my envelope uses a background gas of only hydrogen, with an energy distribution at cryogenic temperatures (30 K). Still, this statement easily holds for room temperature setups

begin this chapter with a digression, giving word of warning that this simplification is not always justified. The remainder of this chapter highlights this notion by investigating one example of this, a phenomenon known as RF heating. In this chapter, RF heating is studied through numerical ion trajectory simulations, is analytically modeled, and experimentally observed.

4.2 BEYOND THE PSEUDOPOTENTIAL APPROXIMATION

It is ubiquitous among works on quantum computation with trapped ions [137] to introduce the core concepts of ion trapping by introducing the *pseudopotential* — a time-averaged approximation of the driving RF electric field potential of a Paul trap. This pseudopotential, often assumed to be harmonic, is a useful approximation as it can be used to describe the motion of an ion without needing to take into account the infinite terms of the solution of the Mathieu equations (see Section 2.1.1). The unperturbed motion of an ion within a harmonic pseudopotential is sinusoidal in time and contains no frequency component at the driving RF frequency (known as *micromotion*). This sinusoidal motion is known as *secular* motion.

The pseudopotential approximation is an immensely useful simplification of the electric field potential, since it allows us to easily describe equilibrium positions and motional modes of ions. This in turn, hugely simplifies the treatment of atom-light interaction, which constitutes the foundation of trapped-ion quantum computation. The interactions discussed Section 2.2.6, indeed, assume ions are trapped in a static confining potential. It is therefore a lucky break that one needn't worry about any ion motion with frequencies higher than the secular frequencies.

However, the influence of the RF field cannot always be neglected. The simulations used in this section cover regimes where the dynamics of ions do not obey the pseudopotential approximation. We note some phenomena that lead to a deviation from this approximation:

-Doppler cooling-

The interaction of an ion and a Doppler cooling beam is fundamentally dependent on the ion's velocity (see Eq. 4.24). In typical ion trap operation the amplitude of an ion's periodic motion is dominated by its secular motion³, which makes it tempting to dismiss the effect of micromotion. However, an ion's velocity amplitude due to the RF drive is similar to its secular velocity amplitude. Figure 4.1(a) demonstrates the notable difference between the velocity distribution of an ion if simulated in a pseudopotential versus a time-dependent potential generated by an RF drive. Since the functionality of Doppler cooling depends on velocity, it is evident that micromotion cannot be neglected if one includes Doppler cooling in simulations. In fact Blümel et al. [138] note that with a sufficient amount of micromotion, a Doppler cooling beam can *add* energy to ions.

-Coulomb interaction-

A second reason that the pseudopotential approximation might fail to adequately describe ion motion is an interaction of forces due to the Coulomb potential between multiple ions and the time-dependent RF potential. When ions are far from the motional ground state (several hundreds of meV) this interaction leads to an aperiodic energy exchange between the driving RF field and the ions. This energy exchange, termed RF heating, is the focus of the majority of this Chapter and will be further defined, described, modeled and measured in Section 4.3. For now, we ask for the reader's patience for our vagueness in the usage of terms such as "RF heating" and "energy" in the following paragraph.

It is worth mentioning that it is not obvious that the RF potential should induce energy changes, even for a case as simple as two trapped ions and is therefore often overlooked. It has previously been recognized [138] that RF heating can play a dominant role in energy exchange in few-ion systems, though the rate at which RF heating occurs depends on the ion's energy. This heating rate, as a function of ion energy, is shown in Figure 4.1(b). The data in this plot are obtained using ion trajectory simulations, detailed in the following section. For comparison, using a pseudopotential in the simulation does not show any heating, demonstrating that, indeed, the pseudopotential fails to encompass a dominant energy exchange mechanism. Despite this, the dependence of ions' motion on

³ For typical trapping parameters, the ratio of amplitudes of micromotion a_{RF} to secular motion a_{sec} is closely related to the ratio of the secular frequency ω to the trap drive frequency Ω_{RF} and is also similar to the Mathieu q-parameter: $a_{mm}/a_{sec} \approx \sqrt{8}\omega/\Omega_{RF} \approx q/2$, see Section 2.1.1.

the RF potential for systems with low numbers of particles is often assumed to be negligible [65, 130]. This discrepancy suggests that a more solid foundation on ion dynamics within an RF potential is required.




Figure 4.1: Examples where the pseudopotential does not accurately portray ion dynamics. (a) The velocity distribution and (b) the energy change rate of ions simulated in a pseudopotential versus ions in a time-dependent RF drive. Velocities in (a) are scaled with respect to the amplitude in velocity of the ion's secular motion. The simulation parameters, definitions, and results in (b) are similar to those in Ref [138].

4.3 RF HEATING OF NON-CRYSTALLIZED TRAPPED IONS

In this section, we focus our attention onto the investigation of RF heating. The content of this section is adapted from the publication **RF-induced heating dynamics of non-crystallized trapped ions** [139]. We start with an introduction where terms used in the previous section finally get some well deserved definition and motivational context.

4.3.1 RF heating: an overview

In trapped-ion based quantum computation, qubits are assumed to be well-localized, separated, trapped ions that share common motional modes due to their Coulomb interaction. A collection of such ions represents a register of qubits, and is characterized by a regular and predictable spatial structure, known as a Coulomb *crystal*⁴. Quantum computation demands that ions remain in their crystal state.

A prominent event that disturbs the crystal structure is a collision with a particle from the residual background gas [140]. Such a collision can transfer enough energy to ions such that the crystal structure is destroyed. The ions undergo a transition described as *melting*, to a phase colloquially named an ion *cloud* [138, 141]. The defining feature of a cloud is that ions do not have unique average positions, unlike in a crystal. By this definition, as few as two ions can constitute an ion cloud, despite that the term cloud often carries the connotation of involving an ensemble of ions. Most of the discussion in this section involves a two-ion cloud.

Melted ions are subjected to a change in energy that is not present in the crystal phase: Energy can be transferred to the ions from the RF field. This process is known as *RF heating*, and has been alluded to in the previous section. RF heating occurs when the time-dependent Coulomb force between ions in an RF field is irregular⁵. It is the dominant source of energy gain in ion clouds. RF heating has previously been studied in the context of interactions of ions with ultra-cold buffer gasses [142, 143]

⁴ Though some would argue that to be called a crystal it should exhibit periodicity, which is not the case.

⁵ Note 16 in Appendix e

and ion ensembles [144–149], but has not been considered in the context of crystals with low ion number, as, for example, registers of ionic qubits.

Ions in a cloud are no longer suitable to be used as qubits for quantum computation, and need to be efficiently returned to the crystal state, a process known as *recrystallization*. While laser cooling techniques can be employed to remove energy from the ions, the opposing increase in energy due to RF heating hinders or even prevents recrystallization. In this section, the properties of RF heating are investigated. As the motion of melted ions in an RF field is chaotic⁶ it is inconceivable to attain generalized analytic descriptions of the ion motion [138, 150, 151]. However, with our simulation we can numerically analyze dynamics of melted ions, allowing us to focus on the energy exchange mechanism of RF heating in detail.

The investigation of RF heating is carried out in the following steps:

- 1. **Numerical simulations:** a particle motion simulator is used to track the position of ions in an RF field. The simulation results show that RF heating can be described by discrete events, ion-ion collisions.
- 2. **Simplified model:** Results from the numerical simulations are used to develop approximate analytical expressions of RF heating, based on models for ion-ion collisions. These expressions are the basis for a simplified simulation of ion cloud energy dynamics, that avoids the computational overhead involved in tracking the motion of all ions in an RF field.
- 3. **Experimental validation:** We experimentally produce controlled melting events and estimate cloud energy by monitoring changes in the cloud's fluorescence. These results are compared to results from the full numerical and simplified simulations

Before detailing on this investigation, a more detailed description of ion clouds and RF heating is required, and is provided in the following section.

4.3.2 Ion clouds

We are unaware of literature in which the term "ion cloud" is clearly defined. However, from an intuitive point of view, the fundamental difference between an ion crystal and a cloud is well-known. In the crystal phase, ions have a uniquely defined average position, whereas in the cloud phase, all ions share an average position, the center of the trap. Viewing the fluorescence of trapped ions on a camera makes this distinction abundantly obvious: one either sees a grid of points, or just one elliptical blur. Even in the case of two trapped ions, one cannot simply visually distinguish two melted ions. For most of the work presented here, a more thorough definition of ion cloud is not necessary. The intuitive understanding of what it means for ions to be melted is sufficient. Still, it is worth mentioning two cases of ambiguity:

Firstly, it is generally agreed upon that a single trapped ion cannot form an ion cloud. This holds even at high energy, despite that one would observe a blur on a camera. Unlike in an ion cloud, a hot ion's motion is periodic, and there is no phase transition between a ground-state cooled ion⁷ and a hot one

Secondly, the distinction between cloud and crystal is not obvious in the energy regime corresponding to the phase transition between the two. For example, consider a two-ion crystal that due to a background-gas collision has gained some energy — just enough that the two ions can potentially swap positions [140]. In the absence of laser cooling, the two ions mostly remain in unique positions, but occasionally switch positions. One cannot claim that ions have unique average positions, thus cannot be considered crystallized. On the other hand, a camera readout would still suggest two distinct regions of increased fluorescence, not at all like a cloud. In Ref. [138], this energy range is dubbed the "quasi-periodic" regime, which separates the crystal and chaotic phases. This regime represents an area of hysteresis [152] between periodicity⁸ and chaos: a crystal that heats up into this energy range

⁶ Note 17 in Appendix e

⁷ Of course, a single cold ion is also not described as being "crystallized."

⁸ Or regularilty, at the very least

can remain crystallized, while a cloud that cools down into this range can remain melted. Overall, the quasi-periodic range is a gray area, both from a semantic and physical point of view.

Despite the lack of precision in qualitatively defining an energy that separates the crystal and cloud phase, we can analytically estimate this energy: the formation of a cloud is expected when the energy associated with ion motion (i.e., its kinetic energy) is much larger than an ions Coulomb energy. The ratio of these energies is known as the Coulomb coupling constant, Γ , and is usually written as [145]

$$\Gamma = \frac{1}{4\pi\epsilon_0} \frac{q^2}{a_{\rm ws} k_B T} \tag{4.1}$$

with k_BT the ion's thermal energy, q the ions' charge, and

$$a_{\rm ws} = \left(\frac{4}{3\pi\rho}\right)^{1/3} \tag{4.2}$$

the Wigner-Seitz radius. Here, ρ is the particle density. The density is ill-defined for low particle numbers. Previous works [145, 153, 154] have found the transition from crystal to cloud to occur for a coupling constant between $\Gamma=150$ and 200, but note that this is valid for large ion numbers (> 50). For low ion numbers, where $a_{\rm ws}$ is approximated to be the average distance between neighboring ions d, the crystal-cloud transition occurs⁹ in the range of $\Gamma\approx0.5$, which suggests that the mean energy per ion required to "break" an ion crystal is much higher for small crystals.

The separation between two ions aligned along the RF-free axis is given by

$$d = \left(\frac{q^2}{2\pi\epsilon_0 m\omega_z^2}\right)^{\frac{1}{3}},\tag{4.3}$$

with m the ion's mass and ω_z the axial motional frequency. The thermal energy k_BT in Eq. 4.1 can be defined in terms of the ion's kinetic energy $(1/2)m\langle v^2\rangle\approx (3/2)k_BT$, where the mean square-velocity $\langle v^2\rangle$ includes the contribution from the RF field. The mean velocity can be derived using the equation for an ion's position $r_k(t)$ given in Eq. 2.6 of Section 2.1, where we introduced the ion's secular motional amplitude a_k , and the Mathieu q-parameters, q_k , in three directions $k \in \{x,y,z\}$. The average velocity can be approximated [91] as

$$\left\langle v^2 \right\rangle = \lim_{T \to \infty} \frac{1}{T} \int_0^T (\dot{r}(t))^2 dt \approx \sum_k \left(\frac{a_k^2 \omega_k^2}{2} + \frac{a_k^2 q_k^2 (\Omega_{RF}^2 + \omega_k^2)}{16} \right). \tag{4.4}$$

Combining and rearranging Eqs. 4.1, 4.3, and 4.4, allows us to attain an expression for the oscillation amplitudes at the melting phase transition for a two-ion crystal:

$$a_k \approx \frac{2^{7/6}}{\sqrt{\Gamma}\sqrt{8\omega_k^2 + q_k^2(\omega_k^2 + \Omega_{RF}^2)}} \left(\frac{\omega_z q^2}{\sqrt{2}\epsilon_0 m\pi}\right)^{\frac{1}{3}}.$$
 (4.5)

Plugging in 'typical' trap parameters (40 Ca $^+$, $\omega_{\{x,y,z\}}/(2\pi)=\{3,3,1\}$ MHz, $\Omega_{RF}/(2\pi)=35$ MHz), we find the crystal-cloud transition, $\Gamma=0.5$ to correspond to average oscillation amplitudes of 8 µm in the axial direction and 2 µm in the radial directions. This is comparable to the ion-ion separation, $a_{WS}=5.6\,\mu\mathrm{m}^{10}$.

⁹ This is based on our own simulations. I am unaware of a literature value.

¹⁰ Note 18 in Appendix e

4.3.3 Heating in time-dependent potentials

As previously described, RF heating is an energy exchange that takes place due to interaction of the RF field and ions' Coulomb interaction. Here we provide a qualitative description of the physical process involved in this energy exchange. For this, we now (finally) define *energy* and *energy exchange*.

At any moment in time, the energy E_{tot} of a set of trapped ions is given by the sum of their kinetic energies V_{kin} and potential energies. The sources of potential energy are the trap's DC field potential V_{DC} , the RF field potential V_{RF} , and the Coulomb potential V_{Coul} . The total energy at any point in time t is thus given by:

$$E_{\text{tot}} = \sum_{i} \left[qV_{\text{DC}}(\vec{r}_{i}) + qV_{\text{RF,0}}(\vec{r}_{i})\cos(\Omega_{\text{RF}}t) + V_{\text{kin}}(v_{i}) + \frac{1}{2} \sum_{j \neq i} V_{\text{Coul}}(\vec{r}_{i}, \vec{r}_{j}) \right]$$
(4.6)

for ions (indexed i) with positions \vec{r}_i and velocities v_i , and with charge q. $V_{RF,0}$ is the amplitude of the RF potential, and Ω_{RF} its frequency.

In the context of describing an ion's energy change over time, also known as ion heating, this total energy is not very helpful: Ions experience driven motion due to the time-dependent RF potential $V_{\rm RF}$, so their energy is constantly changing. This affects *all* components of Eq. 4.6, not just the $V_{\rm RF}$ term, since the positions \vec{r}_i and velocities v_i are driven by the RF field. It is therefore convenient to separate the motion of trapped ions into two distinct timescales: *micromotion* which describes the oscillation synchronous with the RF field, and *secular motion*, which describes the motion in an effective static potential, the *pseudopotential*. The total energy can likewise be separated into the energies corresponding to these two types of motion, as

$$E_{\text{tot}} = E_{\text{RF}} + E_{\text{sec}} \tag{4.7}$$

where E_{RF} is the RF component of energy, and E_{sec} the secular component. When we refer to energy changes, heating and cooling, of trapped ions, we imply changes in secular energy, E_{sec} .

The secular energy is given by:

$$E_{\text{sec}} = \sum_{i} \left[qV_{\text{DC}}(\vec{r}_i^{(\text{sec})}) + qV_{i,\text{ps}}(\vec{r}_i^{(\text{sec})}, t) + \right]$$

$$(4.8)$$

$$V_{\text{kin}}(v_i^{(\text{sec})}) + \frac{1}{2} \sum_{j \neq i} V_{\text{Coul}}(\vec{r}_i^{(\text{sec})}, \vec{r}_j^{(\text{sec})})$$

$$(4.9)$$

which looks similar to the total energy, but full ion motion is replaced by secular position and velocity, $\vec{r}_i^{(\text{sec})}$ and $v_i^{(\text{sec})}$. The time-dependent RF potential is replaced by the pseudopotential, given by:

$$V_{i,ps}(\vec{r}_i^{(sec)}) = \frac{q}{4m_i \Omega_{PF}^2} \left| \nabla V_{RF,0}(\vec{r}_i^{(sec)}) \right|^2. \tag{4.10}$$

In the crystal phase, the secular motion can simply be expressed as normal modes of motion, with distinct frequencies, as detailed in Section 2.1.4. Obtaining the secular energy analytically is then straightforward, as the ion's motion contains a limited number of sinusoidal contributions and can be described and analyzed in the pseudopotential approximation. The RF energy does not couple to the secular energy, so, in the absence of other external influences (e.g. laser cooling, electrode surface noise, etc.), secular energy is conserved in the crystal phase.

If a collection of crystallized ions melts, ions experience aperiodic motion due to irregular ion-ion Coulomb interactions. The frequency spectra corresponding to the secular motion and micromotion broaden and overlap due to this aperiodic motion. This allows energy from micromotion E_{RF} to be transferred to the secular energy E_{sec} . RF heating¹¹ is this energy exchange between secular and RF energy, mediated through non-periodic Coulomb interaction. Secular energy is, in this case, not conserved.




Figure 4.2: Schematic representation of two ions' secular energy change due to Coulomb interaction in an RF field. While following a trajectory governed by Coulomb repulsion, ions additionally experience a force from the RF field, which is unequal due to the field's position dependence. In this example, the RF phase is such that during the moments of increased Coulomb interaction, the relative RF field is aligned with the direction of motion, both during the approach and withdrawal of the two ions. This leads to an increase of ΔW in motional energy after the collision. On the right, the filled circle represents the position of ion 1 as it increases and decreases in the Coulomb potential with the aid of the RF force, whereas the dotted circle is the position if no RF is applied.

The mechanism of energy exchange can be outlined with a simple one-dimensional model, shown in Figure 4.2. In a static potential, two ions approach each other, experiencing opposing Coulomb forces, and repel, as denoted by the dashed lines. In an oscillating potential, ions deviate approximately sinusoidally (solid lines) from this path. Since the amplitude of the RF field is dependent on the position in the trap, the two ions experience different forces from the RF field. In the example in Figure 4.2, directly before the moment of closest proximity, the difference in RF force reduces the distance between the ions, compared to the distance if no RF field was applied. The ions in this example therefore have more Coulomb energy at small distances than they would have in a static potential (this difference is denoted by $1/2\Delta W$). As the ions begin to repel, the RF field has switched sign, and now the difference in RF force aids in separating the ions. The added velocity that the ions gain from this force, in addition to the now transferred extra Coulomb energy, contribute to a total energy of ΔW added to $E_{\rm sec}$. This process would *remove* energy from $E_{\rm sec}$ if the phase of the RF field had been shifted by π .

The schematic in Figure 4.2 provides a qualitative description of the mechanism of energy transfer. In practice, such "head-on" collisions do not occur in three dimensions, and the RF phase will generally not line up with the Coulomb force as schematically presented. Still, this schematic description gives a qualitative intuition of how the RF field induces energy change when ions experience a time-dependent Coulomb interaction. In the following sections, a more quantitative description of RF heating is provided, by analyzing simulations of the trajectories of ions in a time-dependent trapping potential. First though, an overview of how these simulations work is provided.

4.4 SIMULATION OVERVIEW

The simulation is a numerical integration of the motion of ions (indexed i), with mass m_i , at positions $r_{i,k}(t)$ and velocities $v_{i,k}(t)$, under the influence of a total force $F_{i,k}^{(T)}$. The index k denotes the Cartesian direction, $k = \{x, y, z\}$. The classical equations of motion that need to be solved are then

$$\ddot{r}_{i,k}(t) = \frac{F_{i,k}^{(T)}(v_{i,k}(t), \vec{r}(t), t)}{m_i},\tag{4.11}$$

where the vector \vec{r} denotes the set of $r_{i,k}$ for all i and k. The energy regimes covered in this section, pertaining to non-crystallized ions, are more than five orders of magnitude higher than that of motional quanta, so a classical approach is justified.

Simulations of ion dynamics have been carried out with Matlab, using a built-in ordinary differential equation (ODE) solver and is publicly available [155]. The Matlab ODE solvers find numerical solutions for first order initial value differential equations of the form

$$\dot{\vec{r}} = f(\vec{r}(t), t),\tag{4.12}$$

where $\vec{r}(t)$ is a set of time-dependent parameters. This solver solves only first-order differential equations (DEs). The set of second-order equations of motion, Eq. 4.11, must be expanded into a set of first-order equations. This is done by introducing the substitution $v_{i,k} = \dot{r}_{i,k}$. We thus solve the set of equations

$$\dot{r}_{i,k}(t) = v_{i,k}(t)$$
 (4.13)

$$\dot{v}_{i,k}(t) = \frac{F_{i,k}^{(T)}(v_{i,k}, \vec{r}(t), t)}{m_k}.$$
(4.14)

For three dimensions and N ions, a system of 6N differential equations are numerically solved. The solver provides a list of times $\{t_n\}$ and parameters $\{r_{i,k}(t_n), v_{i,k}(t_n)\}$, for a set of initial parameters $\{r_{i,k}(t_0), v_{i,k}(t_0)\}$. The returned times do not necessarily coincide with time-steps used within the solver. Therefore, returned parameters are interpolated from the internally calculated parameters.

4.4.1 Contributing forces

There are four contributing forces that are included in the simulation:

- The force from the static DC electric field generated by the trap electrodes
- The time-dependent force from the RF electric field generated by the trap electrodes.
- The Coulomb force between ions
- Interaction with the Doppler cooling beam

The implementation of these forces is discussed below.

-DC electric field-

The force on an ion with charge q in a static potential V_{DC} is given by

$$F_{i,k}^{(\mathrm{DC})} = -q \nabla V_{\mathrm{DC}}.\tag{4.15}$$

We determine the position dependent ∇V_{DC} analytically, using the description of the DC potential discussed in Section 2.3.1, in which the full DC potential defined as a linear combination of potentials in a spherical harmonic basis, $V_{DC} = \sum_{l,n} v_{l,n} Y_{l,n}$, where $v_{l,n}$ is a scaling coefficient¹². For most

¹² It's difficult to assign a more constructive name than 'scaling coefficient' for $v_{l,n}$, since its dimension, and thus physical interpretation, depends on the order l, as $[v] = V/m^l$. For l = 1, this term represents the electric field strength, in $V m^{-1}$. For l = 2, it is the magnitude of the field curvature, in $V m^{-2}$

of the simulations described in this thesis, two components from this basis are used: the primary 'end-cap'-type DC potential

$$Y_{2.0} = 2z^2 - x^2 - y^2, (4.16)$$

and a potential that lifts degeneracy in the radial direction,

$$Y_{2,2} = x^2 - y^2. (4.17)$$

Here, z is chosen to be the direction along the axial trap axis. The corresponding electric fields $\{E_x, E_y, E_z\}_{l,n} = -\nu_{l,n} \nabla Y_{l,n}$ are then

$$E_{x,2,0} = 2x\nu_{2,0},$$
 $E_{y,2,0} = 2y\nu_{2,0},$ $E_{z,2,0} = -4z\nu_{2,0}$ $E_{x,2,2} = -2x\nu_{2,2},$ $E_{y,2,2} = 2y\nu_{2,2},$ $E_{z,2,2} = 0.$ (4.18)

Static homogeneous fields, given by the terms $Y_{1,\{-1,0,1\}}$, can be included as

$$E_{y,1,-1} = -\nu_{1,-1},$$
 $E_{z,1,0} = -\nu_{1,0},$ $E_{x,1,1} = -\nu_{1,1},$ (4.19)

though are generally kept at zero. Such terms can be included to investigate the effects of excess micromotion caused by uncompensated stray charges (see Section 2.3.3).

To summarize: the types and magnitudes of field potentials are selected by choosing values of $v_{l,n}$. At each time-step in the simulation, a force on each ion i, in direction k is calculated using $F_{i,k} = q_i \sum_{l,n} v_{l,n} E_{k,l,n}$, where all values of $E_{k,l,n}$ are determined analytically. The solver includes expressions for electric fields of all spherical harmonic potentials up until the $l = 4^{\text{th}}$ degree.

-RF electric field-

The force on an ion due to the trap's RF potential V_{RF} is

$$F_{ik}^{(RF)} = -q\nabla V_{RF}, \tag{4.20}$$

noting that in contrast to the DC case, the RF potential is time-dependent. The same harmonic expansions as for the DC case is used. The simulations in this section almost exclusively use only the term $Y_{2,2}$, which represents the saddle-type RF potential associated with most linear Paul traps. Similar to the DC case, we define the magnitude of the RF potential through a factor $v^{(RF)}$, where the l, n subscript is dropped since they are generally kept at l=2 and n=2. The factor $v^{(RF)}$ is half the amplitude of the RF potential field curvature¹³ in V m⁻². The force on an ion in the RF potential at time t is then

$$\begin{cases}
F_{i,x} \\
F_{i,y} \\
F_{i,z}
\end{cases} = \begin{cases}
-2x \\
2y \\
0
\end{cases} q_i \nu^{(RF)} \cos\left(\Omega_{RF}t + \phi_{RF}\right) \tag{4.21}$$

with Ω_{RF} the trap drive frequency, and ϕ_{RF} the trap drive's phase.

The simulation allows for the time-dependent RF potential to be substituted by an effective static potential, the *pseudopotential*, with forces given by:

$$\begin{cases}
F_{i,x} \\
F_{i,y} \\
F_{i,z}
\end{cases} = \begin{cases}
-2x \\
-2y \\
0
\end{cases} \frac{1}{m} \left(\frac{qv^{(RF)}}{\Omega_{RF}}\right)^2 \tag{4.22}$$

¹³ Ordinarily, one would use $\nu^{(RF)}$ to denote the amplitude of the curvature of the RF potential. However, the definition of $\nu^{(RF)}$ used here swallows up the factor 1/2! in the second order Taylor expansion of the potential, so $\nu^{(RF)}$ is half of the amplitude. I prefer this notation because it is more consistent with $\nu_{l,n}$ in the DC case

Using the pseudopotential in simulations neglects RF motion of the ions and provides an approximation for the remaining secular motion. This can be immensely useful for numerical simulations where one is confident that the pseudopotential approximation is accurate, since the integration step size can be relaxed by an order of magnitude, reducing simulation time.

For our purposes, as explained in Section 4.2, the time-dependent RF potential cannot be neglected. Still, the option to simulate ion trajectories within a pseudopotential is a useful tool for debugging, performance testing, and comparative studies.

-Coulomb interaction-

The Coulomb force on ion i due to the co-trapped ions j is

$$F_{i,k}^{(\text{Coul})} = \frac{1}{4\pi\epsilon_0} \sum_{j \neq i} \frac{q_i q_j}{|\vec{r_i} - \vec{r_j}|^2} (r_{i,k} - r_{j,k}). \tag{4.23}$$

with ϵ_0 the vacuum permittivity, and q_i and q_i the charge of ion i and j.

-Doppler cooling-

Typical ion trap experiments use Doppler cooling to remove energy from ions. The force that the Doppler cooling beam imparts upon ions thus plays a fundamental role in their dynamics.

Doppler cooling is a stochastic photon absorption and emission process, typically spanning a manifold of many electronic levels. For example, a ⁴⁰Ca⁺ ion is typically Doppler cooled in an eight-level manifold using 397 nm and 866 nm light. The stochastic and discrete nature of emission and absorption suggests that one must include randomized impulses to ions when numerically integrating their motion. Such discrete events are not well-accommodated by the built-in ODE solvers provided by Matlab, many of which are optimized to use internal variable time-steps. That is not to say that such processes cannot be adapted in numerical integration, but this requires specialized numerical methods.

A convenient solution is to substitute the discrete random impulses of emission and absorption by an average continuous damping force, and to approximate the many-level cooling manifold by a two-level system, with an effective spontaneous decay constant, Γ . Given a beam with an on-resonance coupling strength Ω and a detuning from resonance δ , an ion with velocity \vec{v} experiences an effective force given by [81]

$$F^{(D)} = \frac{\Gamma}{2} \frac{\Omega^2 / 2}{\Omega^2 / 2 + \Gamma^2 / 4 + (\delta - \vec{k}_{\rm D} \cdot \vec{v})^2} \hbar \vec{k}_{\rm D}, \tag{4.24}$$

where $\vec{k}_{\rm D}$ is the Doppler cooling beam's wavevector and \hbar is the reduced Planck constant.

We do note here that this method fails to accurately portray ion dynamics when they are cooled near the Doppler limit [81], since quantization of motion and photon recoil are not included. Ions are always orders of magnitude above the Doppler cooling limit in the simulations presented in this chapter.

4.4.2 Analytic description of surface trap potentials

In general, the simulation models the RF and DC fields as 'ideal' harmonic potentials. The potentials of real ion traps deviate from this ideal, having higher-order terms, which are more prominent further from the trap center. For a few crystallized ions in a 3D Paul trap, such anharmonicities can generally be neglected, as the range of the positions of the ions is small with respect to the size of the trap. In smaller traps (i.e. traps with lower ion-electrode separation such as surface traps) such anharmonicities may not necessarily be neglected and can affect the normal modes of motion [66]. Experimentally, the most detectable indication that the trap potential of a surface trap is anharmonic is its finite trap depth: Perfectly harmonic traps would be infinitely deep, which, judging by our rate of ion loss, is clearly not the case with surface traps. The simulation includes the option to take the non-ideal potential of a surface trap into account, as outlined below.

For analyzing ion motion in anharmonic trapping potentials, one could include higher-order spherical potential terms, using values of $v_{l,n}$ with l > 2. However, to accurately describe the potential of a surface trap, also far from the trap center, one requires many high-order terms¹⁴, which is mathematically cumbersome. A more realistic way to analytically model the potential of a surface trap is to follow the method of [156]: Assuming an infinitely large surface trap composed of square electrodes with no gaps, the potential V_e generated by applying a voltage v_e on an electrode e is given by

$$V_e(x,y,z) = \frac{v_e}{2\pi} \sum_{i,j=\{0,1\}} \left((-1)^i (-1)^j \tan^{-1} \left[\frac{(x_i - x)(z_j - z)}{y\sqrt{y^2 + (x_i - x)^2 + (z_j - z)^2}} \right] \right). \tag{4.25}$$

Here, the electrode's opposing corners have coordinates $(x_0, 0, z_0)$ and $(x_1, 0, z_1)$.

We generate an electrode layout similar to the trap discussed in Section 3.2.1.2 (though ignoring the gaps between electrodes and the optical access slot) and calculate the voltage set $\{\nu_e\}$ required to produce a desired trap potential, using the methods discussed in Section 2.3.1. The full potential is then given by the sum of electrode potentials, $\sum_e V_e$. Expressing the potential in the simulation in this fashion is useful because it faithfully models the anharmonicity and trap depth inherent to surface traps and thus provides an effective tool to make a side-by-side comparison of an ideal harmonic trap and surface traps. We make such a comparison and analysis in Section 4.8.

4.5 FULL ION DYNAMICS SIMULATIONS

In this section, numerically calculated ion trajectory data is used to analyze properties of RF heating. The simulator is described in Section 4.4. We refer to these simulations as "full," to contrast them with simplified simulations later in this work.

4.5.1 Simulation overview

In our simulations, we track the dynamics of two trapped $^{40}\text{Ca}^+$ ions. We use trapping parameters that match typical experimental values [106], with motional frequencies of $\{\omega_x, \omega_y, \omega_z\} = 2\pi\{3.3, 3.6, 1.1\}$ MHz, where z is the direction with no RF potential. At the start of the simulation, ions are placed in their crystallized equilibrium positions¹⁵. One ion has zero initial velocity, while the other is given an initial velocity in a random direction, mimicking a collision with a background gas particle. An initial kinetic energy of 1.4 meV ($\approx 80\,\text{m s}^{-1}$) is chosen, as it is marginally more than the required energy to melt the crystal [157]. This initial energy is, incidentally, in the same order of magnitude as a typical value of energy transfer induced by a collision with a background gas particle. Most simulations contain 5 ms of trajectory data¹⁶.

4.5.2 Extracting secular energy

To investigate RF heating, we must be able to extract the total secular energy from simulation data, which can be obtained by using Eq. 4.9. The kinetic energy of an ion with mass m_i is $V_{kin} = (1/2)m_iv_i^2$. The Coulomb energy between particles i and j with charge q_i and q_j is given by

$$V_{\text{Coul}} = \frac{1}{4\pi\epsilon_0} \frac{q_i q_j}{|\vec{r}_i - \vec{r}_j|},\tag{4.26}$$

¹⁴ Including terms up to l=4 is not sufficient, and this already contains 25 $\nu_{l,n}$ terms.

¹⁵ The equilibrium positions can be analytically calculated for a two ion crystal, for known trapping parameters. However, in our case, we can numerically obtain them by running the simulation with arbitrary starting points, but with a strong velocity damping force, which quickly brings ions to their equilibrium positions. These positions are used as starting positions in further simulations.

¹⁶ For a two-ion trajectory simulation on my laptop, a 5 ms trajectory takes about 20 minutes of simulation time, and contains roughly 1 GB of output data

with the vacuum permittivity ϵ_0 . $V_{DC}(\vec{r})$ and $V_{RF}(\vec{r})$ can be analytically obtained from simulation parameters (see Section 4.4).

The simulation, however, does not provide the ions' secular motion, $\vec{r}_i^{(\text{sec})}$. This motion thus needs to be extracted from the full simulated ion trajectory. One cannot fairly obtain secular motion by simply low-pass filtering the position data, in the hope of removing the part of the Fourier spectrum that describes the higher-frequency micromotion: The signals corresponding to secular motion and RF motion overlap spectrally. In particular, we wish to retain motional information that results from ion-ion Coulomb interaction, which in ion clouds have timescales similar to the trap drive period¹⁷. The secular motion is obtained from the simulated motion as follows:

The secular motion is obtained from the simulated motion as follows: The simulated time-dependent position of ion i is denoted as $\vec{r}_i^{(0)}$, and is to be separated into a secular and RF component, $\vec{r}_i^{(\text{sec})}$ and $\vec{r}_i^{(\text{RF})}$.

The equation of motion that governs the position $\vec{r}_i^{(0)}$ is given by

$$\ddot{\vec{r}}_{i}^{(0)} = -\frac{q}{m} \nabla V_{\text{RF},0}(\vec{r}_{i}^{(0)}) \cos{(\Omega_{\text{RF}}t)}. \tag{4.27}$$

$$\ddot{\vec{r}}_i^{(\text{sec})} + \ddot{\vec{r}}_i^{(\text{RF})} = -\frac{q}{m} \nabla V_{\text{RF},0} (\vec{r}_i^{(\text{sec})} + \vec{r}_i^{(\text{RF})}) \cos(\Omega_{\text{RF}} t). \tag{4.28}$$

If the secular frequency $\omega_{\{x,y\}}$ is much lower than the RF drive frequency Ω_{RF} , the RF component of position can be described by the equation of motion

$$\ddot{\vec{r}}_i^{(RF)} \approx -\frac{q\nabla V_{RF,0}(\vec{r}_i^{(sec)})}{m}\cos(\Omega_{RF}t). \tag{4.29}$$

This approximation is valid if the amplitude of \vec{r}_i^{RF} in one oscillation period is small enough such that $\nabla V_{\text{RF},0}(\vec{r}_i^{(\text{sec})})$ changes negligibly. Since simulations do not directly provide $\vec{r}_i^{(\text{sec})}$, an iterative approach is used, where initially the simulated positions $\vec{r}_i^{(0)}$ are used as an approximation for the secular motion: $\vec{r}_i^{(\text{sec})} \approx \vec{r}_i^{(0)}$. The RF component of the position is then

$$\vec{r}_i^{(\text{RF})} \approx -\frac{q\nabla V_{\text{RF},0}(\vec{r}_i^{(0)})}{m\Omega_{\text{RF}}^2}\cos(\Omega_{\text{RF}}t). \tag{4.30}$$

The secular motion is iteratively approximated by removing the RF component from the full simulated positions:

$$\vec{r}_{i}^{(1)} = \vec{r}_{i}^{(0)} - \vec{r}_{i}^{(RF)}$$

$$= \vec{r}_{i}^{(0)} + \frac{q \nabla V_{RF,0}(\vec{r}_{i}^{(0)})}{m \Omega_{RF}^{2}} \cos(\Omega_{RF} t).$$
(4.31)

As $\vec{r}_i^{(1)}$ is a better approximation for secular motion than $\vec{r}_i^{(0)}$, we can improve our estimate for $\vec{r}_i^{(\text{RF})}$ in Eq. 4.30. Higher order adjustments to the secular position can thus be found iteratively:

$$\vec{r}_i^{(n+1)} = \vec{r}_i^{(0)} + \frac{q\nabla V_{\text{RF},0}(\vec{r}_i^{(n)})}{m\Omega_{\text{RF}}^2}\cos(\Omega_{\text{RF}}t). \tag{4.32}$$

Note that Equation 4.32 remains an approximation for the secular motion and is not an exact solution even as $n \to \infty$.

Figure 4.3(a) demonstrates how a simulated trajectory (projected in one dimension) is adjusted using 4.32 in several iterations to remove the RF. The trajectories $\vec{r}_i^{(2)}$ and $\vec{r}_i^{(3)}$ are visibly indistinguishable. The remaining motion is approximately secular. We can further quantify the performance of the iterative RF removal method: for this, a separate simulation of a single trapped ion is run, which

¹⁷ Another method to remove the RF component is to isolate (interpolated) trajectory data at stroboscopic intervals corresponding to the period of the RF drive [158]. Also here, however, trajectory information relating to ion-ion Coulomb interaction is lost

produces a trajectory that exhibits both secular motion and micromotion, but no unpredictable behavior due to Coulomb interaction. In this case, the theoretical secular motion is purely sinusoidal, and secular energy is constant. We can thus analyze how well a perfect sine fits the numerically obtained trajectory $r^{(n)}$ for various iterations n. The goodness-of-fit is then quantified by the fit's coefficient of determination, i.e., its R-squared value [159]. Figure 4.3(b) shows the R-squared value as function of n (right axis). An additional metric is to gauge the stability of the calculated secular energy using Eq. 4.9, which in the ideal case is constant. The stability of the numerically obtained secular energy is shown in Figure 4.3(b)(left axis), expressed in terms of the ratio of the standard deviation to the mean of the energy. The ion trajectory $\vec{r}_i^{(n)}$ and secular energy $E_{\rm sec}$ both change negligibly for orders higher than n=3. However, these metrics do not reach zero, as this method only provides an approximate solution.




Figure 4.3: (a) Secular motion is extracted from simulated trajectories in iterative steps. Ideal secular trajectories should be sinusoidal, and the secular energy should be constant. (b) shows goodness-of-fit (R-squared) of the calculated secular trajectories fit to a sine curve, and the ratio of deviation in calculated energy to its mean.

The secular energy E_{sec} is obtained using Eq. 4.9 with the secular positions $\vec{r}_i^{(n)} \approx \vec{r}_i^{(\text{sec})}$ and the corresponding velocities $\vec{v}_i^{(n)}$. In this thesis, when referring to the ion cloud's energy, the secular energy E_{sec} is implied, with n=3. In the following section, RF heating is investigated by analyzing the secular energy of a simulated cloud of ions.

4.5.3 Simulation results

Figure 4.4(a) shows traces of the energy $E_{\rm sec}$ over time, for several simulation runs. All simulations start with identical parameters, except for a randomly chosen RF phase, which reflects that a collision with a background particle can occur at any time during the RF-drive cycle. In every trace, energy increases over time, but not necessarily continuously. The change in secular energy is attributed to RF heating. Although the only difference between the individual simulations is the initial RF phase, there is a large variation in the development of energy over time, resulting in energies ranging from about 4 to 30 meV after 5 ms. This variation attests the chaotic nature of melted ion dynamics.

Upon closer inspection, despite this unpredictable behavior, quasi-stable solutions do exist: a prominent example in Figure 4.4(a) is denoted by the light green area. Here, starting from after 1 ms, the secular energy remains comparatively steady. Such areas of stability occur only when the sinusoidal position of the two ions are out of phase by nearly π , in each dimension. We will see in Section 4.6.1.2 why this condition leads to increased stability.

It is clear from the traces in Figure 4.4(a) that one could never hope to accurately predict the course of energy gain of ions of a single melting event in an experimental setting. However, by simulating multiple melting events, an average energy gain can be estimated. The thick dark blue line in Figure 4.4(a) is an average of the individual simulations, which increases approximately with the square-root




Figure 4.4: Secular energy of simulated ion clouds. (a) Multiple simulation runs with nearly identical starting conditions (thin gray lines) exhibit a large spread in energy change. The average of these runs (thick dark blue line) increases smoothly over time, following a square-root-of-time trend. The green area highlights an example of a quasi-stable ion trajectory (b) A close-up of one of the traces from (a) shows that energy changes at discrete moments in time, coinciding with peaks in Coulomb energy, also shown in (b).

of time. This trend of the average increase of energy is a characteristic of RF heating, and is observed regardless of which trap parameters are used (for example, motional frequencies, ion mass, etc).

To further examine the dynamics that lead to energy changes, we take a trace from Figure 4.4(a) as an example, and show a short time interval in Figure 4.4(b). Additionally, the Coulomb energy of the two ions in this interval is plotted. Secular energy does not change continuously, but at discrete points in time. These discrete changes in energy are correlated with moments of increased Coulomb interaction. This is in agreement with the intuitive description of RF heating presented in Section 4.3.3: RF heating occurs due to multiple ions' joint interaction with the RF field and Coulomb force. As a reminder, this interaction does not influence ion crystals because the frequencies of the Coulomb force and RF motion are spectrally separated. In contrast, in an ion cloud the Coulomb force exhibits rapid changes that occur on timescales similar to the RF drive period.

From here on, we will refer to the moments of increased Coulomb interaction that may lead to a change in energy as a *collision*¹⁸. Individual changes in energy as a result of collisions are denoted by ΛW

It is instructive to gather statistics of these energy changes, as it provides insight to the orders-of-magnitude associated with RF heating. We thus collect a set of energy changes $\{\Delta W\}$, by evaluating the secular energy before and after any peak in Coulomb energy. We additionally extract statistics about the ions locations within the trap where they undergo these energy changes, and the times at which they occur. All traces from Figure 4.4(a) are used for this collection.

Figure 4.5(a) shows the distribution of the location that the collision occurs, with respect to the trap center. All collisions occur within 4 µm of the trap center, regardless of the amplitude of the ion motion. For reference, the distribution of ion location with respect to the trap center is shown, which displays ion oscillation amplitudes of $\sim 10-30 \text{ µm})^{19}$.

Figure 4.5(b) shows a polar histogram depicting at which phase of the RF drive cycle a collision occurs. Collision times are not correlated to the RF phase. However, collisions that lead to a large change in secular energy have a strong tendency to occur when the magnitude of the RF field is highest (at extrema in the RF cycle). This is shown in the same polar histogram in Figure 4.5(b), where an arbitrary energy cut-off of $|\Delta W| > 1$ meV has been chosen for visual clarity. The tendency for high changes in energy to require the collision to be at maxima of the RF field, as graphically depicted, aligns well with the schematic depiction of RF heating presented in Figure 4.2. Here secular energy change is depicted as occurring when the RF component of relative velocity has an opposing sign directly before and after the peak of the collision. Times where the RF component of velocity is zero

¹⁸ The fact that energy is not necessarily conserved makes it an *inelastic* collision.

¹⁹ Note 20 in Appendix e

correspond to times where the RF field is at a maximum or minimum. The simulated data therefore corroborates our intuitive description of RF heating provided in Section 4.3.3.

Figure 4.5(c) shows a histogram of the energy changes that occur in the example trace of Figure 4.4(b). The majority of collisions lead to near-negligible changes in energy²⁰. The distribution in energy changes is qualitatively symmetric. However, the mean of the distribution ($\mu = 5.2 \,\mu\text{eV}$) and average collision rate ($\bar{f}_{\text{coll}} = 528$ collisions per millisecond), leads to an overall increase of energy of approximately 14 meV after 5 ms.

Figure 4.5(d) shows the standard deviation of $\{\Delta W\}$, for various bins of peak Coulomb interaction energy. It is apparent that a larger change of secular energy requires higher Coulomb interaction. The inverse is not necessarily true: high Coulomb interaction does not always result in high changes in secular energy. Returning to the intuitive picture of Section 4.3.3, the magnitude and sign of energy change is essentially luck of the draw, depending on how the collision lines up with the phase of the RF field. We observe that collisions with Coulomb interaction energies below 0.5 meV, indicated by the dotted line in Fig. 4.5(d), do not result in noticeable changes in the system's total energy. Thus, from here onward, we will consider an ion-ion collision to be an event where the interaction energy exceeds this threshold.

In summary, the figures and statistics provided in this section give an insight into the parameters and their orders-of-magnitudes that comprise RF heating: hundreds of collisions occur per millisecond, collisions induce energy changes of tens to hundreds of microelectronvolts, and these changes overall lead to an average RF heating rate of several millielectronvolts per millisecond. The results shown in this section have used a specific parameter set, and are thus not generalized. Despite this lack of generality, these simulations provide an important take-away: *RF heating is not a continuous process but occurs at discrete moments of high Coulomb interaction energies.* In the following section, this notion is used as a basis for a simplified simulation that estimates RF heating rates, allowing us to produce generalized results at a low computational cost.




Figure 4.5: Statistics of ion-ion collisions (a) Histogram of collision location with respect to the center of the trap. Collisions close to the center of the trap compared to the ions' motional oscillation amplitudes. (b) Polar histogram of RF drive phases that collisions occur. Collisions that lead to high secular energy change are correlated with the RF phase. (c) Histogram of the magnitude of energy changes ΔW due to collisions. (d) The standard deviation of these energy changes $\sigma_{\Delta W}$ is shown in bins of the collisions' peak Coulomb energies.

4.6 COLLISION MODEL

This section provides a generalized quantitative description of the parameters involved in energy dynamics in RF heating. This description is used as a basis for a simulation that tracks the energy of an ion cloud subject to RF heating. This simplified simulation allows us to estimate RF heating rates at a significantly lower computational cost compared to the full ion dynamics simulation of the previous section.

4.6.1 Model parameters

To model RF heating, events of increased Coulomb interaction – ion-ion collisions – are considered. Although two melted ions continuously experience a Coulomb interaction, in the previous section it was shown that below a certain threshold of Coulomb energy, changes in secular energy are negligible. We can thus consider collisions to be discrete events where the Coulomb interaction surpasses the threshold that is found from the full simulation. Having been separated into discrete events, RF heating can be modeled in two steps: 1) A collision induces a change in secular energy, ΔW , and 2) secular energy remains constant until a new collision occurs, separated by a time $t_{\rm coll}$. The simplified simulation alternates between generating new values of ΔW and $t_{\rm coll}$. Values are sampled randomly from a sampling distribution, based on models presented in the following sections.

4.6.1.1 Collision energy

We draw on results from the full ion dynamics simulation presented in Section 4.5.3 to derive and validate a model that describes the energy change due to a collision, ΔW .

The change in energy of any dynamic system of particles i can be expressed in terms of the forces \vec{F}_i acting on the particles with velocities \vec{v}_i , as

$$\Delta W = \int \sum_{i} \vec{F}_{i} \cdot \vec{v}_{i} dt. \tag{4.33}$$

In a Paul trap, the total force on each ion is the sum of static and RF fields, and the Coulomb force: $\vec{F}_i = \vec{F}_i^{(DC)} + \vec{F}_i^{(RF)} + \vec{F}_i^{(Coul)}$. The velocity of the ions can be expanded into the contributions of secular and RF motion, $\vec{v}_i = \vec{v}_i^{(sec)} + \vec{v}_i^{(RF)}$. The product $\vec{F}_i \cdot \vec{v}_i$ can thus be expanded into six components.

To deduce which of these components carry information about energy change due to RF heating, we remind ourselves of the phenomenological description presented in Section 4.3.3. Here, it was claimed that energy is exchanged from ions RF motion to its secular motion through an interaction with the Coulomb potential. We thus conjecture that a transfer of energy occurs between the components $I_{\rm RF} = \int \vec{F}^{({\rm Coul})} \cdot \vec{v}^{({\rm rf})} dt$ and $I_{\rm sec} = \int \vec{F}^{({\rm Coul})} \cdot \vec{v}^{({\rm sec})} dt$. Figure 4.6(a) displays the evolution of these two integrals over the duration of one collision. The Coulomb energy $E_{\rm Coul}$ and total energy $E_{\rm sec}$ are plotted alongside. These energies are offset to put their initial values at zero. After the collision, the values of the two integrals $I_{\rm RF}$ and $I_{\rm sec}$ are equal but opposite. The magnitude of these changes matches the change in $E_{\rm sec}$. Around the peak of the collision, the moment of highest Coulomb interaction, energy is temporarily stored in $E_{\rm Coul}$ which comes primarily from the energy integral described by the ions secular motion, $I_{\rm sec}$. In fact, in an RF-free scenario, we have $E_{\rm Coul} = I_{\rm sec}$ by definition. In this case, $I_{\rm sec}$ would be zero when integrated over the full duration of the collision (assuming $E_{\rm Coul}$ is negligibly small before and after the collision).

Energy exchange between the RF field and the ions secular motion can be interpreted as a transfer of energy originating from the forces produced by $I_{\rm RF}$. We see in the example in Figure 4.6(a) that a loss of energy in $I_{\rm RF}$ over the duration of a collision corresponds to a gain in energy in $I_{\rm sec}$, which ultimately corresponds to an increase in secular energy $E_{\rm sec}$. The Coulomb force has thus mediated a transfer of energy from $E_{\rm RF}$ to $E_{\rm sec}$. The change in secular energy is given by $\Delta W = I_{\rm sec} = -I_{\rm RF}$. While both integrals describe the secular energy change, $I_{\rm sec}$ contains sharp peaks at moments of high

Coulomb interaction. $-I_{RF}$ is thus a smoother, and therefore more intuitive, representation of ΔW . In summary, ΔW can be expressed as

$$\Delta W = -\int \sum_{i} \vec{F}_{i}^{(\text{Coul})} \cdot \vec{v}_{i}^{(\text{rf})} dt. \tag{4.34}$$

Using $\vec{F}_1^{(Coul)} = -\vec{F}_2^{(Coul)}$ for a two-ion collision, the change in energy reduces to

$$\Delta W = -\int \vec{F}_1^{\text{(Coul)}} \cdot \Delta \vec{v}^{\text{(rf)}} dt, \tag{4.35}$$

with the difference in RF velocity between the two ions $\Delta \vec{v}^{(rf)} = \vec{v}_1^{(rf)} - \vec{v}_2^{(rf)}$.




Figure 4.6: Comparison of secular energy metrics. (a) During a single collision, energy is transferred between the RF and secular components of motion. $I_{\rm sec}$ and $I_{\rm RF}$ respectively gain and lose ΔW energy, which marks an overall increase in $E_{\rm sec}$ (b) Comparison of total secular energy ($E_{\rm sec}$, Eq. 4.9), the energy change integral (Eq. 4.35), and the simplified integral (Eq. 4.42). The integrals are offset to 2 meV at t=0 to reflect the initial energy.

The above description might seem more speculative than physically founded, and should be scrutinized more thoroughly. We validate Eq. 4.35, using numerical data from one of the full ion dynamics simulations from Section 4.5.3, Figure 4.4(b). The energy development of Eq. 4.35 is shown in Figure 4.6(b) as the accumulation of $\Delta W + E_{\rm init}$, with $E_{\rm init} = 2$ meV to reflect the initial energy of the system. For reference, the total energy $E_{\rm sec}$ (see Eq. 4.9) of the system is also plotted, showing close agreement. To quantify this subjective statement, as has been done for Figure 4.5, we collect energy changes due to collisions, defined as the difference in secular energy before and after events where the Coulomb energy exceeds 0.5 meV. These energy differences are extracted both from the absolute secular energy and the integral of Eq. 4.35. Comparing the two sets of energy differences, which ideally are identical, a correlation between the two energy metrics is calculated to be an R-squared [159] value of 90%. This confirms that Eq. 4.35 can faithfully describe the change in secular energy.

Further simplification of this integral allows us to express ΔW in a form that readily translates into a computationally efficient model for predicting energy changes. This simplification can be achieved by approximating $\Delta \vec{v}^{(RF)}$. The relative RF velocity is estimated knowing the ions' positions relative to each other, and the phase of the RF field:

For a saddle-type RF potential

$$V_{\rm RF}(\vec{r}_i, t) = \frac{1}{2} \psi_{\rm RF} \left(r_{i,x}^2 - r_{i,y}^2 \right) \cos \left(\Omega_{\rm RF} t \right) \tag{4.36}$$

with potential curvature ψ_{RF} , the force on an ion i with charge q at position $\vec{r}_i = [r_{i,x}, r_{i,y}, r_{i,z}]$ is given by

$$\vec{F}_i^{(RF)}(\vec{r}_i, t) = -q \nabla V_{RF} \tag{4.37}$$

$$= [-r_{i,x}, r_{i,y}, 0]q\psi_{RF}\cos(\Omega_{RF}t). \tag{4.38}$$

We have shown in Section 4.5.2 that we can estimate an ions motion due to this force by assuming that the ions' positions $r_{i,x}$ and $r_{i,y}$ are approximately constant during an oscillation cycle with frequency $\Omega_{\rm RF}$. Likewise, in assuming that $r_{i,x}$ is constant over an oscillation cycle, one can approximate $r_{i,x} \approx r_{i,k}^{(\rm sec)}$. Integrating $\int \vec{F}^{(\rm RF)} dt = m \vec{v}^{(\rm RF)}$ with mass m allows us to approximate the RF component of the velocity as

$$\vec{v}_i^{(\text{RF})} \approx \left[-r_{i,x}^{(\text{sec})}, r_{i,y}^{(\text{sec})}, 0 \right] \frac{q\psi_{\text{RF}}}{m\Omega_{\text{RF}}} \sin(\Omega_{\text{RF}}t). \tag{4.39}$$

The difference in RF velocity is then

$$\Delta \vec{v}^{(\text{RF})} \approx \left[-\Delta r_x^{(\text{sec})}, \Delta r_y^{(\text{sec})}, 0 \right] \frac{q\psi_{\text{RF}}}{m\Omega_{\text{RF}}} \sin(\Omega_{\text{RF}}t). \tag{4.40}$$

Here $\Delta r_x^{(\text{sec})}$ and $\Delta r_y^{(\text{sec})}$ are the ions' separation in their secular motion.

The Coulomb force is given by:

$$\vec{F}_1^{\text{(Coul)}} = \frac{1}{4\pi\epsilon_0} \frac{q^2}{|\Delta \vec{r}|^3} \Delta \vec{r} \tag{4.41}$$

with $\Delta \vec{r} = \vec{r}_1 - \vec{r}_2$ the ions' separation, and ϵ_0 the vacuum permittivity. The Coulomb force $\vec{F}_1^{(\text{Coul})}$ is dominated by the secular motion of the ions, such that $\Delta \vec{r} \approx \Delta \vec{r}^{(\text{sec})}$. Equation 4.35 can thus be approximated as

$$\Delta W \approx \frac{q^3 \psi_{\rm RF}}{4\pi \epsilon_0 m \Omega_{\rm RF}} \int \frac{\left(\Delta r_x^{\rm (sec)}\right)^2 - \left(\Delta r_y^{\rm (sec)}\right)^2}{\left|\Delta \vec{r}^{\rm (sec)}\right|^3} \sin(\Omega_{\rm RF} t) dt. \tag{4.42}$$

Since this formulation of ΔW has taken steps of approximation, its validity should be examined. Similar to having analyzed full ion dynamics simulation data with Eq. 4.35, we now analyze the same data set with Eq. 4.42. Results, shown in Figure 4.6(b), are in agreement with the results generated with Eq. 4.35, and with the total energy E_{sec} . The energy changes in these results are compared to energy changes derived from E_{sec} , resulting in an R-squared correlation of 82%, confirming that Eq. 4.42 provides a good approximation of energy change in a collision.

It may not be obvious why this formulation is considered to be a 'simplified' version of Eq. 4.35. The benefit of 4.42 is that very little information about the trap and RF dynamics of the ions is required. Notably, to estimate energy change, neither the ions' absolute position within the trap, nor the RF components of their motion are required. In Section 4.6.2, this equation is used in a simulation to estimate ion energy. This simplified simulation benefits from the computational advantage that it does not require information about particle dynamics at time-scales of the RF drive.

4.6.1.2 Collision rate

The previous section discusses an analytic expression for energy change ΔW as a result of an ion-ion collision. To estimate the change in energy over time resulting from multiple collisions, the time between collisions $t_{\rm coll}$ must be determined. We generate a time between collisions by first determining an *average* collision rate $\bar{f}_{\rm coll}$. This rate is the number of collision events that occur for particles with oscillatory motion in three dimensions. Our method for determining $\bar{f}_{\rm coll}$ is outlined in this section.

We start by assuming that each ion moves sinusoidally, characterized by their secular frequencies ω_k , in three Cartesian directions, $k = \{x, y, z\}$. The exercise is then to analytically estimate how often in a given time ions are within a certain range of each other. A secular approximation of the ions motion is used. While in reality ions experience RF driven motion, this additional motion is not dominant: ion motion from simulations and analytic approximations in Section 4.5.3 show that for typical trapping parameters, the amplitude of the RF-driven motion is less than 10% of the amplitude of secular motion²¹. Furthermore, the approximation is justified since the RF component of motion of the two ions is strongly correlated when they are within collision range. Thus, the RF motion can be neglected when considering timing of collision events. Additionally, as discussed in Section 4.5.3, most collisions occur near the center of the trap, where the amplitude of the RF driven motion is minimal.

In the absence of a Coulomb force, the secular position of ions is sinusoidal in time, with distinct frequencies and amplitudes in each dimension. The ion i has amplitudes $a_{i,k}$, and phase $\phi_{i,k}$. Its position at time t is thus given by $r_{i,k} = a_{i,k} \sin{(\omega_k t + \phi_{i,k})}$. We pose a criterion that defines when a collision occurs, as being events where the absolute distance between ions in each dimension separately is less than a chosen collision threshold, r_c , i.e. $\forall k : |r_{1,k} - r_{2,k}| < r_c$. This criterion ensures that the positions of the two ions are both within the bounds of a cube of sidelength r_c , but necessarily within an absolute distance r_c of each other.

To estimate the frequency that the condition is satisfied in all three dimensions simultaneously, we must first determine *how long* ions are within a projected collision range in each dimension. This duration is denoted as Δt_k .

We start by analyzing the one-dimensional (x) case: the separation between ions, $d_x = r_{1,x} - r_{2,x}$ is sinusoidal with frequency ω_x [160], and can thus be expressed as

$$d_x = a_{d,x}\sin(\omega_x t + \phi_{d,x}) \tag{4.43}$$

with amplitude $a_{d,x}$ and phase $\phi_{d,x}$ given by

$$a_{d,x} = \sqrt{a_{1,x}^2 + a_{2,x}^2 - 2a_{1,x}a_{2,x}\cos(\Delta\phi_x)}$$
(4.44)

$$\phi_{d,x} = \tan^{-1} \left(\frac{a_{1,x} \sin(\phi_{1,x}) - a_{2,x} \sin(\phi_{1,x})}{a_{1,x} \cos(\phi_{1,x}) - a_{2,x} \cos(\phi_{2,x})} \right), \tag{4.45}$$

with $\Delta \phi_x = \phi_{2,x} - \phi_{1,x}$. Rewriting Eq. 4.43 for t gives:

$$t = \frac{1}{\omega_x} \left[\sin^{-1} \left(\frac{d_x}{a_{d,x}} \right) - \Delta \phi_x \right]. \tag{4.46}$$

The amount of time Δt_x that the two particles are within collision range r_c is

$$\Delta t_x = t(d_x = r_c) - t(d_x = -r_c) \tag{4.47}$$

$$=\frac{2}{\omega_x}\sin^{-1}\left(\frac{r_c}{a_{d,x}}\right) \tag{4.48}$$

In one dimension, non-interacting ions are within (projected) collision range exactly twice per oscillation period (or be continuously within range). The collision criterion in one dimension, $|d_x| < r_c$, can thus be represented by a pulse wave²²:

$$B_x(t) = \begin{cases} 1, & \text{if } (t \mod T_x) \le \text{Re}\{\Delta t_x\}. \\ 0, & \text{otherwise.} \end{cases}$$
 (4.49)

with the time between collisions $T_x = \pi/\omega_x$, being half the oscillation period. If $r_c > a_{d,x}$, then Δt_x becomes complex, which is unphysical. Fortunately, an easy fix is to use only the real part of Δt_x ,

²¹ The amplitude ratio of RF-driven motion to secular motion in the absence of excess micromotion is approximately given by half of the Mathieu *q*-parameter (see Eq. 2.6).

²² Also known as a pulse train, a rectangular pulse waveform, a rectangular wave or a periodic rectangular function

which caps off at $\Delta t_x = T_x$ when $r_c \le a_{d,x}$. This case corresponds to the two ions continuously being within collision range. Accordingly, B_x is then continuously 1, which is physically accurate.

This model is extended to three dimensions with three pulse waves, $B_{\{x,y,z\}}$, each characterized by periods $T_{\{x,y,z\}}$, and pulse durations $\Delta T_{\{x,y,z\}}$. The collision criterion is satisfied in all three dimensions when all three pulse trains are simultaneously 1; in other words, when $B_x B_y B_z = 1$. The pulses of $B_x B_y B_z$ occur at irregular intervals, but we can determine an average rate at which pulses occur. This average rate of pulses represents the ion-ion collision rate, and is given by²³

$$\bar{f}_{\text{coll}} = \frac{\Delta t_x \Delta t_y \Delta t_z}{T_x T_y T_z} \left(\frac{1}{\Delta t_x} + \frac{1}{\Delta t_y} + \frac{1}{\Delta t_z} \right). \tag{4.50}$$

The average time between collisions is $\bar{t}_{\text{coll}} = 1/\bar{f}_{\text{coll}}$. A derivation of the collision rate is given in Appendix a.

4.6.2 Simplified RF heating model

In the previous sections, models have been presented to estimate the secular energy change ΔW due to ion-ion collisions in a cloud, and the average duration between collisions, $\bar{t}_{\rm coll}$. In this section, these models are used to construct a simplified RF heating simulation. The simulation uses trap parameters and the ions' secular motion amplitudes to alternatingly generate an energy change and a duration until a subsequent collision. This section outlines how the models from the previous section are used to randomly generate ΔW and $\bar{t}_{\rm coll}$ and how these values are used to update the parameters of the simplified simulation. The simulation is used to estimate an RF heating rate for various trap parameters.

4.6.2.1 Simulation overview

As described in the previous section, the motion of each ion is described as being sinusoidal in three orthogonal directions (indexed k), with parameters $a_{i,k}$ the amplitudes of ion i, and a relative phase between the ions $\Delta \phi_k$. These parameters are depicted in Figure 4.7. It is assumed that these parameters are constant between collision events. The total secular energy,

$$E_{\text{sec}} \approx \sum_{i,k} \frac{1}{2} m \omega_k^2 a_{i,k}^2, \tag{4.51}$$

therefore also remains constant between collisions, as is seen in the full ion dynamics simulations shown in Figure 4.4. Collisions result in updates of the parameters $a_{i,k}$ and $\Delta \phi_k$, which may also reflect a change of energy ΔW .

The approximation of Eq. 4.51 is based on the assumption that the Coulomb energy is negligible while ions are far outside of the collision threshold. The simulation is initialized with a chosen energy E_0 , distributed randomly²⁴ over the amplitudes $a_{i,k}$. The initial phases $\Delta \phi_k$ are chosen randomly and uniformly.

The parameters $a_{i,k}$ are applied to Eq. 4.50, using values of Δt_k given in Eq. 4.47, to produce an average collision rate \bar{f}_{coll} . The collision threshold is chosen to be $r_c = 1.44\,\mu\text{m}$, corresponding to a Coulomb energy of 0.5 meV, based on analysis of simulations described in Section 4.5.3. There it is shown that collisions with Coulomb energies less than 0.5 meV do not result in appreciable changes in secular energy.

Eq. 4.50 produces an average duration between collisions ($\bar{t}_{coll} = 1/\bar{f}_{coll}$) for a given ion motion parameter set. This is, however, only an *average* of some distribution function; to generate a new collision time, one could simply assume $t_{coll} = \bar{t}_{coll}$, but a more fair choice can be made by drawing randomly from a realistic probability distribution function. To determine this function, we note that the time between collisions is usually longer than a typical secular motion period. Due to their aperiodic nature, subsequent collisions therefore occur at uncorrelated intervals. The probability distribution function

²³ Note 22 in Appendix e

²⁴ Note 23 in Appendix e

for collision times can thus be modeled as an exponential distribution, $P(t) = \bar{f}_{\text{coll}} \exp(-t\bar{f}_{\text{coll}})$. A discussion on the validity of this choice of function is offered in Appendix a. A random number from a known probability distribution function P(t) can be generated by drawing a random value p, uniformly between 0 and 1, and transforming it with the inverse cumulative distribution function (ICDF) [161] of P(t). A random time t_{coll} is thus generated from the ICDF of the exponential distribution, given by $-\log(1-p)/\bar{f}_{\text{coll}}$.

The soonest a collision can occur is half an oscillation period after the previous collision. Such a case could occur, for example, when $a_{i,k}$ is small ($\leq r_c$) in two dimensions. In that case, ions move apart sinusoidally along only one axis, and meet each other again half an oscillation period later. Earlier collisions cannot occur. Therefore, if the chosen time t_{coll} is less than $\min_k T_k$, this time is discarded and a new collision time is randomly generated. The validity (and usefulness) of this step is also discussed in Appendix a.

Now that a time until a collision is established, we turn to generating a random energy change ΔW . The equation that models energy change is Eq. 4.42. The only requirements for this equation are known trap parameters (the amplitude of the RF potential curvature ψ_{RF} , and the RF frequency Ω_{RF}) and the relative secular motion of the two ions $\Delta \vec{r}^{(\text{sec})}$. The timescale of a collision event is typically much lower than the period of secular motion. We can thus model the dynamics of a collision as being independent of forces from the trapping potentials, and solely dependent on Coulomb forces. The method of attaining $\Delta \vec{r}^{(\text{sec})}$ is then to simulate two particles approaching and move apart from each other, under the influence of Coulomb force, devoid of trapping forces.

This Coulomb collision simulation is run with a collision trajectory based on values derived from $a_{i,k}$ and $\Delta \phi_k$. The results of the simulation, along with a randomized RF phase, are applied to equation 4.42 to produce an energy change ΔW . Based on simulated trajectories and ΔW , a new set of parameters $a_{i,k}$ and $\Delta \phi_k$ are produced.

The Coulomb collision simulation works as follows: Firstly, the velocity of the two particles at the moment of collision is to be calculated. A given set $a_{k,i}$ and $\Delta\phi_{i,k}$ determines the velocity at the time of collision, since these parameters define at what phase in the ions' oscillation the collision occurs. This phase can be found by referring to Eq. 4.43, where the separation between ions is given by $d_k = a_{d,k} \sin(\omega_k t + \phi_d)$, shown in the bottom left panel of Figure 4.7. Since a collision occurs when $d_k \approx 0 \ll a_k$ for all k, we can pose that at the time of collision, $\omega_k t = -\phi_d$. Other solutions, multiples of π , are dropped without loss of generality.

With the ions' position, $r_{i,k} = a_{i,k} \sin{(\omega_k t + \phi_{i,k})}$, the velocity is given by $v_{i,k} = a_{i,k} \omega_k \cos{(\omega_k t + \phi_{i,k})}$, still assuming that the RF contribution to motion is negligible for modeling a collision event. Replacing $\omega_k t$ with $-\phi_d$, and substituting ϕ_d using Eq. 4.45, we find velocities

$$v_{i,k} = a_{i,k}\omega_k \cos\left(-\tan^{-1}\left(\frac{a_{1,k}\sin(\phi_{1,k}) - a_{2,k}\sin(\phi_{2,k})}{a_{1,k}\cos(\phi_{1,k}) - a_{2,k}\cos(\phi_{2,k})}\right) + \phi_{i,k}\right),\tag{4.52}$$

which reduces to

$$v_{1,k} = \frac{\omega_k a_{1,k} (a_{1,k} - a_{2,k} \cos(\Delta \phi_k))}{\sqrt{a_{1,k}^2 + a_{2,k}^2 - 2a_{1,k} a_{2,k} \cos(\Delta \phi_k)}}$$
(4.53)

$$v_{2,k} = \frac{\omega_k a_{2,k} (a_{2,k} - a_{1,k} \cos(\Delta \phi_k))}{\sqrt{a_{1,k}^2 + a_{2,k}^2 - 2a_{1,k} a_{2,k} \cos(\Delta \phi_k)}}.$$
(4.54)

The ions' velocities \vec{v}_i are used as parameters for the Coulomb collision simulation. The simulation is a numerical integrator whose only force is the Coulomb force. A schematic outline of the parameters used and extracted from the simulation are shown in the right panel of Figure 4.7 (in two dimensions, for visual clarity), detailed below. Two points in a cube with side-lengths r_c are chosen at random, denoting their positions as $\vec{\chi}_i^0$. Two particles are initialized outside of the collision range defined by this box, by placing them at positions $\vec{\chi}_i^{(\text{start})} = \vec{\chi}_i^{(0)} - \vec{v}_i t_s$, effectively sending ions back in time by t_s from their collision positions. The time t_s is set to $t_s = 8r_c / \max_{i,k} |v_{i,k}|$, where the value 8 is chosen so that ions are placed far enough from each other such that the Coulomb energy is far below the collision




Figure 4.7: Schematic overview of variables described in the main text, as used in the simplified simulation. (Top left) Ion motion in one dimension. (Bottom left) Relative ion position. A collision occurs at phase ϕ_d in the secular cycle. This phase is used in generating initial positions $\vec{\chi}_i^{(s)}$ and velocities \vec{v}_i for the Coulomb collision simulation, depicted in the schematic on the right. See main text for a description of the variables.

threshold at the start of the simulation. $\vec{\chi}_i^{(\text{start})}$ and \vec{v}_i are the starting parameters for the simulation. The simulation is carried out for a duration $2t_s$. The simulation thus covers the action of two ions moving towards a common region, where they experience Coulomb interaction, and subsequently move away from each other again.

The time-dependent ion positions $\vec{\chi}_i(t)$ produced by the simulation represent a randomized secular path, \vec{r}_i , typical for the secular motion parameters $a_{i,k}$ and $\phi_{i,k}$. The simulated positions can therefore be applied directly to Eq. 4.42 to produce a change in energy ΔW . Since the collision time is uncorrelated with the phase of the RF field, a random phase is added to the argument of the sine.

The total energy $E_{\rm sec}$ of the simulation is defined in terms of $a_{i,k}$, as in Eq. 4.51. The values of ΔW that are produced with Eq. 4.42 are used to derive new values of $a_{i,k}$ and $\phi_{i,k}$, which update the value $E_{\rm sec}$. To update $a_{i,k}$, we use values from the Coulomb simulation: we assume the ions originate from the collision point, $\bar{\chi}_k = (\chi_{1,k}^{(0)} + \chi_{2,k}^{(0)})/2$, and are moving at velocity $v_{i,k}^{(f)}$, the final velocity of the simulation. This uniquely defines a new set of motion parameters, $a_{i,k}^{(n)}$ and $\phi_{i,k}^{(n)}$, as

$$a_{i,k}^{(n)} = \sqrt{\bar{\chi}_k^2 + \left(\frac{v_{i,k}^{(f)}}{\omega_k}\right)^2}$$
 (4.55)

$$\phi_{i,k}^{(n)} = \tan^{-1}\left(\frac{\bar{\chi}_k \omega_k}{v_{i,k}^{(f)}}\right) \tag{4.56}$$

with associated energies $E_{k,\text{sec}}^{(n)}$. The new amplitudes $a_{i,k}^{(n)}$, however, do not reflect the change in energy due to the RF field, since the final velocities $v_{i,k}^{(f)}$ are determined with an RF-free simulation. These velocity values must therefore be adjusted to conform with ΔW . To make this adjustment, Eq. 4.42 is separated into it's two sum components,

$$\Delta W = \Delta W_x + \Delta W_y \tag{4.57}$$

$$= c \int \frac{\left(\Delta r_x^{(\text{sec})}\right)^2}{\left|\Delta \vec{r}^{(\text{sec})}\right|^3} \sin(\Omega_{\text{RF}}t) dt + c \int -\frac{\left(\Delta r_y^{(\text{sec})}\right)^2}{\left|\Delta \vec{r}^{(\text{sec})}\right|^3} \sin(\Omega_{\text{RF}}t) dt \tag{4.58}$$

which define the energy changes in the two radial directions, with c the same prefactor as in Eq. 4.42. The *actual* update, $a_{ik}^{(\text{new})}$, is given by

$$a_{i,k}^{(\text{new})} = \sqrt{\left(a_{i,k}^{(\text{n})}\right)^2 + \frac{\Delta W_k}{m\omega_k^2}},$$
 (4.59)

which assumes that the excess energy ΔW is distributed equally over the two ions. The updated set of parameters can now once again be used to produce a new time until a following collision.

The method of simulating collision energy and duration between collisions readily expands into more than two ions by extending the parameter set $a_{i,k}$ and $\phi_{i,k}$, with $i = \{1,2,3...\}$. In this case, a collision time $t_{\text{coll}}^{(i,j)}$ is generated for all combinations of ion pairs $i \neq j$. The ion pair with the shortest collision time is selected to undergo a simulated collision. The parameters $a_{i,k}$ and $\phi_{i,k}$ of the chosen pair are updated to reflect a collision having occurred only between those two ions, using the method described above. This method is applicable under the assumption that collisions are predominantly between no more than two ions. Independent full ion trajectory simulations (with typical trapping parameters) show that for clouds of three, four, and five ions, about 3%, 4%, and 7% of collisions involve three or more ions. These values depend on the chosen trap parameters, so are merely intended as indications of how often a more-than-two-body collision can be expected to occur. We find that our simulation still works adequately for a five-ion cloud. Larger ion numbers will, however, require a different approach to simulating energy changes.

In summary, a given a set of parameters $a_{i,k}$ and $\phi_{i,k}$ is used as a basis for a randomized Coulomb collision simulation, from which an estimated energy change due to RF heating ΔW is extracted. The results of the collision simulation and this change of energy are used to update the values of $a_{i,k}$ and $\phi_{i,k}$ accordingly. With this updated set of parameters, a new collision time t_{coll} can once again be generated. This process is repeated until the sum of all collision times exceeds the desired simulation duration. E_{sec} is calculated with Eq. 4.51 for every step of the simulation, giving a time-dependent energy.

This method of simulating RF heating reduces the computation duration by more than three orders of magnitude, compared to the full ion dynamics simulation. This speed-up comes mainly from the fact that the simplified simulation does not need to track particle motion continuously at nanosecond timescales (corresponding to RF motion). Particle motion is only simulated at brief intervals, at collision events. Even at those moments, where Coulomb collision simulations are run, the method is still computationally inexpensive, as RF motion is neglected. The numerical integrator of the Coulomb collision simulation adapts its simulation step size according to the acceleration that the particles undergo. Therefore, determining ions' motion as they approach and retract from each other requires very little computation. Only the brief moment when ions are close enough that their mutual Coulomb force significantly modifies their trajectories requires finer simulation step sizes.

4.6.2.2 Simulation comparison

We now finally turn to employing the simplified RF heating simulation presented in the previous section. The first step is to benchmark the performance of the simplified simulation by comparing its results to those of the full ion dynamics simulation. The simplified simulation is executed with identical parameters as those used for Figure 4.4. The results of the two types of simulation are displayed in Figure 4.8. In Figure 4.8(a), the energy gain due to RF heating is shown as an average of 20 runs of the simulation. Also plotted is the average energy of 20 runs of the full ion dynamics simulation. The mean values of the two simulations are in good agreement, though the simplified

model underestimates the variability of all simulation runs, shown by the thin lines, denoting \pm one standard deviation. The collection of energy changes ΔW and duration between collisions $t_{\rm coll}$, taken from all simulation runs, are shown in Figure 4.8(b) and Figure 4.8(c) as histograms (normalized to be displayed as probabilities), shown for both simulation types. The collection of collision data from the full ion dynamics simulation is taken from events corresponding to peaks in Coulomb energy, where ion positions satisfy the collision criterion of $\Delta r_k < r_c = 2.8 \, \mu \text{m}$ in all dimensions k (which was previously determined to be a cutoff above which energy change due to RF heating is negligible). This same threshold is used in the simplified simulation. The distribution and occurrence frequency are similar for both simulation types. We conclude that the simplified energy simulations work reliably as an indicator for average change in energy due to RF heating²⁵.




Figure 4.8: Comparison of full ion dynamics simulation to simplified simulation, displaying (a) the average secular energy $E_{\rm sec}$ of 20 repetitions of both simulations (thick lines), standard deviation of the 20 repetitions (thin lines), and histograms of (b) energy differences in each collision and (c) time between collisions. Histograms are normalized to be displayed as probabilities

4.6.2.3 Analytic RF heating model

The main advantage of the simplified simulation (other than the confirmation that the models for ΔW and $t_{\rm coll}$ are realistic) is the computational speed-up compared to the full ion dynamics simulation in simulating RF heating. This speedup is crucial in order to investigate RF heating rates for a variety of trap parameters, from which more generalized statements about heating rates can be made.

Figure 4.9 displays the average development of energy over time due to RF heating, for various trapping parameters, using the simplified simulation. Unless otherwise specified, the simulations use two 40 Ca⁺ions, with motional frequencies $\omega_{\{x,y,z\}} = 2\pi\{3.4,3.3,1.1\}$ MHz and a 35 MHz trap drive frequency. Each trace is an average of 20 individual simulation results, each with randomly generated initial values of motional amplitudes $a_{i,k}$ and phases $\phi_{i,k}$, though constrained by a fixed initial energy of 3 meV, given by Eq. 4.51. The number of simulation runs per setting, 20, is chosen as a trade-off between determining a statistically meaningful mean rate of energy change, and restricting computational time.

Figure 4.9 shows traces for various (a) radial motional frequencies, (b) axial frequencies, (c) ion masses, and (d) numbers of ions in the cloud. As seen in Figure 4.9(a), lower radial motional frequencies result in a lower gain in energy for a melted crystal. This can be explained by the models in the previous section, where it is shown that both the rate of ion-ion collisions and the energy change per collision increases with increasing RF power (and radial motional frequency). This behavior is often observed experimentally, as many ion trapping experiments purposefully lower the RF voltage to assist in refreezing melted ions [71].

²⁵ Similar comparisons have been made for varying parameters such as motional frequencies and particle masses, to convince ourselves that we didn't just get lucky with the parameter set used in Figure 4.8.

As seen in Figure 4.9(b), a change in axial frequency has a less significant influence on the heating rate, compared to radial frequencies. Evaluating the models for collision energy change and time between collisions in the previous section, we note that the axial frequency does not directly contribute to energy changes, and only has a minor influence on duration between collisions. A lower axial frequency has a higher axial (half) oscillation period T_z , and thus has a reduced average collision rate \bar{f}_{coll} (see Eq. 4.50)²⁶.

In Figure 4.9(c), ion species with higher masses exhibit a higher rate of energy change, though it must be noted that to keep the motional frequencies constant with changing masses, the trapping potentials are adjusted accordingly. Figure 4.9(d) shows that clouds with a higher number of ions exhibit a larger increase in energy. This is understandable, as collisions are more frequent with more ions in the cloud. The slower initial onset of energy increase a higher ion number is because the initial energy of 3 meV is quickly distributed over all the ions, thus individual ions have lower average initial energies. Collisions are therefore initially less energetic, and thus are less susceptible to RF-induced energy changes.




Figure 4.9: Two-ion cloud energy dynamics, obtained with the simplified simulation. Unless otherwise specified, $m=40\,\mathrm{u}$, $\omega_z/(2\pi)=1.1\,\mathrm{MHz}$, $\omega_r/(2\pi)=3.3\,\mathrm{MHz}$, and $\Omega_\mathrm{RF}/(2\pi)=35\,\mathrm{MHz}$. We vary (a) the radial frequency, (b) the axial frequency, (c) the ions' mass, and (d) the number of ions. In (c), the trapping fields are adjusted to ensure the same motional frequencies for all masses.

The goal now is to generalize the dependence on various parameters, such as those displayed in Figure 4.9, into a single generalized analytic expression, based on the ansatz that energy change due to RF heating follows a diffusive model. For this, the simplified simulation is repeated with randomized values of the DC confining potential, RF potential, drive frequency, ion number, and ion mass. For each random setting, the simulation is repeated 10 times, from which an average energy is produced, as function of time $E_{\rm sec}(t)$. This amount of repetitions is, as before, a trade-off between attaining a statistically significant average, and minimizing computational time. The resulting energy curves, akin to those shown in Figure 4.9, can be characterized by a one parameter that quantifies the RF heating rate. Here, we draw an analogy between melted ion energy transfer and random-walk processes, such as diffusion due to Brownian motion. In such processes, if a variable is subjected to randomized

²⁶ The one-dimensional collision duration Δt remains approximately constant with decreasing ω_z , since $a_{d,z} \overset{\alpha}{\sim} \omega_z^{-1}$

changes, the probability distribution of its value spreads out over time, where the rate of spreading is characterized by a diffusion constant. In our model, the energy $E_{\rm sec}$ over time t follows a trend of $E \sim \sqrt{Dt}$, where D is the diffusion constant [144, 162]. We perform least-squares regressions between the model $E_0 + \sqrt{Dt}$ and our simulated average energy $E_{\rm sec}(t)$ to determine the diffusion constant. In this model the initial energy $E_0 = 3$ meV is not a fitting parameter, since this is fixed by the simulation.

The simulation is repeated for 500 random parameter settings, in which we alter the RF and DC field potential curvatures, the number of ions n, and their mass m. Each parameter setting produces a particular value of the diffusion constant D, which is used to derive a generalized expression for D, in terms of ion mass m, axial frequency ω_z , radial frequency ω_r , trap drive frequency Ω_{RF} , and number of ions n.

A polynomial model for *D* is used:

$$D = am^b \omega_r^c \omega_z^d \Omega_{RF}^e n^f \tag{4.60}$$

with estimated parameters a - f, displayed in Table 4.1. The uncertainty values represent a $\pm 34.1\%$ confidence interval in the least-squares regression between the polynomial model of Eq. 4.60 and the simulated results.

$E \sim \sqrt{Dt}$					
$D = am^b \omega_r^c \omega_z^d \Omega_{RF}^e n^f \text{ (eV}^2/\text{s)}$					
	Fit value	Uncertainty			
a	60	18			
b	1.0	0.05			
c	2.45	0.05			
d	0.52	0.04			
e	0.00	0.06			
f	2.96	0.04			

Table 4.1: Fit results for RF-induced energy diffusion model

This simple model provides an effective method to quantify the energy dynamics resulting from RF heating without resorting to numerical methods. Notably, this model is generally applicable to any linear Paul trap, and is easily applied by inserting known trapping parameters. Additionally, the model can be used to estimate an ion cloud heating rate at for a given ion cloud energy E with $E'(t) = \sqrt{D/(4t)} \approx D/(2E)$ (this is valid under the assumption that the initial energy E_0 is small compared to E, which is generally the case, as can be seen in Figure 4.9. The results in Table 4.1 agree with the previously drawn conclusions about RF heating rates:

For a fixed number of ions, the diffusion coefficient is most sensitive to changes in the radial motional frequency ω_{RF} , reinforcing the notion that reducing this parameter in an ion trap experiment (by reducing the RF voltage) is the most effective method of reducing the RF heating rate. The heating rate is to a lesser extent dependent on the axial motional frequency. The heating rate is strongly dependent on the number of ions in the cloud, due to the increased collision rate.

To summarize, we have developed a simple simulation that is capable of efficiently estimating an RF heating rate for a low number of trapped ions. By estimating these rates for various trap parameters, a generalized model for RF heating is established. The strength of this model is that it takes a chaotic system that to-date does not have a reliable analytic description, and attempts to bring a reasonable structure and quantification to it. While analytic models for RF heating have been presented previously [145, 146, 148], these models are only applicable for large ensembles of more than one-hundred ions, and use a mean-field theory to estimate heating rates. RF heating rates in those works are considerably higher, by many orders of magnitude, and do not readily allow extrapolation to determine heating rates at low ion numbers. In the context of quantum computation, where individual qubit registers are expected to contain low numbers of ions, such analytic descriptions are therefore not directly applicable. Realizing that the theory presented here limits itself to 5 ions, it remains to be emphasized

that investigation of higher ion numbers is warranted. This work provides a solid foundation to pursue such investigation, as it details models and tools that make ion dynamics simulations approachable.

4.6.3 Recooling an ion cloud

An obvious question that arises from the previous section is how to combat RF heating, such that an ion cloud recrystallizes. Melted ions are in an energy regime where the most effective method of removing energy is Doppler cooling. Doppler cooling is used universally in trapped ion (and neutral atom) experiments to extract energy from the particles. It allows trapped ions to become and remain crystallized, even in the presence of external heating processes, for example caused by electric field noise [29]. The rate of Doppler cooling can, in its most simple terms, be experimentally controlled with two parameters: the beam's power, in other words its coupling strength Ω to the cooling transition, and its detuning Δ from resonance of that transition. It is often the case in trapped ion experiments that these two parameters are chosen to bring ions as close as possible to the lowest possible energy that can be attained with Doppler cooling, the Doppler limit [163], given by the mean energy $\hbar\Gamma/4$, with Γ the spontaneous decay rate of the cooling transition. Reaching this limit requires a low beam power ($\Omega \ll \Gamma$) and a detuning of $\Delta = \Gamma/2$. These settings are, however, not ideal for efficiently cooling ions with energies far above the Doppler limit, since a high cooling rate typically requires high beam power, and the ions' Doppler shifts are much larger than the beam detuning. Melted ions are in this high-energy regime where standard Doppler cooling parameters are not optimal. When melted, RF heating and Doppler cooling are two competing mechanisms of energy change. If set improperly, Doppler cooling is outmatched by RF heating, and recrystallization will not occur. In this section, we estimate which Doppler cooling parameters enable efficient recrystallization, by implementing a simple cooling model into the simulations. The results in this section are limited to a cloud of two ⁴⁰Ca⁺ions.

Doppler cooling is a stochastic photon absorption and emission process, typically spanning a manifold of many of an ion's electronic levels. For example, a 40 Ca $^{+}$ ion is typically Doppler cooled in an eight-level manifold using 397 nm and 866 nm light, consisting of two $4S_{1/2}$, two $4P_{1/2}$, and four $3D_{3/2}$ levels. Proper implementation of Doppler cooling in a numerical ion dynamics simulation requires tracking the ions electronic state among these eight levels, and subjecting it to momentum kicks corresponding to photon absorption and emission. Transitions between states are dependent on spontaneous decay rates and transition selection rules, which in turn are dependent on beam polarization, power, detuning, and ion velocity. A full simulation that includes all such dependencies is outside the scope of this work 27 .

As described in Section 4.4.1, in our simulations we elect to simplify the stochastic dynamics of this eight-level system by applying a continuous time-averaged force acting on an effective two-level system. This can easily be considered quite a leap in approximation, so it is worth emphasizing that this section is intended to provide guidelines of laser cooling parameters, and not exact values. In our experiment, we typically blue-detune the 866 nm repump laser, which ensures that dark resonances due to Raman interaction with the red-detuned cooling beam [164] are avoided, and apply an abundance of beam power, to ensures that there is minimal population in the $4D_{3/2}$ state. We can then assume that the spontaneous decay rate is dominated by that of the $4S_{1/2}$ and $4P_{1/2}$ cooling transition. To avoid having to flip back to page 76: In a two-level system with a spontaneous decay rate Γ , the average Doppler cooling force on ion i with velocity \vec{v}_i is given by [81]:

$$\vec{F}_i^{(D)} = \frac{\Gamma}{2} \frac{\Omega^2 / 2}{\Omega^2 / 2 + \Gamma^2 / 4 + (\delta - \vec{k} \cdot \vec{v}_i)^2} \hbar \vec{k}, \tag{4.61}$$

with Ω the on-resonance coupling strength, δ the detuning of the Doppler beam from resonance (in radians per second), \vec{k} the wavevector of the beam, and \hbar the reduced Planck constant. This force is implemented directly as one of the continuously contributing forces in the full ion dynamics simulation.

²⁷ Though might have been included if the Covid lockdown had been longer.

As argued before, studying a broad parameter range is computationally costly when using the full simulation. It is therefore beneficial to include Doppler cooling in the simplified simulations, so that cooling parameters that lead to recrystallization can be efficiently found. Incorporating Doppler cooling in the simplified simulation is not as straightforward as for the full simulation, since the simplified model does not continuously track time-dependent velocities, as required for calculating $\vec{F}^{(D)}$. Doppler cooling is implemented as follows:

As described in Section 4.6.2, a duration t_{coll} is determined, after which the motional amplitude and phase parameters $a_{i,k}$ and $\phi_{i,k}$ are updated to reflect a change in energy due to a collision. In this update, the additional energy change due to interaction with the Doppler force is to be included. This energy is given by

$$\Delta W^{(D)} = \int \sum_{i} \vec{F}_{i}^{(D)}(\vec{v}_{i}) \cdot \vec{v}_{i} dt. \tag{4.62}$$

To calculate $\Delta W^{(D)}$ while avoiding simulating ion velocities in RF fields, an analytic expression for the ions' velocities \vec{v}_i is required. The RF contribution to motion cannot be neglected when assessing ion velocity, as seen in Figure 4.1. A secular approximation for \vec{v}_i is therefore not valid. To obtain an analytic expression for velocity that includes the influence of the RF field, we turn back to Eq. 4.32,

$$\vec{r}_i^{(n+1)} = \vec{r}_i^{(0)} + \frac{q\nabla V_{\text{RF},0}(\vec{r}_i^{(n)})}{m\Omega_{\text{RF}}^2}\cos(\Omega_{\text{RF}}t),\tag{4.63}$$

where it was assumed that $\vec{r}_i^{(n+1)}$ approximates the secular motion $\vec{r}_i^{(sec)}$, and $\vec{r}_i^{(0)}$ the full motion (secular and RF). Using $\nabla V_{\text{RF},0} = \psi_{\text{RF}}[r_x, -r_y, 0]$, this expression expands in each dimension as

$$r_{i,k}^{(n+1)} = r_{i,k}^{(0)} \left(1 + f_k \cos(\Omega_{RF}t) + f_k^2 \cos^2(\Omega_{RF}t) + \dots + f_k^n \cos^n(\Omega_{RF}t) \right), \tag{4.64}$$

with $-f_y = f_x = q\psi_{RF}/(m\Omega_{RF})^2$ and $f_z = 0$. In practice, secular motional frequencies are easier to measure than the RF field curvature ψ_{RF} , so it is useful to write $f_{x,y}$ in terms of motional frequencies, as $-f_y = f_x = \sqrt{\sum_k \omega_k^2}/\Omega_{RF}$.

In the limit that the sum in Eq. 4.64 is taken to infinity, $n \to \infty$, the equation simplifies to²⁸

$$r_{i,k}^{(\text{sec})} \approx \lim_{n \to \infty} r_{i,k}^{(n+1)} = \frac{r_{i,k}^{(0)}}{1 - f_k \cos(\Omega_{\text{RF}} t)}$$
 (4.65)

The secular motion is sinusoidal, with amplitude $a_{i,k}$ and frequency ω_k . The full motion of the ion can then be written as

$$r_{ik}^{(0)} = a_{ik} \sin(\omega_k t) \left(1 - f_k \cos(\Omega_{RF} t) \right),$$
 (4.66)

which was a roundabout way of getting to the familiar expression of the approximate ion motion as in Eq. 2.6 and Refs. [55, 69]. The velocities are the time-derivative of position,

$$v_{i,k} = a_{i,k} w_k \cos(\omega_k t) \left(1 - f_k \cos(\Omega_{RF} t)\right) \tag{4.67}$$

$$-a_{i,k}f_k\Omega_{RF}\sin(\omega_k t)\sin(\Omega_{RF}t). \tag{4.68}$$

This expression for velocity, which is a function of ion motional amplitudes $a_{i,k}$, is used to determine the change in energy from Doppler cooling, $\Delta W^{(D)}$. Within the span of time between collisions, the change in $a_{i,k}$ due to continuous Doppler cooling is small, which allows us to approximate it as

²⁸ Valid if f_k is smaller than 1. This is guaranteed by the trap's stability criterion, $f_k \approx q_k/2 \ll 1$, with Mathieu parameter q_k

being constant for this duration. The change in energy can be expressed for each ion and dimension separately

$$\Delta W_{i,k}^{(D)} = \delta t \sum_{n}^{N} F_{i,k}^{(D)}(\vec{v}_i(t_n)) v_{i,k}(t_n)$$
(4.69)

where the velocities have been discretized into N time-steps of δt , and are calculated over a timespan of $t_{\rm coll} = \delta t N$, the duration between collisions. Prior to analyzing energy change due to RF heating during a collision, the values of $a_{i,k}$ are updated to reflect the change in energy due to Doppler cooling for the duration leading up to the collision:

$$a_{i,k}^{\text{new}} = \sqrt{\frac{2\Delta W_{i,k}^{(D)}}{m\omega_k^2} + a_{i,k}^2}.$$
 (4.70)

We find that a time-step of $\delta t = 10$ ns is sufficiently small for the purpose of our simulation; smaller steps do not affect the simulated outcome, and result in longer computation time.

The simplified simulation including Doppler cooling is used to study which beam parameters are required to overcome RF heating to recrystalize an ion cloud. Some representative results of the simplified simulation are shown in Figure 4.10(a), shown for coupling strengths $\Omega/(2\pi)$ ranging from o to 80 MHz, for a detuning of $\delta/(2\pi)=-40$ MHz. These values are chosen because they represent reasonable beam powers in experimental conditions, and an interesting detuning where the changes in beam power have a noticeable effect on cooling power (unlike the typical experimental setting of $\Delta=\Gamma/2\approx 11$ MHz, where the cooling beam becomes rather useless for any power). The wavevector is chosen to be $\vec{k}=(2\pi/\lambda)[0.07,0.71,0.71]$, with $\lambda=397$ nm, which is similar to the wavelength and angle of incidence in our experimental setup. An initial ion cloud energy of 15 meV is chosen, which is a typical ion cloud energy after 5 ms of RF heating and reflects a bad-case-scenario that the Doppler beam is only switched on several milliseconds after a melting event had occurred. Trap parameters are identical to the default values used in Figure 4.9 ($\omega_z/(2\pi)=1.1$ MHz, $\omega_r/(2\pi)=3.3$ MHz, $\Omega_{RF}/(2\pi)=35$ MHz, and m=40 u).

The thick lines in Figure 4.10 are averages of 20 simulation runs, and the thin lines are the standard deviation of all runs. Results from the simplified simulation are compared to results from the full ion dynamics simulation with identical trapping and cooling parameters (dotted lines). As in the previous section, the average trends of the two simulations are in good agreement, although the simplified model underestimates the total spread of energy in individual runs. This incites confidence that the simplifications used when implementing Doppler cooling in the simplified model are justified²⁹. The various plots show that with increasing coupling strength, the Doppler cooling rate overcomes the RF heating rate.

To determine the requirements for the Doppler beam parameters required for recrystalizing, we run the simulation for detunings ranging from from 0 to $-400\,\mathrm{MHz}$, and coupling strengths ranging from 0 to $100\,\mathrm{MHz}$, and analyze the final energy after 5 ms of cooling, a typical time spent on Doppler cooling in an experimental sequence. The final energies of the simulations are plotted in Figure 4.10(b), as a function of the Doppler beam's coupling strength and detuning. The figure is subdivided into three regions: in region (i) the final energy is higher than the initial energy, and in regions (ii) and (iii) it is lower. In region (iii), the final energy is low enough for the ions to have recrystallized within the 5 ms simulation time. Region (i) represents a parameter regime where the RF heating rate is higher than the Doppler cooling rate. In this regime, ions remain melted, regardless of cooling duration. In region (ii), the cooling rate is higher than the heating rate, but ions likely do not recrystallize within the 5 ms cooling time. In typical experimental sequences used for ion-based quantum computation, ions are not continuously cooled, as sequences contain non-cooled computation steps. If Doppler cooling is not efficient enough to recrystallize ions before these steps, as would be the case if parameters from region (ii) were used, the cloud can subsequently reheat while the cooling beam is off. In such a cycle, ions can indefinitely remain melted. Therefore, in such an experimental sequence, Doppler cooling

²⁹ Though this does not prove that the two-level continuous force model for Doppler cooling is accurate, since both the full and simplified simulations use it.

parameters that ensure recrystallization are more stringent, corresponding to region (iii). Figure 4.10(b) shows that recrystallization in 5 ms is only ensured in experimental sequences if the Doppler beam is detuned between -100 and -300 MHz, with a coupling strength > 50 MHz. If the cooling beam is on continuously, this requirement is more relaxed, with $\Omega \gtrsim 20$ MHz and $\delta \gtrsim 20$ MHz, corresponding to regions (ii) and (iii) in Figure 4.10(b).




Figure 4.10: Simulated energy dynamics with Doppler cooling. (a) Ca - Ca cloud energy for various cooling coupling strengths, with detuning $\delta/(2\pi)=-40\,\mathrm{MHz}$. For comparison, results from both the full ion dynamics simulation and simplified simulation are shown. (b) Energy after 5 ms of Doppler cooling, for various detunings and coupling strengths.

4.7 EXPERIMENTAL VALIDATION

The simulations of the previous sections beg the question "Can we see these RF heating effects in the lab?" In our experiment, we have indirectly seen the simulation results in effect. The inclusion of additional hardware that implements a "refreeze" beam in every cycle of an experimental sequence has resulted in improved ion lifetimes, for both single- and mixed-species ion crystals. The settings for this refreeze beam were chosen based on the guidelines presented in the previous section.

It is, however, useful to attain a more quantitative experimental confirmation of the RF heating (and recooling) rates that the simulations have provided. In this section, we validate the RF heating simulations by comparing them to experimental ion cloud energy measurements. In these experiments, ion clouds are generated deterministically, and their energy is inferred by monitoring their fluorescence.

4.7.1 Experiment Overview

The experiments are performed on two $^{40}\text{Ca}^+$ ions, in the surface Paul trap described in Section 3.2.2. Fluorescence detection and Doppler recooling is done by off-resonantly exciting the $4S_{1/2}\leftrightarrow 4P_{1/2}$ transition at 397 nm. Undesired population in $3D_{3/2}$ is repumped with 866 nm. The repumper is blue-detuned by approximately 10 to 30 MHz, and its power is adjusted such that the ions' fluorescence rate is saturated, which ensures that an ions cooling dynamics is dominated by the 397 nm transition. The two-level Doppler cooling model used in the simulations is then a realistic approximation, with coupling strength Ω and detuning δ . The effective spontaneous decay rate Γ is assumed to be dominated by the decay rate of the $4P_{1/2} \to 4S_{1/2}$ transition for $^{40}\text{Ca}^+$, $\Gamma/(2\pi) = 21.6\,\text{MHz}$ [165].

In the experiment, the detuning of the 397 nm beam is adjusted by changing the set-point of the wavemeter's frequency lock. The coupling strength is adjusted by changing the input power into the double-pass AOM that enables the in-sequence switching. The values of Ω and δ require calibration,

linking them to experimental settings of wavemeter frequency and AOM power. This is done in our experiment by scanning these parameters, and determining the fluorescence rate R_f from a single 40 Ca $^+$ ion. This rate is then numerically compared by least-squares regression with the model:

$$R_f = A \frac{\Omega^2}{2\Omega^2 + \Gamma^2 + 4\delta^2} \tag{4.71}$$

where A is a constant proportionality factor. This model is a simplification of Eq. 4.61, where it is assumed that the velocity \vec{v} is negligibly small. To ensure this condition, the ion is first cooled with a red-detuned beam for $15\,\mathrm{ms^{30}}$. After cooling, the ion is probed with the Doppler beam for a duration of $500\,\mu\mathrm{s}$, with variable detuning and power. The short probe time ensures that even when the probe beam is blue-detuned, the ion's velocity remains low, as it changes negligibly in this duration. The ions velocity can, in that case, be neglected in the model.

To examine RF heating, the two-ion crystal is deterministically forced to melt. The goal is to transfer enough energy to the ions to generate a cloud, but ensure that the initial cloud energy is low enough that RF heating dynamics can be observed, without the risk of ion loss. To generate the cloud, a periodic force is exerted on the ions by applying an oscillating voltage on nearby trap electrodes, near resonance with two radial motional frequencies. This process is colloquially termed "tickling" [166], but is in practice way less cute than it sounds. The signal used is *two*-toned, since ions excited in the two radial directions require less total energy to guarantee a phase transition to a cloud, compared to excitation in one dimension. The initial cloud energy is thus lower, allowing us to observe a larger overall increase of energy after melting. The frequencies of the tickling pulse are detuned by about $-100\,\mathrm{kHz}$ with respect to the motional mode frequencies. This helps to avoid recooling the crystal before it melts [167, 168]: Motional modes in our anharmonic trapping potential decrease with increasing oscillation amplitude. The oscillation frequency therefore approaches resonance with the excitation field as the ions' energy increases.

The DC electrodes of the surface trap would be the logical choice for where to apply the tickle voltage. However, in our setup, signals applied to these electrodes are heavily filtered at the required frequencies. Also, the signal lines are shielded and require inconvenient bypassing for an external signal to be applied. We therefore elect to superimpose the tickle signal onto the RF trap drive, which results in a tickle voltage on the RF electrodes. While the RF resonator does filter frequencies unequal to the trap drive frequency, this effect is small compared to that of the DC filters.

The downside of using the RF electrode for tickling is that ions are ideally placed near the RF-null, the inflection point of the saddle-shaped RF potential, where the field strength is at a minimum. The tickle field couples poorly with ions whose micromotion is well-compensated. We therefore intentionally apply a bias field of about $\nabla V_k^{(\text{bias})} = 100\,\text{V}\,\text{m}^{-1}$, with $k = \{x,y\}$. This field displaces the ions by $q\nabla V_k^{(\text{bias})}/(m\omega_k)$ from the trap center, which is less than a micrometer for our trap parameters. Independent simulations confirm that this displacement negligibly affects the RF heating and Doppler cooling rates.

An experimental sequence consists of a tickle pulse, during which cooling beams are off, followed by a short (500 μ s) Doppler cooling pulse. During this pulse, fluorescence is collected by either a PMT or a CCD camera. Ion melting is observed by a drop in fluorescence count rate of the PMT and a lack of spatial resolution of the two ions on the camera. After melting, the fluorescence count rate is used as an indicator of ion energy. The cooling pulse duration is kept short enough to avoid a significant change in ion energy during the detection period. Short detection times, however, result in a lower signal-to-noise ratio in the fluorescence count rate. The fluorescence rate of an ion cloud in our experiment is on the order of 10^3 counts per second. Thus, on average, less than 1 count is detected in the 500 μ s detection time. To obtain an acceptable signal-to-noise ratio, an average count rate is taken from ~ 1000 repetitions of the sequence. All sequences are concluded with a 20 ms high-power (> 100 MHz), far-detuned (> 120 MHz) Doppler cooling beam, which ensures that ions are crystallized for following sequences.

Figure 4.11(a)(i) displays an example of measured fluorescence rate, as function of pulse duration of the tickle field. The drop in fluorescence rate is an indication that the ions gain energy. Figure 4.11(b)

³⁰ Back-of-the-envelope calculation shows that $\vec{k} \cdot \vec{v} < \Gamma/20$ at the Doppler limit

displays corresponding CCD images. In these images, the horizontal axis is the axial direction and the vertical axis is a radial direction, parallel to the trap surface. With increased tickle pulse duration, the ion images become increasingly broader as a result of their increased motional amplitude. The fourth image in Figure 4.11(b), taken after 100 µs, shows a single blur of fluorescence, indicating that ions have melted.

After melting, the cloud is allowed to evolve for a fixed duration, without applying a cooling beam. During this time, ions undergo RF heating. After this time, the fluorescence is probed with the PMT and camera. An example of such a measurement is shown in Figure 4.11(a)(ii). The fluorescence count rate decreases over time, which, as before, indicates an increase in ion energy. Accordingly, the cloud size, seen in the CCD images in Figure 4.11(b), increases as function of wait time.

To extend the usefulness of the measured fluorescence rate beyond being just an "indication" that RF heating is taking place, the rate is to be quantified in terms of ion cloud energy. This requires a mapping of detected fluorescence to energy. We produce this mapping through a set of simulations that estimate fluorescence rate as function of ion energy. The simulation works as follows:

For a given total secular energy $E_{\rm sec}$, a random set of motion parameters $a_{i,k}$ is produced, constrained by Equation 4.51. From this set of parameters, ion velocities $v_{i,k}(t)$ can be estimated using Eq. 4.67, which are in turn used to find the Doppler cooling force $\vec{F}_i^{(D)}$, using Eq. 4.24. An average over time of the Doppler cooling force is calculated. This average force, $\langle \vec{F}_i^{(D)} \rangle_t$ is directly linked to the amount of 'kicks' received by an ion due to absorption and spontaneous emission, and is therefore proportional to the fluorescence rate. However, this rate depends on the randomly chosen values of $a_{i,k}$. This process is therefore repeated 20 times for a fixed $E_{\rm sec}$, each with random value of $a_{i,k}$, and each producing an average force $\langle \vec{F}_i^{(D)} \rangle_t$. From this set, we take an average of average forces, $f = \langle \langle \vec{F}_i^{(D)} \rangle_t \rangle_a$, which is proportional to the fluorescence rate, for a given value of $E_{\rm sec}$.

This procedure of generating f is repeated for values of $E_{\rm sec}$ ranging from 0 to 30 meV. We thus have a value that is proportional to a fluorescence rate as function of ion energy, $f(E_{\rm sec})$. This result is scaled such that $f(E_{\rm sec}=0)=1$. We assume that in our experiment, the fluorescence rate from a crystallized pair of micromotion-compensated ions corresponds to $f(E_{\rm sec}=0)$, where for this particular fluorescence rate calibration, the previously mentioned bias field is disabled. The scaled measured fluorescence rate can therefore be mapped directly to ion energy. An example of such a mapping is shown in Figure 4.11(c). The mapping of fluorescence rate to energy is, however, not unique for the entire domain: values of $f\approx 1$ can be produced by two distinct energies, as can be seen on the left side of the plot. For our experimental settings, energies associated with melted ions are outside of this range of ambiguity.

4.7.2 Experimental RF heating results

Figure 4.12 shows the ion cloud energy, inferred from measurements, as function of wait time, for various radial motional frequencies. These frequencies are adjusted by changing the amplitude of the RF voltage. It is difficult to adjust tickle frequency and duration in such a way to precisely choose an initial cloud energy. Therefore, the initial energies for the various parameters are not necessarily equal. Despite not having perfect control over the initial energy, the limited spread and drift in cloud fluorescence rate of multiple repetitions of the experiment suggests that the initial energy is at least stable and reproducible.

The lines represent the lower and upper boundary of the standard deviation of multiple simulation runs, using the simplified RF heating model presented in Section 4.6.2. The displayed error bars of the measured data are the uncertainty in the mean of the ~ 1000 measurement repetitions. Systematic uncertainty, likely dominated by errors in the model that maps fluorescence to energy, is not included.

The initial energies used in the simulations are set to the measured initial energies. The motional frequencies used in the simulation are the same as experimental values. Measured and simulated data are in agreement for both the time-evolution of energy and the motional frequency dependence. However, the measured data at the highest radial frequencies noticeably deviates from the simulated curve. This deviation is attributed to a systematic error in determining energy from fluorescence: our fluorescence-to-energy mapping is not well-defined, meaning that a small uncertainty in fluorescence




Figure 4.11: Overview of fluorescence measurements in ion melting experiments, detected by (a) a PMT detector, and (b) a CCD camera. (i) Ions are driven to melt by tickling. (ii) While melted, ions undergo RF heating. (iii) Ions can be recooled and recrystallized by Doppler cooling. (c) Numerical simulations allow us to map measured fluorescence to ion cloud energy.

could represent a large uncertainty in calculated ion cloud energy. In particular, note in Figure 4.11(c) that at higher energies this mapping becomes increasingly more poorly defined, which could explain the larger deviation from simulated results at higher energies in 4.12. Despite this deviation, by estimating the diffusion constant from experimental data ($E \sim \sqrt{Dt}$, see Section 4.6.2.3), it is still possible to conclude that higher RF power, thus higher radial motional frequencies, lead to an increased heating rate³¹.

4.7.3 Experimental recrystallization results

The increasing ion energy due to RF heating can be opposed by applying Doppler cooling. We experimentally investigate the cooling efficiency of a Doppler beam with various parameters, which allows us to confirm the recrystallization trends shown in Figure 4.10.

In the experiment, following a 20 ms waiting time during which an ion cloud gains energy, a Doppler cooling beam is applied. Variable parameters are the beam's coupling strength Ω , detuning δ , and the cooling duration. As before, the energy of the cloud (or potentially crystal, if cooling has been efficient enough) is probed by applying a short detection pulse. In our experiment, the hardware is limited such that the detection pulse has the same coupling strength and detuning as the cooling pulse. The measured fluorescence rate is therefore not independent of the settings of the cooling pulse, which restricts the accuracy of the energy estimate. Furthermore, for higher beam detuning, the fluorescence rate is less dependent on ion energy, which makes estimating it less precise. We therefore elect not to use the fluorescence rate as an indicator for ion energy, but instead use the measured cloud size from fluorescence collected on the CCD camera.

The size of the ion cloud on the CCD image could be translated into ion energy using a similar method as before, when we mapped fluorescence rate to ion energy. However, since the perceived cloud size depends on Doppler beam detunings and strengths, a unique mapping and calibration would be required for each beam setting. We wish to avoid such a large look-up-table with yet another




Figure 4.12: RF heating experimental results, with a Ca - Ca cloud. Energies have been measured for three sets of radial motional frequencies. Error bars represent the statistical spread in the acquired data. The thick lines are an average of simulated energies, and the thin lines are the standard deviation.

possible source of systematic error. Therefore, instead of displaying data with an absolute scale of ion energy or absolute cloud size, the results are presented as values of *relative* cloud size, scaled from the cloud size at the moment Doppler cooling is initiated to the size of an ion crystal. In this simplified presentation of the data, without attributing an exact cloud energy, we still determine to which degree the beam parameters have reduced the cloud size, and potentially recrystallized, and thus provides an indication of how well cooling has worked.

The measured ion cloud size after Doppler cooling is displayed in Figure 4.13, for various beam coupling strengths $\Omega/(2\pi)=\{44,\,58,\,80\}$, detunings $\delta/(2\pi)=$ -20 to -300 MHz, and pulse durations up to 10 ms. The results suggest that for coupling strengths below 80 MHz, recrystallization is only achieved after 10 ms if the Doppler beam is detuned by -100 to -200 MHz. This agrees well with the conclusion drawn from the recooling simulations shown in Figure 4.13.




Figure 4.13: Cloud size after applying a cooling beam to an ion cloud for various times, for experimental measurements and simulations. In the experiment, cloud size is determined from CCD imaged fluorescence. We vary the detuning and coupling strength of the Doppler beam. Colors are scaled with the mean cloud size at 0 ms and crystal size as reference.

We directly compare the measured data with the simplified RF heating simulation (see Section 4.6.2). The cloud size is reconstructed from the simulations by correlating the ions' position and

fluorescence rates at the specific ion motion parameters $a_{i,k}$ within the respective detection window. The simulation uses the same cooling beam detunings, pulse durations, and trap parameters, as in the experiment. The coupling strength, however, is varied from 10 to 120 MHz. From this set, we calculate which value of simulation coupling strength has the best agreement with the experimental data, in terms of least-squares difference. The simulations with the best agreement are displayed alongside the respective experimental results in Figure 4.13.

The value for coupling strength at which this agreement occurs differs by about a factor of two. This discrepancy is attributed to model imperfections that arise from simplifying the Doppler cooling process in our simulations. These imperfections affect the simulated results, but also the way that beam parameters are calibrated in the experiment. For example, in this calibration, experimental data is fit to a model that assumes a spontaneous decay rate of $\Gamma/(2\pi)=21.6\,\mathrm{MHz}$, which neglects possible decay to the $3D_{3/2}$ level. The modeled value of spontaneous decay Γ is thus an upper bound for the effective two-level spontaneous decay. A possible discrepancy in Γ leads to an incorrect determination of experimental values of Ω .

Barring this discrepancy, the simulated and measured data have a good qualitative agreement, reaffirming the conclusion about the most efficient recooling settings: ensuring recrystallization of a Ca - Ca cloud in 5 ms requires a Doppler cooling beam with $\Omega/(2\pi) > 80\,\mathrm{MHz}$ and $\delta/(2\pi) \approx 150\,\mathrm{MHz}$. Recrystallization is delayed or unattainable with a lower beam power and/or incorrect detuning. For example, typical Doppler cooling parameters for reaching the Doppler limit, $\Omega \ll \Gamma$ and $\Delta = \Gamma/2$, deviate by more than an order of magnitude from the optimal recrystallization settings.

The experimental results, and their comparison to simulation data, are an indication of the magnitudes of the scales of relevant parameters involved in RF heating: melted ions gain energy on the order of *electron volts per millisecond*. Ion clouds can be efficiently recrystallized in *several milliseconds*, if the recooling beam has a coupling strength and detuning that are both several times larger than the cooling transition's spontaneous decay rate. These conclusions are supported by the full ion dynamics simulation, the simplified simulation, and the experimental results.

The conclusions are intentionally kept general, avoiding attributing specific values. For an experimental physicist working with trapped ions, it is important to be aware of the orders of magnitude of RF heating, not specific values. The results represent *average* heating and cooling rates. As has been shown with the full ion dynamics simulation, the energy development of an energy cloud changes notably from one cloud to the next. Owing to this chaotic variability, one could never accurately predict what the exact energy development after a melting event in one's experiment will be. It is therefore much more constructive to have a general understanding of what average rates can be expected from a particular set of trap parameters³².

4.8 OUTLOOK: FURTHER CONSIDERATIONS OF RF HEATING

The majority of this chapter has covered simulations of RF heating for the case of two melted ⁴⁰Ca⁺ions, in ideal trapping conditions (i.e., no micromotion, harmonic trapping potentials, infinite trap depth, etc.). Needless to say, there are many avenues for further investigation on melted ion motion. In this section, as an outlook, we present simulations and offer brief discussions of other example scenarios that warrant further investigation.

The simulation results are summarized in Figure 4.14. As in Figure 4.4, these plots display the change in secular energy $E_{\rm sec}$ of 20 simulation runs (thin gray lines) using the full ion dynamics simulation. Each run has a random RF drive phase, which creates a marginal variation in the initial conditions. The thick dark line is the average energy of the runs. Additionally, for comparison, the energy gain of the typical Ca - Ca ion cloud settings used throughout this chapter is shown (light blue line).

-(a) Surface trap-

Realistic anharmonicity is introduced into the simulated trapping potentials by modeling the potentials of a surface trap. The simulated trap has a similar electrode layout as the Sandia HOA trap (see Section 3.2.1), though does not include the central slot. The electrode voltages have been set such




Figure 4.14: Simulated ion energy gain under various conditions. The gray lines are secular energies of individual simulation runs, though have been marked red for runs where an ion has escaped the trap. The thick dark blue line is an average of these energies. The light blue line is a reference, showing RF heating of the 'standard' settings used throughout this chapter: two ⁴⁰Ca⁺ions in a harmonic trapping potential with motional frequencies 3.3, 3.6 and 1.1 MHz, and a trap drive frequency of 35 MHz.

that the motional frequencies match those used for the reference simulation. The trap's escape point is 54 µm away from the trap center and has an energy barrier of 33 meV for the chosen voltage settings. Initially, the energy gain is similar to the reference. At low energies, oscillation amplitudes are low, and the trap potential 'looks' harmonic to the ion. With increasing ion energy, the oscillation amplitudes increase, and the anharmonic potential is no longer negligible when describing the ions' motion. In Figure 4.14(a), making the trap potential anharmonic appears to increase the rate of energy gain, compared to the ideal harmonic trap, though not significantly. This increase is counter-intuitive, as the motional frequencies decrease with increasing energy; as shown in Section 4.6.2.3, lower motional frequencies are expected to decrease the RF heating rate. However, the increase in RF heating rate is likely attributed to a higher susceptibility to *parametric resonance* as the degree of anharmonicity increases, defined and described later in this section when subplot (g) is discussed.

-(b) Multiple-ion cloud-

The unfavorable scaling of RF heating as function of ion number has been discussed in 4.6.2.3, as determined by the simplified model. Running the full simulation with more than two ions in a cloud allows us to confirms this trend. The simplified model predicts an average energy of 30 meV after 5 ms for a three-ion cloud. The *full* simulation, shown in Figure 4.14(b), matches this prediction very well. Only individual runs for the three-ion simulations are shown. Average values are shown for three to five ions.

-(c) Higher mass-

Figure 4.9 displays that higher ion masses are subject to higher RF heating rates. This was, however, demonstrated under the constraint that the trapping fields were adjusted to maintain constant motional frequencies. In Figure 4.14(c), we show how the RF heating rate is affected by a change in mass, *without* adjusting the trapping fields. Since the RF heating rate is most sensitive to the radial motional frequency (see Table 4.1), and that the frequency is inversely proportional to ion mass, it follows that a higher ion mass will result in a lower heating rate. Figure 4.14(c) confirms that the RF heating rate decreases with respect to the lower-mass reference, if the trapping potentials are kept identical.

-(d) Mixed-species-

It is not trivial to intuitively reason how the presented RF heating models extend to mixed-species clouds. The models in the simplified simulation for obtaining the collision time $t_{\rm coll}$ and collision energy ΔW rely on the symmetry of both ions' frequency and mass being identical, and are therefore not directly applicable to mixed-species³³.

As an example of RF heating of a mixed-species crystal, we run simulations of 40 Ca $^+$ and 88 Sr $^+$ clouds. One could postulate that the heating rate of Ca - Sr should lie somewhere between that of the a Ca - Ca cloud and an Sr - Sr cloud. Figure 4.14(d) shows that, indeed, the mixed species ion cloud initially performs favorably in terms of heating rate compared to the reference, though is higher than its higher-mass counterpart. However, after 5 ms the energy gain is similar to the reference. It seems that on timescales of several milliseconds, the rate of energy gain is dominated by that of the lighter species.

-(e) Mixed-species on a surface trap-

Placing the Ca - Sr cloud into a surface trap, thus exposing it to anharmonic potentials, increases the RF heating rate, as shown in Figure 4.14(e). The rate is comparable to the single species surface trap heating rate, as in (a). This does, however, present a more significant issue for trapping mixed species on surface traps: the heating rate for a Ca - Sr cloud is comparable to that of Ca - Ca, but the trap depth, (i.e., the energy an ion requires to escape the trap), is more than twice as low for the Sr ion. The higher mass of trapped mixed-species ions therefore is more likely to gain enough energy to escape the trap, compared to a single-species case. Indeed, this is confirmed by simulations, as one of the twenty runs resulted in a 88 Sr $^+$ ion escaping the trap. This scan is marked red in Figure $^{4.14}$ (e), and terminates prematurely due to ion loss.

This simulation highlights the importance of aiming to maximize the trap depth for surface traps intended to be used for mixed species operation. In typical experimental sequences, sideband cooling and state preparation already take more than 5 ms, let alone the rest of the experimental sequence

during which Doppler cooling is switched off. If one out of twenty melting events lead to ion loss within 5 ms, and background gas collisions occur a few times per minute, one could expect ion lifetimes of tens of minutes. If Doppler cooling parameters are not set properly to ensure recrystallization after each experimental cycle, this lifetime will be significantly lower (which we can qualitatively confirm from experience in the lab). This stresses the importance of ensuring a sufficiently deep trap, and proper recrystallization settings for the Doppler cooling beams.

-(f) Excess micromotion-

Excess micromotion arises when an undesired field forces ions away from the ideal center of the trap, where the RF potential curvature is zero³⁴. Ions that are on average displaced from the RF-null experience a higher average motional amplitude at the RF frequency relative to their secular amplitude. In a cloud, ion-ion collisions are expected to occur on average further away from the trap center, thus are more likely to occur in regions of higher RF field. One could thus speculate that the RF heating rate would be enhanced under the influence of excess micromotion. However, the model for energy change due to collisions in 4.6.1 allowed us to conclude that the collision energy is independent of where in the trap the collision occurs. Instead, the magnitude of energy change is dependent on the RF potential curvature, which is constant for an ideal harmonic potential, and does not change when an additional DC field is applied. Our model therefore suggests that micromotion does not influence the RF heating rate.

We have applied a DC stray field of $1\,\mathrm{V}\,\mathrm{mm}^{-1}$ in one of the radial directions. Simulated results are shown in Figure 4.14(f). This field induces a micromotion modulation index of $\beta\approx0.4$, far above a desirable experimental value. The resulting energy trend is very comparable to the micromotion free reference, even perceptibly marginally lower. Excess micromotion appears to play only a minor role in RF heating dynamics.

-(g),(h) Resonance-

It is well-known that applying an additional RF field near or at one of the secular frequencies of trapped ions will add energy to their secular motion (This is precisely the 'tickling' method used to melt ions in Section 4.7). Applying an RF field with a frequency f that is a *multiple* of the secular frequency, $f = n_k \omega_k$, similarly results in motional excitation. Here, ω_k is the secular frequency, and n_k an integer. The condition for resonant excitation of the ion's motion extends to specific combinations of multiples of the secular frequency and the trap drive frequency $\Omega_{\rm RF}$, as $f = n_{\rm RF}\Omega_{\rm RF} + n_k \omega_k$, with $n_{\rm RF}$ another integer [167, 169–171].

Parametric resonance might occur at values of Ω_{RF} and ω_k where the condition $f=0=n_{RF}\Omega_{RF}+n_k\omega_k$ can be realized for certain integer values of n_{RF} and n_k . This effectively means that resonant excitation can occur without the application of an external RF field, which is why parametric resonance is sometimes referred to as 'self-resonance'. In ideal harmonic traps, operating within the usual Mathieu stability regime in the RF-null, parametric resonance does not pose a problem: higher order Fourier components of secular motion vanish, and thus do not interfere with multiples of the trap drive frequency. In anharmonic traps, however, these higher order terms are not negligible, especially with the large amplitudes associated with ion clouds. The Mathieu stability diagram is in this case dissected with regions of instability [167, 172–174].

In Figure 4.14(g), the DC confinement has been increased with respect to the reference, such that the axial motional frequency is changed from 1.1 to 1.72 MHz. The corresponding increased anti-confinement in the radial direction brings one of the radial motional frequencies, from 3.64 to 3.5 MHz, which is an exact factor of ten lower than the trap drive frequency. There is no perceptible change in the average RF heating rate.

The simulations shown in plot 4.14(h) have the same trap frequencies as in 4.14(g), but use the potentials of a surface trap, which include anharmonicity. The average increase in secular energy is higher than the reference, despite the lower radial motional frequency. Though the difference in RF heating rate does not seem prominent, one should note the displayed energy trend is an average of the energies of ions that have not escaped from the trap, which by the end of the scan is ten out of the initial twenty. The displayed RF heating rate may therefore seem misleadingly low, as it exhibits a

strong case of survivor bias. The individual energy traces of ions that escape the trap within 5 ms are highlighted red in 4.14(h).

In summary, an unfortunate combination of trap settings can cause ions' motion to become unstable. While this instability also occurs with just a single trapped ion, with multiple ions the combination of parametric resonance instability and RF heating accelerates the increase in energy, and is detrimental to the ions' lifetime. Scans with identical trap parameters, but only a single ion (not shown), do result in non-constant secular energy, but not in ion expulsion from the trap.

4.9 CONCLUSIONS

In this chapter, we have considered the properties of RF heating of ion clouds in Paul traps. We have analyzed simulations, produced heating models, and witnessed RF heating in our experimental setup.

The results convey the necessity of having a recrystallization setting in experimental hardware in addition to the typical trapping and Doppler cooling settings, which are optimized for achieving a low ion energy, but not for rapidly cooling down excessively energetic ions, such as those in an ion cloud. Recrystallization strategies include lowering the RF drive power and applying a far-detuned, high-power Doppler cooling beam. Although these strategies are conceptually known as a rule-of-thumb, this work presents an indication of the orders-of-magnitude involved in RF heating and Doppler cooling, and gives an intuitive understanding of the underlying processes.

Efficiently overcoming RF heating is imperative in ion trap systems with low ion escape barriers, such as surface traps. Realizing scalability in ion-trap-based quantum computers will likely require segmented surface trapping architectures [24]. With increasing numbers of ions, collisions with background particles become more frequent, and therefore also the number of melting events. Even if the energy transferred in such collisions is lower than the trap depth of surface traps (typically tens or hundreds of meV), energy gain from RF heating can lead to loss of ions from the trap, possibly in tens of milliseconds. Therefore, to avoid persistently reloading ions, experimental sequences should include a refreeze phase in every cycle. Our results suggest that for recrystallization of a melted ion crystal, the Doppler cooling beam should have a detuning of roughly $\delta \approx -6\Gamma$ and should be applied for more than 5 ms, with at least $\Omega > 3\Gamma$ coupling strength. Lowering the power of the RF drive field during this refreeze phase will aid recrystallization by decreasing RF heating. At a lower RF power, the benefit of a lower heating rate outweighs the disadvantage of the resulting lower trap depth.

The considerations of trap and cooling parameters become more stringent when operating with multiple ions of mixed species [65], which is, in addition to the use of surface traps, another expected requirement for scalable ion-based quantum computers [175]. For single-species clouds in fixed trapping fields, RF heating rates decrease as the ions' mass increases. This can be seen in Table 4.1, noting that for fixed trap parameters, $m \propto \omega_r^{-1}$. On the other hand, the trap depth E_d , usually lowest in the radial direction perpendicular to the trap surface, is approximately proportionally lower for higher masses, $E_d \propto m^{-1}$. Simulations show that mixed-species crystals suffer from a worst-of-both-worlds: RF heating rates are dominated by the lower mass ion in the cloud, while the trap depth remains low for the higher mass ion. Additionally, one must take extra care in reducing the amplitude of the RF trap drive voltage as a strategy to lower RF heating, since a decrease of RF power also leads to a decrease in trap depth. The higher-mass ion is then even more prone to escaping the trap. It is therefore beneficial to operate ion traps with species of similar mass. However, regardless of the mixed-species mass ratio, efficient recrystallization is imperative.

The existence of RF heating in Paul traps has been recognized as long as the existence of Paul traps [56]. However, due to its chaotic nature, attempts to model and estimate RF heating rates for a few ions has been sparse. This work presents a first effort into producing simple models that describe RF heating of low numbers of trapped ions. We have analyzed a dynamic chaotic system and developed a simplified model to characterize it. The RF heating model can be further extended to include effects of multi-species operation, larger numbers of ions, excess micromotion, and anharmonicities in trapping potentials. Furthermore, recrystallization with Doppler cooling can be further investigated, accounting for the effective dynamics of the eight-level cooling scheme. A Doppler beam with chirped detuning [176] can, for example, be an effective method to ensure recrystallization.

TRANSPORT OPERATIONS OF TRAPPED-ION QUBITS

It's simple, so says the captain / Face forward, move slow, forge ahead.

— Guster, The Captain

5.1 INTRODUCTION

One of the current challenges in trapped-ion quantum computation is scaling up the number of ionic qubits. That is not to say that it is difficult to store many ions simultaneously: Paul traps can store hundreds of thousands, if not millions [144], of ions in one trapping region. However, in such collections, ions cannot individually be employed as qubits, as they cannot be individually addressed and measured, lack a faithful channel of communication, and cannot maintain coherence at useful timescales. Even collections of tens of well-localized (crystallized) trapped ions in the same trapping potential present technical challenges for quantum computation, for, among others, the following reasons:

- The number of motional modes increases proportionally with the number of ions *N*. The high density of modes in frequency space, termed 'mode-crowding,' impedes the ability to coherently manipulate individual motional modes, which serve as the medium through which entanglement between qubits is generated.
- Coupling to motional modes is weaker with increasing numbers of ions, which increases the time it takes to perform gates that rely on the ions' motion (see, for example, the MS gate, Section 3.2.8). Gate times increase proportional to \sqrt{N} . Extended gate times can reduce the fidelity of quantum operations, if they cannot be performed much faster than timescales associated with qubit decoherence.
- With an increasing number of ions, maintaining a linear chain of ions along the RF-null requires a increasingly low ratio of axial to radial confinement. Since RF power and Mathieu stability limit how much the radial confinement can be increased, attaining longer ion chains requires that axial confinement must be decreased. Therefore, axial motional frequencies are reduced as *N* increases, which comes at the cost of higher sensitivity to external noise sources [94].
- Larger ion crystals are more susceptible to collisions with the background gas. Although a collision with a background-gas particle occurs with only a single ion, it affects the motional state of the entire ion crystal [140].

All sources of error considered, storage of more than one-hundred ions in one trapping region will likely not allow for operations with high enough fidelities to be practical for the purpose of quantum computation [25].

These obstacles are mitigated (partially, at least) if the collection of qubits can be isolated in multiple independent potentials, thus forming separated ion crystals. Scalability in quantum information processing using trapped ions will therefore inevitably require arrays of separated trapping regions [177]. Among the proposed designs to realize scalable quantum computation is the QCCD [24, 27]. Such a device consists of a set of voltage-carrying electrodes laid out in a planar architecture. A set of electrode rails carries an RF voltage that provides radial confinement around a one dimensional channel (or interconnected set of channels), the RF-null. A set of segmented electrodes carrying DC voltages allow multiple trapping regions to be generated along the RF-null channel. The chip can then support regions dedicated to various tasks, such as memory storage, qubit interaction, recooling, and loading.

A cornerstone of the functionality of a segmented trap for quantum computation is the ability to manipulate the configuration of the ion crystals, in other words to control the position of ions. These manipulations can broadly be subdivided into three types of operations:

- Shuttling operations are required to move an ion crystal between various dedicated trap regions.
- Merging and splitting operations join and separate ion strings, which allows for coupling and decoupling between separate ion registers.
- **Rotation** operations are required to allow ions traveling along an RF-null to pass each other without interference, enabling enhanced connectivity of separated ion registers.




Figure 5.1: Schematic overview of ion crystal transport operations.

These operations are schematically shown in Figure 5.1.

These transport operations are enabled by altering the trap's electrostatic potentials, achieved by sequentially adjusting voltages of the trap electrodes. As is the case for all quantum operations, physical ion manipulation operations should be carried out on timescales much shorter than qubit coherence times. On the other hand, for high fidelity gate operations, the ions' motional state must remain near the ground state, and therefore must be minimally affected by the rapidly changing electric fields that enable the transport. Since rapid changes in electric fields can excite the ion's motion, there is an effective upper bound on how fast operations can be carried out. Additionally, transport speeds are limited by the rate at which electrode voltages can be changed, which in turn is limited by an experimental setup's electrical filters and the voltage signal generator. These filters are imperative in ion trap setups, in particular those employing surface traps, to reduce electronic noise at frequencies corresponding to the ions' secular motion.

Precise manipulation of an ion crystal configuration requires a pre-determined sequence of voltages to be applied to multiple electrodes of a segmented trap. Instead of characterizing transport operations in terms of a list of voltages, a clearer representation is to parametrize the operation in a set of 'control knobs' that quantify an intuitive property of the ion configuration throughout a transport sequence. For example, ion shuttling can be described as a one-dimensional control parameter that defines where along the trap axis the axial potential minimum is placed. This control parameter is linked to a set of electrode voltages (through a mathematical model or through a look-up-table). This link can, for example, be a linear interpolation of two known voltage sets for two trapping sites at course grained positions [124]1. Of course such a description of voltage sequences oversimplifies actual experimental conditions, since stray fields and trap manufacturing tolerances limit the precision of this operation. Transport sequences require careful calibration to overcome errors stemming from these limitations. Also, appropriate voltage ramp rates need to be determined, to ensure ion acceleration and deceleration do not induce motional excitation. This can be prevented by transporting ions adiabatically [62], in which the timescale of ion movement is much larger than the secular motion period. Faster transport operations [177, 178] can be realized, but typically require excessively careful calibration, and are much more sensitive to changes in field potentials over time, compared to adiabatic transport.

¹ While generally valid for three-dimensional traps, we see later that linear interpolation of voltage sets for two-dimensional traps requires more intermediate voltage solutions than just those at neighboring electrodes

Ions are generally required to remain close to the RF-null throughout a transport operation. Electric potentials that enable transport must be carefully simulated and calibrated to adhere to this requirement. This makes physical ion manipulation on segmented *surface* traps technically more challenging than on their 3D counterparts. One dimension less of symmetry is one dimension extra of headache: Typical 3D traps have electrodes placed symmetrically around the RF-null. By ensuring that opposing electrodes undergo symmetric voltage changes during a transport sequence, ion positions are kept near the RF-null, even when the applied voltages waveforms are distorted due to electrical filters. In contrast, to keep ions near the RF-null in a surface trap, in particular perpendicular to the trap surface, there are no symmetry shortcuts and the electrode potentials must be well understood.

Further complications arise when considering transport of mixed-species ion crystals. In ion shuttling, for example, dynamic fields result in unequal accelerations of ions of different masses. Therefore, to avoid motional heating, shuttling at useful speeds requires careful tailoring of voltage ramps [179]. Ion crystal rotations with mixed species are also more complicated than single-species operation, since ions of unequal mass are forced away from the RF-null at unequal magnitudes. As we will see in Section 5.4.4, this results in much narrower tolerances in applied voltages for successful rotations.

This chapter details practical considerations in the design of voltage sequences for various ion transport operations in surface ion traps. Presented values and results are specific to the linear surface traps used throughout this thesis, but the prescribed methods for determining voltage sequences are universal. Sections 5.2 and 5.3 discuss strategies for generating shuttling and splitting voltage sequences for surface traps. Section 5.4 discusses ion crystal rotations, which is adapted from Ref. [88].

5.2 SHUTTLING

Ion shuttling is the action of displacing an ion along the 1D channel defined by the trap's RF-null. Shuttling allows ions to be moved to various regions of the trap. In the context of the QCCD architecture for quantum computing, such trap regions can each have unique functions, such as ion loading, storage, manipulation and readout. Additionally, certain regions on the trap may be designed specifically to enable other types of physical ion manipulation, such as splitting (Section 5.3) and crystal rotations (Section 5.4). Ion shuttling is therefore an essential component for scalable quantum information processing. Shuttling has been studied extensively in various trap architectures [27, 124, 178, 180–182], and is routinely employed in many quantum algorithms [180, 183–185]. This section provides general guidelines for designing voltage sequences for ion shuttling.

Following the conventions introduced in Section 2.1, we consider a linear RF trap, with an axial direction z, defined by axis along which the trap's RF-null extends. The remaining axes, x and y, are radial directions. For now, we ignore the specific layout of the trap electrodes and their voltages, and focus only on the requirements of the electric potential V, and how it is to be controlled in order to facilitate ion shuttling. In particular, the ion position is displaced by adjusting where along the trap axis the minimum of the axial confining potential resides. The axial component of the trap potential V_z , assumed to be harmonic, during a transport sequence is given by:

$$V_z(z,t) = \frac{1}{2}\psi_z(t) (z - z_p(t))^2,$$
 (5.1)

where a possible constant offset in the potential is neglected. The parameter $z_p(t)$ is the location of the axial potential minimum. The field curvature $\psi_z(t)$ is often expressed in terms of ion mass m, charge q and axial motional frequency ω_z , as $\psi_z(t) = mq^{-1}\omega_z(t)^2$. It is desirable to maintain constant mode frequencies during the transport, $\omega_z(t) = \omega_z(0)$ for all t, as rapidly changing frequencies are a source of ion heating [186]. Ion shuttling involves moving the potential minimum, from an initial position $z_p(0) \equiv z_0$ to a final position $z_p(t_f) \equiv z_f$, with duration t_f . As with all qubit operations, core requirements of the shuttling operation are that the duration is sufficiently short to avoid loss of coherence because of lengthy computation sequences, and that the number of motional quanta added to the ion(s) is kept to a minimum. The total transfer duration t_f and the shape of the signal $z_p(t)$ for $0 < t < t_f$ must be chosen in such a way that these requirements are met, as is discussed in the

following section. Afterwards, when a suitable potential ramp $z_p(t)$ has been designed, the desired potential sequence $V_z(z,t)$ needs to be mapped to electrode voltages, V_e .

5.2.1 Potential ramp design

The design of $z_p(t)$ has been studied in detail in [186] and [187], analyzing both classical and quantum dynamics to estimate the ion's energy change due to a transport operation. It has further been studied in [181], where transport operations are optimized for speed in the presence of electric field noise. However, in practice, numerical minimization of energy gain and transport duration in ideal trapping conditions are overshadowed by experimental limitations, such as stray fields, undetermined higherorder fields [188], and precision of applied voltages. While it has been demonstrated that shuttling operations can be achieved on timescales of a few motional oscillation periods [178] (i.e. several microseconds), the strict demands that are placed on voltage calibration and sensitivity to changes in stray fields have limited the use of fast transport as a component of larger quantum information sequences. Additionally, fast switching of electrode voltages must be allowed [188], and is made technically challenging when dealing with electrical filters of segment voltages². Instead, adiabatic ion transport, comparatively slow, is more commonly utilized. Here, typical shuttling velocities may reach some meters per second, which equates to tens to hundreds of microseconds of transport time per electrode segment. Indeed, in practice, shuttling remains one of the largest time overheads for many quantum computation sequences [27]. In the following, we will discuss this safer and currently more practical adiabatic approach to ion shuttling.

A naive approach to simple adiabatic transport is to move an ion from z_0 to z_f at a constant rate, given by $z_p(t) = (z_f/t_f)t$, setting the z-axis such that $z_0 = 0$, for convenience. This trajectory and corresponding acceleration is shown in Figure 5.2(a), in the left panels. Solving the classical equations of motion for an ion in a potential given by Eq. 5.1,

$$\ddot{z}(t) = -\frac{1}{m} \frac{\partial V_z(z, t)}{\partial z} = -\omega_z^2 \left(z(t) - z_p(t) \right) \tag{5.2}$$

with

$$z_p(t) = z_f \frac{t}{t_f}, \quad z(0) = 0, \quad \dot{z}(0) = 0$$
 (5.3)

results in an ion position during transport given by

$$z(t) = z_f \left(\frac{t}{t_f} - \frac{\sin(\omega_z t)}{\omega_z t_f} \right). \tag{5.4}$$

The final energy of the ion after the transport operation is given by the sum of its kinetic energy and potential energy. Assuming that the position of the potential minimum is constant at the end of the transport, $z_p(t) = z_f$ for $t \le t_f$, the added energy E_f is given by

$$E_f = \frac{1}{2}m\omega_z^2(z(t_f) - z_f)^2 + \frac{1}{2}m(\dot{z}(t_f))^2$$
(5.5)

$$=\frac{z_f^2 m}{t_f^2} \left(1 - \cos(\omega_z t_f)\right) \tag{5.6}$$

and is shown in the right panel of Figure 5.2(a), using parameters $m=40\,\mathrm{u}~(^{40}\mathrm{Ca}^+)$, $\omega_z/(2\pi)=1\,\mathrm{MHz}$, and $z_f=200\,\mathrm{\mu m}$. Energy is expressed in terms of phonon number, $E_f/(\hbar\omega_z)$ to put the energy gain in a relevant context, but note that an exact description of phonon occupation requires a quantum mechanical analysis of ion transport.

The dips in the right plot of Figure 5.2(a) correspond to times at which the deceleration at the end of the transport results in an energy loss that exactly balances out the gain in energy after the

² Which in our case are in-vacuum and in-cryo, thus difficult to work around.




Figure 5.2: Overview of shuttling strategies. Left panels show the set position and acceleration of the minimum of the axial trap potential. The energy gained by an ion in this potential depends on the total duration of the transport operation t_f , shown in the right panels, where the following parameters are used: $z_f = 200 \, \mu\text{m}$, $m = 40 \, \text{u}$, $\omega_z = 2\pi \, 1 \, \text{MHz}$

initial acceleration. With careful calibration and timing, these dips³ can be utilized to enable low duration transport with low added energy [188, 189], and is known as *bang-bang* transport. This type of transport is extremely sensitive to drifts in electric potentials, for example caused by changes in stray charges. The stringent requirement of repeated recalibration makes use of bang-bang transport within QC sequences impractical, and is therefore up until now rarely applied.

Assuming the worst case scenario, Eq. 5.6 suggests that if less than a phonon of heating is required, transport times are to be on the order of $t_f \approx \sqrt{z_f^2 m/(\hbar \omega_z)} = 2\,\mathrm{ms}$, which exceeds acceptable times. The time at which the < 1 phonon goal is achieved is marked with a dashed line in Figure 5.2. This is provided as an indication of the general performance, but we note that a realistic experimental goal might be considerably lower. Applying this intended potential sequence to an actual trap will likely result in less added energy than plotted in Figure 5.2(a), since electrical filtering of voltage ramps will smooth out the initial acceleration and final deceleration⁴. Regardless, a more suitable transport strategy is desired.

We can smooth out the acceleration and deceleration purposefully by ramping up to a constant transport velocity, and ramping back down before the destination z_f is reached. Using a quadratic ramp, the potential minimum follows:

$$z_{p}(t) = \begin{cases} \frac{1}{2}at^{2}, & 0 \leq t < t_{i} \\ at_{i}\left(t - \frac{1}{2}t_{i}\right), & t_{i} \leq t < t_{f} - t_{i} \\ -\frac{1}{2}a(t - t_{f})^{2} + z_{f}, & t_{f} - t_{i} \leq t \leq t_{f} \end{cases}$$
(5.7)

with $a = z_f/(t_i(t_f - t_i))$, and t_i the duration of the ramp. The position $z_p(t)$ and its corresponding acceleration are shown in Figure 5.2(b) (left). As before, the classical equations of motion, Eq. 5.2, can be solved with this polynomial ramp, though is best left outsourced to *Mathematica*. The energy of the ion after shuttling with a polynomial ramp, calculated using Eq. 5.5, is approximately given by

$$E_f \approx \frac{3z_f^2 m}{(t_f - t_i)^2 t_i^2 \hbar \omega_z^3},\tag{5.8}$$

which boasts a much lower energy, and therefore better performance, compared to a constant transport. The denominator in Eq. 5.8 is maximized when $t_i = t_f/2$, suggesting that the best performance is reached when the constant transport (second line in Eq. 5.7) is eliminated entirely. In this case, the final energy is given by

$$E_f = \frac{128z_f^2 m}{t_f^4 \hbar \omega_z^3} \sin^4 \left(\frac{\omega_z t_f}{4}\right) \tag{5.9}$$

and is shown in Figure 5.2(b) (right). The time in which the one-phonon goal is realized has been reduced to 60 µs, nearly two orders of magnitude, by smoothing out the transport.

Further improvements can be made by including even higher-order smoothing to the acceleration and deceleration of $z_p(t)$, as is done in Refs. [181, 189]. There, one sees a resemblance between the acceleration profile and a sine curve, from which we pose that a simplification of the higher-order polynomial representation of $z_p(t)$ can be derived from a sinusoidal acceleration $\ddot{z}_p(t) = c \sin(2\pi t/t_f)$, with a constant c. Using the boundary conditions $z_p(0) = 0$, $\dot{z}_p(0) = 0$, and $z_p(t_f) = z_f$, we find that the potential displacement follows

$$z_p(t) = -\frac{z_f}{2\pi} \sin\left(\frac{2\pi t}{t_f}\right) + z_f \frac{t}{t_f}.$$
 (5.10)

 $_{\rm 3}\,$ Dips do not reach zero in the figure due to finite numerical precision.

⁴ This constitutes a rare case where experiments outperform theory.

Another round of application of Eqs. 5.2 and 5.5 in Mathematica produces a final ion energy given by

$$E_f = \frac{16z_f^2 m \pi^4}{(t_f^3 \omega_z^2 - 4\pi^2 t_f)^2 \hbar \omega_z} \left(1 - \cos(\omega_z t_f) \right), \tag{5.11}$$

shown in Figure 5.2(c). Using a sinusoidal acceleration brings the theoretical transport time down to $14\,\mu s$ to traverse a distance of $200\,\mu m$, a more-than three orders of magnitude improvement over a constant ramp. While higher-order polynomial models theoretically outperform the sinusoidal approach, the latter is numerically simpler and more intuitive. The following section describes how the desired potential sequence for ion shuttling is translated to a voltage solution in our experimental setup.

5.2.2 *Voltage solution for shuttling potentials*

Section 2.3.1 has shown how a desired DC trapping potential is expressed as a linear combination of a basis of spherical harmonic potentials, as $V_{DC} = \sum_{l,n} v_{l,n} Y_{l,n}$. A transformation matrix **V** maps the desired potentials $v_{l,n}$ to a set of electrode voltages V_e .

The potential basis $Y_{l,n}$ is defined around a central point x = y = z = 0. Thus, the primary confining potential term $Y_{2,0} = 2z^2 - x^2 - y^2$ is a harmonic potential that has a minimum in the axial (z) direction at z = 0. Ion shuttling demands the displacement of this axial minimum. One method to shift the ion position along the axial direction is to add a homogeneous electric field, given by the term $Y_{1,0} = z$. The axial potential minimum is then $z = -\nu_{1,0}/(4\nu_{2,0})$. This is, however, an ineffective strategy: The potential-to-electrode mapping V is a solution for potentials in a volume of $20 \times 20 \times 20 \ \mu m^3$ around the trapping center, and leaves the potentials outside of this range undetermined. Shuttling, on the other hand, requires hundreds of micrometers to be traversed. One could then produce a solution V for a volume of, for example, $20 \times 20 \times 200 \,\mu\text{m}^3$. While such solutions are technically possible, it is unnecessary to produce a set of voltages that generates a well-defined potential spanning several electrode widths, far away from the ion. Instead, at any given point in time, we only require a voltage solution that produces a harmonic potential locally around the current intended minimum. We therefore instead pre-compute separate multipole-to-voltage solutions $V(z_p)$, for multiple axial positions $z_p \in \left[0, z_f\right]$, where each solution uses simulated potential data in a $20 \times 20 \times 20 \ \mu\text{m}^3$ volume, centered around x = y = 0 and $z = z_p$. The applied voltages, as function of intended ion position z_p , are then $V_e(z_p) = \mathbf{V}(z_p)\vec{v}(z_p)$. Here, $\vec{v}(z_p)$ is a vector containing the terms $v_{l,n}$, which in our case are limited to l = 1, 2.

Ideally, $\vec{v}(z_p)$ is not a function of position z_p , as trapping fields around the ions should not change while shuttling. In practice, however, stray charges and inconsistencies between simulated and actual potentials produce unpredictable deviations in the potential. The motional frequencies are dominated by the field potential produced by voltages applied to electrodes, and are expected to deviate only minutely due to stray charges. Excursions of the ion away from the RF-null, however, can be excessive, and can be more detrimental to an ion's motional state [190], and therefore its coherence. Such excursions can be compensated by setting the multipole terms \vec{v} to be a function of position z_p , in particular the components of $\vec{v}_{1,-1}$ and $\vec{v}_{1,1}$. One can experimentally determine for multiple points along the transport path which values of \vec{v} should be applied such that micromotion is compensated (see Section 2.3.3). \vec{v} can then be interpolated over the full range between $z_p = 0$ and $z_p = z_f$.

Figure 5.3 shows electrode voltages required to place the axial potential minimum at z_p , shown for a range of $z_0 = 0$ to $z_p = z_f = 110 \,\mu\text{m}$, using the Golden Gate trap as an example. These voltage curves represent the waveform that transports an ion from one segment to a neighboring one. For visual clarity, voltages of only 5 nearby electrodes are shown as waveforms on the left, where electrode labeling is shown on the trap schematic in the center. The rest of the (further away) electrodes have similar waveforms, though with lower amplitude, shown as an image plot on the right, with color corresponding to electrode voltage. Voltages are shown for the matching electrodes in the central schematic. Voltages of opposing electrodes (i.e. mirrored in the x-direction) are identical.




Figure 5.3: Voltages required to transport an ion chain from one trap segment to a neighboring one. Voltages shown for electrodes on the right are identical for opposing electrodes, mirrored on the *x* axis. Voltage waveforms on the left are for nearby electrodes, as labeled in the central schematic.

The voltages shown in Figure 5.3 are shown as a function of position z_p . The final step in generating a voltage sequence is by determining the timing of bringing z_p from its initial to final position. The time-dependent voltage sequence is produced by determining time steps $V_e(z_p(t))$, with the choice of $z_p(t)$ as determined in the previous section.


5.2.3 Experimental implementation of shuttling

As a basic functionality test, we apply the methods described above to our experiment, using a trapped 40 Ca $^{+}$ ion in the Golden Gate trap. The sinusoidal voltage ramp was used for this demonstration. Figure 5.4(a) shows CCD camera images of the ion during the various stages of a shuttling trip, in which the ion is being transferred from one electrode to a neighboring one. Each image is acquired with a long exposure (2 s) with the ion placed at a stationary point along the shuttle sequence. In Figure 5.4(a) these images are spliced together, so the "time" axis is not continuous but refers to the discrete time steps.

To estimate heating induced by the shuttle sequence, the ion is transported from its initial position to a neighboring segment, brought to rest, and then transported back towards its initial position. The ion is said to have made two shuttling trips. At the end of this transport sequence, the phonon number of the axial mode is probed using sideband thermometry (see Section 2.3.4). The measured phonon number is shown in Figure 5.4(b) for various numbers of trips. One trip is set to take 170 µs. For all data points, the total time between the start of the shuttling sequence and start of the phonon number detection pulse is kept the same, which ensures that the anomalous heating *not* caused by the transportation is the same for all data points. From a linear fit through the data, we estimate an increase of phonon number due to shuttling to be 0.05(2) phonons per trip.

5.3 SPLITTING

A distinguishing feature of the QCCD architecture is the ability to have multiple simultaneous trapping sites, each occupied with an ion crystal composed of a handful of ions. In order to maintain full connectivity throughout the trap, ions must retain the possibility to interact with each other, even if located in separate trapping sites. Ion-ion interaction is achievable by bringing trapping sites close






Figure 5.4: Experimental demonstration of shuttling sequences. (a) Camera images of an ion in a shuttling sequence, using a sinusoidal ramp. Separate images, each taken with the ion kept stationary at discrete moments along the shuttling sequence, are spliced together. (b) Mean phonon number of the axial mode after after an ion is shuttled back-and-forth between two neighboring electrodes, shown as a function of number of trips. A "trip" is a transfer from one electrode to its neighbor, which is done in 170 µs.

together, while keeping individual wells separated [191–193]. However, since the coupling between ions has an inverse cubic scaling with ion separation, storing ions in separated trapping sites limits how fast interaction operations can be carried out, to an order of milliseconds for typical surface trap architectures. For the purpose of quantum information processing, in which minimizing operation duration is paramount for optimizing process fidelity, coupling usually requires ionic qubits to be co-trapped in a single well. Full connectivity of qubits is then achieved by splitting and merging ions from and into individual trapping sites. This section covers the groundwork for determining potentials and voltages for ion chain separation⁵.

Ion chain separation has previously been theoretically analyzed [194–196] and experimentally investigated [177, 180, 197]. The core concept is to smoothly transform the standard axial quadratic potential, in which ions are co-trapped, into a double-well potential, where ions are located in the two spatially separated potential minima. In order to have two potential minima, the axial potential must contain a quartic component. Generally, as described in Section 2.3.1, the potential around the trapping region are described by a linear combination of spherical harmonic potentials. A full description of a potential with a double well requires a spherical harmonic expansion up to at least the fourth degree. Such an expansion is feasible, but to produce the potentials for ion chain splitting constrained to that basis of spherical harmonics requires excessively high electrode voltages, certainly beyond the range of what the voltage source in our setup can produce. Instead, we consider only the one-dimensional quartic and quadratic components of the potential in the axial direction, along the RF-null, which significantly reduces the voltage requirements. In the following, the design and experimental implementation of the axial potential for ion chain splitting is discussed. The discussion is limited to a two-ion crystal, though the extension to higher ion number strings is straightforward.

5.3.1 Splitting potentials

Initially, two ions are co-trapped in a harmonic potential centered at axial position z=0. The goal is to transform this quadratic potential into a quartic potential with two minima, separated by a distance $2s_f$ after a time t_f . We assume the minima are symmetrically separated, such that their final positions are $z_{f,2}=-z_{f,1}=s_f$. The axial potential, up to fourth-order, is given by

$$V_z(z,t) = \alpha(t)z^4 + \beta(t)z^2.$$
 (5.12)

Odd-order terms are zero, since the splitting action is symmetric around z=0. The zeroth-order contribution can be ignored without loss of generality. The coefficients α and β represent field potential strengths.




Figure 5.5: Schematic depiction of the potentials required for splitting an ion chain, such that ions of mass m are confined in separate potential wells, separated by a distance s_f . β is the quadratic term in the polynomial expansion of the axial field potential, and α is the quartic term. Initial and final values are set to ensure an axial secular motional frequency of ω_z for both ions.

The development of the axial potential that enables ion chain splitting is schematically shown in Figure 5.5. Ions, with mass m and charge q, are initially in a quadratic potential and have an axial common mode frequency of ω_z , which constrains the initial parameters to be $\alpha(0) \equiv \alpha_0 = 0$ and $\beta(0) \equiv \beta_0 = 1/2mq^{-1}\omega_z^2$. The final conditions, at time t_f , are that ions are separated by $2s_f$ and that the ions' individual motional frequencies are the same as their initial frequency, ω_z . These requirements set the final curve parameters $\alpha(t_f) \equiv \alpha_f$ and $\beta(t_f) \equiv \beta_f$ to:

$$\alpha_f = \frac{m\omega_z^2}{8qs_f^2} \tag{5.13}$$

$$\beta_f = -\frac{m\omega_z^2}{4q} = -\frac{\beta_0}{2}. (5.14)$$

We thus have a start and end point for both curve parameters, in which β has switched sign. At some point during the splitting sequence, the quadratic component β must be zero, at which point the trap has a purely quartic potential α_m , whose value is yet to be determined.

We must thus establish how to traverse the two-dimensional $\{\beta,\alpha\}$ parameter landscape in order to join the start and end point. A rule-of-thumb of any ion transport operation is that it is favorable to maintain constant mode frequencies throughout the transport. At the very least, mode frequencies should not change rapidly, to minimize the amount of motional quanta induced by the transport operation. Additionally, since ions are more susceptible to motional heating at lower mode frequencies [29], sequences should be designed to keep them high.

Figure 5.6(a) shows a contour plot of in-phase mode frequencies as a function of the curve parameters α and β , for two trapped $^{40}\text{Ca}^+$ ions. The ions' mutual Coulomb potential is included in the calculation of mode frequencies, though is negligible at t_f . For reference, an initial and final parameter set, $\{\beta_0, \alpha_0\}$ and $\{\beta_f, \alpha_f\}$ are marked with pink discs, using the values $s_f = 100\,\mu\text{m}$ and $\omega_z/(2\pi) = 1\,\text{MHz}$ as example values, reflecting typical experimental conditions. One could choose $\alpha(t)$ and $\beta(t)$ to parametrically follow the contour of constant mode frequency, as denoted by the dashed curve in Figure 5.6(a). In this strategy, α is ramped up, while β is decreased, until β crosses zero, after which α is ramped down again. Unfortunately, this procedure is unrealistic in practice. An experimental limitation to physical ion manipulation is the maximum and minimum voltages that can be applied to the trap's electrodes. It is therefore useful to investigate voltage requirements to enable this transport.

5.3.2 Splitting voltages

Voltage solutions for splitting are obtained using a similar method as in Section 2.3.1. However, instead of expanding potentials contained in a volume around the trap center into spherical harmonics, we now expand potentials along the one-dimensional trap axis into a fourth-order polynomial.




Figure 5.6: Axial mode frequency as a function of quadratic and quartic potential contributions, β and α , for two $^{40}\text{Ca}^+\text{ions}$. Ion chain splitting is achieved by translating from the initial to final parameter set, $\{\beta_0,\alpha_0\}\to\{\beta_f,\alpha_f\}$, marked with pink discs for $\omega_z/(2\pi)=1\,\text{MHz}$ and $s_f=100\,\mu\text{m}$. During this translation, the motional frequency could be conserved, as shown in (a). However, adhering to our setup's voltage limits, a trajectory in which mode frequency drops is unavoidable, as shown in (b)

Using the electrode layout of the Golden Gate trap as an example (see Section 3.2.2), we find that splitting ion chains while maintaining a constant mode frequency requires maximum voltages of more than 3000 V, which far exceeds the 40 V limit of our DACs, and presumably the trap's breakdown voltage⁶. Figure 5.6(b) shows the same data as (a), but zoomed in to highlight the parameter regime where voltages are within the ± 40 V range. In the upper grayed out area, at least one of the electrode voltages exceeds this limit. In order to transfer from the starting parameters to the final parameters within the voltage limit, the axial motional frequency will need to drop below about 0.25 MHz. While unavoidable, the low mode frequency around $\beta=0$ is undesirable, as anomalous heating rates are higher at low mode frequencies [34, 94]. The duration of the splitting operation should therefore be minimized. On the other hand, quick operations also lead to motional heating, introduced by changes in the electric field with frequency components matching that of the motional mode. The sequence time and trajectory through parameter space can be numerically [196] and/or experimentally optimized. The dotted line in Figure 5.6(b) shows an example parameter trajectory. This example has not been numerically optimized, but is based on the intuition that the trajectory should be kept short, smooth, and that the motional frequency should be maximized while adhering to the voltage limit.

To display typical voltage ramps and resulting motional frequencies, we parameterize $\alpha(t)$ and $\beta(t)$ with a polynomial expression, noting that this is an empirically chosen trajectory for the sake of example, and has not been numerically optimized. The trajectory through parameter space can be described in terms of the function $\alpha(\beta(t))$. There are three known points that this trajectory must pass through, the start point $\{\beta_0, \alpha_0\}$, the end point $\{\beta_f, \alpha_f\}$, and point where the quadratic curvature is zero $\{0, \alpha_m\}$, with α_m the maximum allowed quartic coefficient within the voltage constraints. This last point is an inflection point, $\dot{\alpha}(\beta=0)=0$. Altogether, $\alpha(\beta(t))$ can be expressed as a unique third-order

⁶ See Philip Holz's thesis [191] to see what exceeding the breakdown voltage might look like.

polynomial. Letting $\beta(t)$ change linearly in time from $\beta(t=0)=\beta_0$ to $\beta(t=t_f)=\beta_f$, we find the following parametric equations:

$$\alpha(t) = a_3 \left(\frac{t}{t_f}\right)^3 + a_2 \left(\frac{t}{t_f}\right)^2 + a_1 \left(\frac{t}{t_f}\right) \tag{5.15}$$

$$\beta(t) = \frac{m\omega_z^2}{2q} \left(-\frac{3}{2} \left(\frac{t}{t_f} \right) + 1 \right) \tag{5.16}$$

with

$$a_3 = \frac{9}{8} \left(-6\alpha_m + \frac{m\omega_z^2}{qs_f^2} \right) \tag{5.17}$$

$$a_2 = \frac{3}{4} \left(9\alpha_m - \frac{2m\omega_z^2}{qs_f^2} \right) \tag{5.18}$$

$$a_1 = \frac{m\omega_z^2}{2qs_f^2}. ag{5.19}$$

Figure 5.7(a) shows the electrode voltages that are required for splitting an ion chain trapped in a single potential centered on one electrode, into two trapping sites, each centered on the two neighboring electrodes. These voltage solutions are for $\alpha_m=0.55\,\mathrm{kV}\,\mathrm{mm}^{-4}$, $\omega_z/(2\pi)=1\,\mathrm{MHz}$, and $s_f=100\,\mathrm{\mu m}$, using the Golden Gate trap electrode layout, schematically shown on the left. At the chosen value of α_m , two electrode voltages reach a limit of $-40\,\mathrm{V}$. Figure 5.7(b) shows the development of motional frequency throughout the splitting operation, which reduces to 0.22 MHz at the lowest point. Figure 5.7(c) shows the trapping potential along the axial direction throughout the splitting sequence, using the voltages in (a) and simulated potentials for the Golden Gate trap electrodes. For visual clarity, at each point in time, the potential is offset such that its minimum is at 0 meV, and the color scaling is cut off at 20 meV. The emergence of the double-well potential can be seen after about $t/t_f=0.6$, which coincides with $\beta(t)<0$.




Figure 5.7: (a) Electrode voltages for splitting an ion string on one electrode into two separate trap sites on two neighboring electrodes. (b) axial mode frequency and (c) potential during a splitting operation.

Once ions are separated by two segments widths, as is the case at $t/t_f=1$ in the example of Figure 5.7, it is useful to no longer consider the full quartic potential, but instead to treat the potential around each ion separately, as individual local harmonic potentials. This allows us to manipulate individual ions further, without putting unnecessary constraints on the potential between the two ions. Potentials should not be solved independently for the two trapping sites, since a harmonic potential solution around one site creates an undesired field gradient in the other. Instead, both sites are considered simultaneously in finding voltage solutions.

5.3.3 Splitting field sensitivity

An important experimental consideration for ion chain splitting is its sensitivity to stray fields along the axial direction. To consider this, we modify the trap potential to include an additional axial field γ , which may come from static stray fields or from induced fields caused by imperfect modeling and calibration of potentials generated by electrode voltages (and thus possibly time-dependent). The modified potential is

$$V_z = \alpha(t)z^4 + \beta(t)z^2 + \gamma(t)z. \tag{5.20}$$

Ion chain splitting may become unsuccessful if the magnitude of this undesired field is above a certain threshold, as depicted schematically in Figure 5.8(a). The plots in (b) show numerically calculated ion positions during a splitting operation for axial fields of $\gamma=0$ and $\gamma=5\,\mathrm{V\,m^{-1}}$, using the same example parameter trajectory for ion chain splitting from Section 5.3.1. The cutoff value of γ below which successful splitting occurs, is numerically determined⁷ to be $\pm 4.5\,\mathrm{V\,m^{-1}}$. When γ exceeds this limit, the two ions do not split, and end up in the same trapping region.




Figure 5.8: (a) Schematic depiction of splitting potentials, with an additional field γ . Even though there is a double well in the last panel, the offset field forces both ions to settle in a single trapping site. (b) Numerically calculated ion trajectories, for $\gamma=0$ and $5\,\mathrm{V\,m^{-1}}$, show that a small offset field can cause a splitting operation to fail.

To put that value into a practical context: uncompensated stray fields may be on the order hundreds of volts per meter, and may drift by tens of volts per meter on a daily basis⁸. In fact, the act of loading ions, during which some trap surface coating may occur, can alter stray fields by tens of volts per meter. The ability to add a compensating axial field throughout the splitting sequence is therefore essential. The outcome of a splitting operation should be regularly checked (for example by CCD or PMT detection), and the compensating field must be adjusted accordingly.

It may be the case that the stray field γ is time-dependent throughout a sequence, due to imperfectly calibrated potentials generated by the trap's electrode voltages. It is, however, not necessary to ensure $\gamma(t) \approx 0$ for the entire duration of the sequence. The restriction on γ is most crucial during the

⁷ In [196], an analytic approximation of this cutoff value is given by $\gamma < \pm (q/(4\pi\epsilon_0))^{3/5}\alpha^{2/5}$. Using our value of $\alpha_m = 0.55\,\mathrm{kV}\,\mathrm{mm}^4$, this equates to $4\,\mathrm{V}\,\mathrm{m}^{-1}$.

⁸ These estimates are based on our usual values of compensation fields applied to correct for micromotion

'critical point' [196] of the splitting sequence, which is the point at which the quadratic component of the potential is zero, $\beta=0$. The compensation field that is required to produce a successful splitting operation can therefore be a constant value, and does not need to be parametrically adjusted throughout a sequence.

5.3.4 Experimental demonstration of ion chain splitting

The voltage sequence for ion chain splitting, described in the previous section, is applied in our experiment, using the Golden Gate trap. We concern ourselves with demonstrating that the trap can produce the required potentials to separate an ion chain, and leave further measurements of induced heating and decoherence for future investigation. We thus elect to simplify the trajectory through the $\{\alpha,\beta\}$ parameter space shown in Figure 5.6(b): α is set to a fixed value of $0.2\,\mathrm{kV/mm^4}$ and β is linearly ramped from 2 to $-1\,\mathrm{V\,mm^{-2}}$. Applying the resulting voltage sequence directly in our experiment does not result in a successful ion chain splitting operation; both ions are observed (through CCD imaging) to be present in a single well, though are displaced from their initial position by about 140 µm. As described in Section 5.3.3, this is due to an undesired electric field along the axial direction, which skews the splitting procedure. We correct for this by applying a compensation field. A scan over various compensation field strengths reveals that below a value of $47\,\mathrm{V\,m^{-1}}$, the shuttling sequence places ions jointly in the left well, above a value of $56\,\mathrm{V\,m^{-1}}$ ions are placed together in the right well, and only between those values do ions separate into the left and right wells.

Figure 5.9 shows images of the ions throughout a splitting sequence, with the appropriate compensation field. The figure illustrates the position of the ions as a function of the parameter β , which follows the linear ramp described above. Images of the ions are taken at discrete steps along the sequence and are spliced together.




Figure 5.9: Camera images of two ions that are separated by a splitting sequence. Separate images, each taken with the ions kept stationary at discrete moments along the sequence, are spliced together. The sequence is parametrized by the quadratic curvature β .

These results demonstrate that the Golden Gate trap is capable of producing the required potentials to split ion chains. However, it remains to be shown that this process can be done with sufficiently low motional mode heating. Although we have not measured this value, it is unlikely that the calculated voltage sequence results in an optimally low phonon number. This notion is exemplified in the following section, where it is found that applying a calculated voltage sequence for a different type of physical ion manipulation operation results in tens of added phonons. The voltage sequences described in that section require further fine-tuning to minimize mode heating induced by the operation. We expect similar methods will be required for the ion chain splitting sequences.

5.4 ION CRYSTAL ROTATIONS

As described in the previous sections, within the framework of a functional quantum computer chip, ion qubits are required to connect to and share information with other ions. This holds true even for ions in non-neighboring registers (i.e. trap regions). However, the design of the trap puts restrictions on the path that can be taken to bring two ions (or sets of ions) together. As ions travel along a one-dimensional RF-null, they may find another ion or group of ions in the way. Qubit information could still pass through this physical blockade by carrying out a SWAP gate [198], which is the swapping of logical information of two qubits, described by the operation

$$\mathbf{SWAP} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{5.21}$$

with respect to the two-qubit basis ($|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$). Such a gate requires multiple entangling operations, and is prone to errors due to gate time and hardware constraints.

A more straightforward method of transferring an ion's qubit information past another ion in the way is by physically moving those ions around each other. This can be made possible by the inclusion of junctions [199–201] in the segmented trap design, which effectively branches out the network of one-dimensional RF-null channels. Due to the complicated electrode structure required to generate the RF field to confine the ions in the junction region, shuttling individual ions through a junction is challenging and typically results in unwanted heating of the ions' motional mode [200, 202] or hundreds of microseconds transfer duration [62, 190]. This type of disturbance is detrimental to the fidelity of subsequent quantum operations, and should therefore be avoided. Another approach of reordering ion crystals is by physically rotating ions in a crystal around each other within a single trapping region. Rotations, along with splitting and shuttling techniques, allow ion reordering without the need for junctions, and can therefore simplify trap design requirements.

The feasibility for ion crystal rotations has previously been shown in 3D traps [203]. Rotations in a planar trap, however, come with more technical challenges compared with 3D traps [204]. As discussed in the beginning of this chapter, care must be taken to ensure that the voltage sequences applied to the trap electrodes do not produce undesired fields that drive ions away from RF-null perpendicular to the trap surface.

In this section, we investigate ion crystal rotations in a planar RF trap. We focus on methods of simulating, calibrating, and optimizing potential sequences that enable rotations. Experimentally, we aim to minimize rotation-induced motional heating and maintain coherence, with the prospect of using rotations in quantum computational algorithms. Furthermore, we investigate rotations of multi-species ion crystals, as many proposed algorithms in quantum computation [205] benefit from the use of multiple ion species.

5.4.1 Rotation overview

We start by providing an intuitive description of how trap voltages can be manipulated to allow ions to rotate around each other. Initially, two ions are aligned along the trap axis, shown schematically in Figure 5.10(a) as the *z*-axis. The goal is to rotate the two-ion crystal around a radial axis, keeping the center-of-mass of the two ions at a fixed location. In principle, any radial direction can be chosen as the axis of rotation, but we elect to use the *y*-axis, with the ion crystal rotating in the *xz*-plane. The primary reason for this choice is that the *y*-axis of our surface trap has the highest anharmonicity, including third-order contributions to the potential. Restricting the ions' motion to the *xz*-plane allows us to neglect anharmonicity in determining potentials that facilitate rotations.

Referring to the generic trap layout in Figure 5.10(a), a rotation is initiated by applying a positive(negative) bias to the corner electrodes E_3 , $E_4(E_1, E_6)$, which tilts the ion chain from the trap axis. The confinement of the weakest trapping axis (z) is increased by increasing the voltage of the corner

electrodes and decreasing the voltage on the central electrodes (E_2 , E_5), until the ions are aligned along the x-axis, which occurs when the confinement along that axis is weaker than that of the z-axis. The polarity of the diagonal bias voltages of the first step is then reversed and the along-axis confining potential is relaxed. After removing the diagonal bias, the ions are back in their original equilibrium positions, albeit with switched positions.

Naively applying voltages on electrodes as prescribed above typically results in additional undesired electric fields, displacing the ions from their equilibrium position, inducing excessive micromotion, potentially enough to expel ions from the trap. As in the previous sections, improved control is necessary and is achieved by parametrising the potential generated around the trapping region by each electrode, thus giving control in terms of trapping potentials instead of in terms of electrode voltages. Once again, we rely on the spherical harmonics representation of trapping potentials described in Section 2.3.1, as they contain a set of potentials $Y_{l,n}$ that are natural for ion trapping and manipulation, and can be controlled independently. A linear combination of $Y_{l,n}$ makes up the DC trapping potential V_{DC} , as $V_{DC} = \sum_{l,n} v_{l,n} Y_{l,n}$, and can be dynamically adjusted to create the time-dependent potential that moves the ions along the desired path. Three of these terms are well-suited to describe an ion crystal rotation, $Y_{2,0}$, $Y_{2,1}$, and $Y_{2,2}$. To avoid the need to flip back-and-forth in pages, those potentials were given by:

$$Y_{2,0} = 2z^2 - x^2 - y^2, (5.22)$$

$$Y_{2,1} = xz,$$
 (5.23)

$$Y_{2,2} = x^2 - y^2. (5.24)$$




Figure 5.10: Schematic overview of potentials used for ion crystal rotations. (a) Trap layout around the trapping region. (b) Upper plots show the spherical harmonic potentials in the *xz*-plane that are instrumental in ion crystal rotations and the lower images indicate electrode voltages to produce them. Red is a positive voltage and blue is negative, though colors are otherwise not to scale.

The potential fields corresponding to $Y_{2,\{0-2\}}$ are shown schematically in Figure 5.10(b), in the plane of rotation (xz). As a guide to the eye, field gradients, $-\nabla V_{DC}$ are shown as arrows. Additionally, the voltages applied to nearby electrodes to produce these potentials are shown below, with red indicating a positive voltage and blue a negative one. Note, these images are schematic representations; the color coding is not to scale, and the displayed potentials are idealized and do not represent actual trap potentials.

The electrode voltage description of ion rotations can be intuitively translated in terms of the harmonic multipole potentials. First, increasing $v_{2,1}$ tilts the ions off of the *z*-axis in the *xz*-plane. Then, increasing $v_{2,0}$ increases the confinement along the trap axis and lowers it along the *x*-axis, to the point at which orienting the ion chain along the *x*-axis becomes favorable. Decreasing $v_{2,2}$ facilitates the decrease of confinement along the *x*-axis. The polarity of $v_{2,1}$ is then reversed, and $v_{2,0}$ and $v_{2,2}$ are decreased. After all potential terms are back at their initial values, ions are back in the original

locations, but have swapped positions. Note that in addition to this sequence of potentials, the trap also produces an RF pseudopotential (not shown in Figure 5.10).

Applying the above procedure in our experimental setup, it does not take too much trial-and-error to find a set of voltages that successfully swap the positions of two ions. However, this hand-waving level of control of potentials and electrodes is not sufficient for applying ion rotations in the context of quantum computation sequences. In that context, ion crystal rotations are conceptually identical to the quantum SWAP gate. In this regard, the requirement of a rotation is not just to switch the physical position of two ions, but to maintain coherence of the two-qubit system.

This poses the following requirements:

- Rotations should have 100% **success rate**, meaning that after each rotation operation, ions should have switched positions.
- The **duration** of a rotation should be minimized, such that extended quantum computation sequences that depend on crystal rotations are not limited by decoherence.
- Excitation of motional modes should be minimized, since the fidelity of operations following rotations degrade with increasing phonon numbers. For mixed-species crystals, it should at least be low enough that sympathetic cooling can still be used to return to the motional ground-state of relevant modes within a reasonable time.

The sequence of multipole potentials needs to produce rotations that meet these requirements and must therefore be accurately determined. The following section details our method of generating a sequence of potentials for ion crystal rotations. The potential sequence is then to be converted to a voltage sequence that can be applied to the ion trap. Similar as for the shuttling and splitting sequences of the previous sections, the mapping of potentials to voltages is determined using the method in Section 2.3.1, which relies on simulations of potentials generated by trap electrodes. Due to manufacturing and simulation tolerances and stray fields, however, this mapping may not perfectly reflect reality. To ensure that the appropriate voltages are applied to produce the desired potentials, calibration is required. This calibration was discussed in Section 2.3.2.

5.4.2 Rotation potentials

We keep mode frequencies constant for the entire duration of a rotation sequence, in an attempt to minimize motional mode heating induced by the rotation. There are two motivations for this strategy: Firstly, when mode frequencies are too low, they are more susceptible to heating due to spurious electric field noise [29, 113]. Secondly, when frequencies of different motional modes are near degeneracy, exchange of phonons is more likely [206], through which well-cooled modes used for computation can gain energy from other modes that might be less-well cooled. The goal therefore is to generate a sequence of potentials that switch the positions of ions while keeping the crystal's mode frequencies constant.

A given trap potential, a combination of the static DC potential and the RF pseudopotential, determines an ion chain's center-of-mass position, spacing, and mode frequencies. One can thus solve which potentials are to be applied to place ions at given rotation angles, at given motional frequencies. The methods of calculating these positions and frequencies are given in Section 2.1.4. To reintroduce the notation and nomenclature: Ion positions and motional frequencies are determined as a function of the trap's RF pseudopotential $V_{\rm RF}$ and the DC potentials $V_{\rm DC}$. The curvature of the pseudopotential is denoted as $v_{\rm RF}^{(i)}$, which depends on the ion's mass (i refers to the ion index). The DC potential is a linear combination of spherical harmonic potential terms, $V_{\rm DC} = \sum_{l,n} v_{l,n} Y_{l,n}$. We cast the coefficients $v_{l,n}$ into the vector \vec{v} .

We are now tasked with finding solutions of the vector $\vec{v}^{(n)}$, at discrete steps n, that facilitate the rotation of an ion crystal. We keep the value of $v_{RF}^{(i)}$ fixed. In our experiment, a suitable starting potential $\vec{v}^{(0)}$ that allows for stable storage of ions is characterized by motional frequencies on the order of $\omega_{\{x,y,z\}}/(2\pi) \approx \{3,3,1\}$ MHz, though avoiding degeneracy in ω_x and ω_y by about 200 kHz.

This corresponds to $\nu_{2,0}^{(0)} \approx 10 \, \mathrm{V \, mm^{-2}}$ and $\nu_{\mathrm{RF}}^{(i)} \approx 160 \, \mathrm{V \, mm^{-2}}$. Setting $\nu_{2,2}^{(0)} \approx 10 \, \mathrm{V \, mm^{-2}}$ ensures that the degeneracy in ω_x and ω_y is lifted.

The goal is to numerically find values for $\vec{v}^{(n)}$ that minimize a cost function $C(\vec{v})$ at steps n of an ion crystal rotation, characterized by a desired rotation angle θ_n . Here, θ_n is the intended angle between two ions in a chain and the z-axis, in the xz-plane, at step n, homogeneously distributed between 0 and π . The cost function is foremost composed of the restriction that $\vec{v}^{(n)}$ should place the two ions (indexed i and j) at positions $\vec{x}_0^{(i)}$ and $\vec{x}_0^{(j)}$ such that they have the required angle θ_n . Additional restrictions are imposed to improve the quality of the operations. Since we aim to minimize mode heating, this entails keeping motional mode frequencies high and avoiding mode crossing. After a rotation, the potential should be identical to the initial potential.

Placing ions at the required angle θ_n is achieved by minimizing the cost function

$$C_0(\vec{v}^{(n)}) = 1 - \frac{(\sin \theta_n, 0, \cos \theta_n) \cdot (\vec{x}_0^{(j)} - \vec{x}_0^{(i)})}{\left| \vec{x}_0^{(j)} - \vec{x}_0^{(i)} \right|},\tag{5.25}$$

noting that the ion equilibrium positions, $\vec{x}_0^{(i)}$ and $\vec{x}_0^{(j)}$, are a function of the spherical harmonic coefficients $\vec{v}^{(n)}$. We use the solution $\vec{v}^{(n-1)}$ as an initial value for the minimization of $C_0(\vec{v}^{(n)})$, to reduce computation time and ensure a smooth rotation. For a single-species two-ion crystal, the equilibrium positions $\vec{x}^{(i,j)}$ can be analytically calculated [88], thus reducing computational cost.

Degeneracy in mode frequencies ω_i is avoided with the cost function

$$C_1(\vec{v}) = \sum_{i} \sum_{j>i} \frac{1}{(\omega_i - \omega_j)^2}.$$
 (5.26)

Maintaining high frequencies can be implemented by the cost function

$$C_2(\vec{v}) = \sum_i \frac{1}{\omega_i^2}.\tag{5.27}$$

Ensuring that the final potential (at $\theta=180^\circ$) is identical to the initial potential ($\theta=0$), can be attained by keeping frequencies as close as possible to their initial values $\omega_i^{(0)}$ throughout the entire rotation, using the cost function

$$C_3(\vec{v}) = \sum_{i} (\omega_i - \omega_i^{(0)})^2. \tag{5.28}$$

This cost function also helps avoid rapid changes in motional frequency during a rotation sequence, which would lead to motional excitation. An indirect consequence of C_3 is that it helps achieve the conditions imposed by C_1 and C_2 . The full cost function to be minimized is given by

$$C(\vec{v}) = \sum_{i=0}^{3} \lambda_i C_i(\vec{v})$$
(5.29)

with λ_i a set of weights.

For single species rotations, there exists a unique solution that keeps all mode frequencies constant throughout the rotation sequence, for given starting parameters. Setting λ_0 and λ_3 to a non-zero value is sufficient for determining this solution uniquely. If $\nu_{RF}^{(i)}$ is not a controllable parameter during a sequence, multi-species rotations cannot maintain constant frequencies and might thus benefit from λ_1 and λ_2 contributions. In the multi-species case, we determine the values of the weights λ_i empirically, looking for values that produce potential sequences with no discontinuities that uphold the imposed restrictions. We find that the weights $\lambda_0 = 10\lambda_1 = 10\lambda_2 = \lambda_3$ provide satisfactory sequences.

In summary, we obtain a list of desired potentials, in terms of multipole contributions, that facilitate an ion crystal rotation, parameterized by a rotation angle θ_n . For the experimental results discussed

later in this section, we have determined rotation potentials at *uniform* rotation angle intervals between 0 and π . One could alternatively choose angle intervals such that the change in rotation angle is not constant, potentially reducing sudden changes in velocity of the ions as they initiate rotation. A suitable distribution of rotation angles could be determined following the methods in Section 5.2.1, where potential ramps were discussed in the context of ion shuttling. In preparing our ion crystal rotation experiments, we calculated that the electrical filters strongly suppress rapid changes to the confining potential that would otherwise induce high motional heating. We therefore expect that a non-uniform distribution of the rotation angles is not a necessity in our experiment.

5.4.3 Experimental implementation

In this section, we describe how ion crystal rotations are implemented in our experiment. Presented experiments make use of the 'Berkeley trap', a planar segmented surface trap (see Section 3.2.1.1). Rotations are carried out on either a two- 40 Ca⁺crystal, or a 40 Ca⁺- 88 Sr⁺crystal.

5.4.3.1 Calculated rotation voltages

Following the methods of the previous section, we calculate ideal rotation potentials. The calculated potentials as function of rotation angle are shown in Figure 5.11(a). To obtain these results, the rotation angle θ_n was varied from 0° to 180° in 100 steps, which ensures a smooth change in electrode voltages. We display the coefficients of the spherical harmonic potentials⁹ that contribute to the ion crystal rotations, $Y_{2,0}$, $Y_{2,1}$ and $Y_{2,2}$. The resulting waveforms follow the intuitive description given in Section 5.4.1.

In Figure 5.11(b), the corresponding voltages of nearby electrodes are shown, following the indexing from Figure 5.10 (which, incidentally, is the electrode layout of the Berkeley trap). The development of voltages also follows the previously provided intuitive description.

Figure 5.11(c) shows the mode frequencies of the crystal during a rotation. Displayed are in-phase (com) and out-of-phase (str) modes for the three mode axes, x, y, and z. These are local axes with respect to the ion crystal: z is aligned with the trap axis at the start and end of a rotation but follows the angle of rotation, i.e. the axis of the ion crystal. The y-axis is always perpendicular to the trap surface. Ca - Ca rotations can be done while keeping all mode frequencies constant throughout. The mixed-species rotation, Ca - Sr, cannot maintain constant mode frequencies, unless the RF power can be adjusted during a rotation sequence. It is possible to have solutions where mode frequencies do not cross. However, the resulting voltage waveform does appear qualitatively less smooth compared with the single-species case.

Figure 5.11(d) shows the predicted ion trajectory during a rotation sequence. In the single-species case, since the axial in-phase mode, z_{com} , is constant, so is the ion separation. The trajectories of the two ions therefore together trace out a circle. The mixed species trajectory, on the other hand, looks comparatively less circular, as mode frequencies are not kept constant, and since Ca and Sr experience unequal confining forces due to the mass-dependent pseudopotential.

Also in contrast to single-species crystals, stray electric fields are much more detrimental to the success of a mixed-species rotation. Due to the pseudopotential's mass dependence, excess fields drive ions in a multi-species crystal away from the RF-null unequally. Any field in the x direction (radial, parallel to the trap plane) therefore tilts a mixed-species crystal diagonally in the xz-plane. In typical ion trapping potentials, an offset field on the order of $E_x = 10 \,\mathrm{V\,m^{-1}}$ produces a diagonal tilt similar to what would be caused by a potential curvature on the order of $v_{2,1} = 1 \,\mathrm{V\,mm^{-2}}$. This deviation can disrupt the success of a rotation, from which it is apparent why these potential terms need to be calibrated. If the E_x field is well understood, an intentional offset in this field can be used advantageously in assisting uni-directional rotations, as has been previously demonstrated [66, 207] for use in state preparation. Minimizing heating, decoherence, and swap duration of such rotations is not crucial and has not been considered previously for use within quantum computational sequences.




Figure 5.11: (a) Simulated contributions of second order potentials, $Y_{2,0}$, $Y_{2,1}$, and $Y_{2,2}$. The potentials are optimized to minimize changes in trap frequency during the rotation. (b) Voltage sequence applied to nearby electrodes using the labeling convention from Figure 5.10. Further electrodes have been omitted, but behave similarly. Required voltages are based on simulations, and do not include simulation imperfections and effects of signal filtering. (c) Resulting trap frequencies of the in-phase (com) and out-of-phase (str) modes, where the z modes follow the axis of the ion crystal. (d) Ion positions during an optimized rotation.

Both single- and mixed-species crystal rotations thus require accurate control over a time-varying potential, generated by applying a sequence of voltages to the trap electrodes. The effect of electrical filters on these voltage sequences is discussed in the following section.

5.4.3.2 Electrical filters

Voltage noise on trap electrodes presents a major contribution to mode heating of trapped ions [29], and needs to be minimized. Electrical filters attenuate voltage noise at ions' trap frequencies, and are therefore an essential component of an ion trap setup. In addition to suppressing electrical noise, filters will attenuate and delay an applied voltage sequence. Thus, the effect of filters needs to be taken into account when designing voltage sequences to realize trap potentials for physical ion manipulation operations. In this section the employed filter system and its effect on rotation sequences is discussed. The filter system differs from the system described in Section 3.1.1.3, as the rotation measurements were performed in an earlier version of the experimental setup. The systems are fundamentally similar.

Our filter setup for a single electrode is schematically shown in Figure 5.12(a). The filter system consists of a two-stage LC low-pass filter (2 \times 4.7 μ H and 47 nF), per electrode, incorporated in the cryostat's outer heat shield (at a temperature of \sim 150 K). Additional filtering is performed in the vicinity of the trap with 8 parallel 4.7 nF capacitors. The cut-off frequency of 47 kHz has been chosen so that it sufficiently filters noise around the motional mode frequencies (42 dB attenuation at 1 MHz, and 80 dB at 3 MHz), while allowing voltage manipulation of the electrodes on timescales of tens of microseconds. Inside the vacuum chamber are two connections to every electrode, each with separate but identical filter components. In practice, the second connection is only used to test the connectivity throughout the entire system, but cannot be neglected for filter analysis. The output amplifier of the AWG is included in the filter model as an effective series resistance (20 Ω), which is taken as an estimate based on measuring the frequency response of the AWG directly at the output.

Furthermore, the non-zero resistance of the electrical wiring is included in the filtering model. In a cryogenic system, electric wiring represents a major part of the heatload and thus there is a trade-off between low thermal conductivity and good electrical connectivity. We employ phosphor-bronze cables, with a resistance of 4Ω per meter at cryogenic temperatures. The ground connections for all electrodes consist of a twisted pair of these wires that have a non-negligible resistance of about 4Ω , both from the vacuum feedthrough to the first filter stage and from the filter stage to the trap PCB. This wire resistance has the detrimental side-effect that the reference potential as supplied at the AWG output is not identical to the potential at the filters, or at the trap, if current is flowing through the capacitors.

The filter's transfer function at the trap, as determined by calculations based on ideal components, is shown in Figure 5.12(b). Typical voltage sequences for rotations with a duration of 25 μ s have their main frequency component at 20 - 40 kHz. Thus the applied electrode voltage waveforms are attenuated by 1.5 - 2 dB and have a phase delay of about 5 μ s. Figure 5.13(a) shows how the waveforms of the significant rotation potentials $Y_{2,\{0,1,2\}}$ are affected by filtering (Note: simulated results include electrode crosstalk, which is discussed in the following paragraph, but is not dominant in Fig 5.13). The distortion in the potential waveforms can impair the performance of a rotation or even impede rotations altogether.

As a result of the finite ground wire resistance, supplying a time-dependent voltage to one electrode will change the reference potential at the trap. The voltage difference between AWG and trap ground levels results in an effective crosstalk between electrodes and is calculated to be the dominant source of crosstalk (compared with, for example, crosstalk due to parasitic capacitance between neighboring electrodes). The simulated crosstalk transfer function between any two electrodes is displayed in Figure 5.12(c). Upon applying a typical voltage sequence to a single electrode, we measure a crosstalk voltage of at most 3% on all other electrodes, with a settling time of about 20 μ s. However, a full rotation sequence employs all electrodes, with a combination of positive and negative voltages. During a typical rotation, the effective crosstalk voltage is between 0.5 and 1.0 V, resulting in a voltage error between 3% and 10% for different electrodes. This voltage error perturbs the intended potentials, most notably adding an undesired electric field perpendicular to the trap plane ($Y_{1,-1}$) of up to 6 V/mm, as can be seen in Figure 5.13(b).




Figure 5.12: (a) Electronic schematic of a single electrode's filter network. (b) Simulated filter transfer function at the trap electrode. (c) Simulated filter crosstalk, showing voltage pickup on other electrodes due a voltage applied to a single electrode. The gray regions indicate the frequency range of voltage sequences for rotations in our experiment.

The inclusion of this undesired field due to filtering highlights the complexity of physical ion crystal manipulation on surface traps, compared with 3D traps. In 3D traps, electrodes on opposing sides of the trapping region have identical voltages applied throughout a rotation sequence and remain similar even when filtered. By symmetry, ion crystals will therefore remain near the RF-null. The geometry of surface traps does not have symmetry around the trapping region perpendicular to the trap plane. Therefore, if a calculated voltage sequence is designed to keep micromotion compensated throughout, the *filtered* sequence will likely generate an additional electric field perpendicular to the trap plane, the direction that lacks mirror symmetry of electrodes around the ion.

The unintended offset field pushes the ions further from the RF-null, thus substantially increasing micromotion. While some micromotion is not necessarily detrimental to an ion's motional state¹⁰, at the simulated 6 V mm⁻¹ of offset field during a rotation sequence, the ions do not remain crystallized, let alone remain near the motional ground state. For some of the calculated voltage sequences, the unwanted field is sufficient to expel ions from the trap, since along the perpendicular axis the energy barrier for escape is the lowest. Therefore, a method of compensating electrode voltage sequences in order to achieve low-heating rotations is required. Through simulations, we deduce which *applied* filtered voltage sequence produces the *intended* voltage sequence at the trap electrodes.

The transfer functions displayed in Figure 5.12(b) and (c) are based on ideal electrical components. In practice, tolerances in components, as well as other sources of inductance, capacitance, and crosstalk, causes the modeled transfer function to deviate from the physical electrical response. Nevertheless, the model represents an approximation that is necessary for creating filter-corrected voltage sequences for ion crystal rotations. It is worth emphasizing that without the application of filter and potential

¹⁰ Independent experiments where we intentionally induce small excursions of micromotion by placing ions in a static position away from the RF-null have resulted in negligible heating. In these experiments, we measured an increase in an ion's motional mode occupation of less than 0.1 phonons after being subjected to a field of $\pm 0.5\,\mathrm{V\,mm^{-1}}$ for 200 µs.




Figure 5.13: Comparison of set potentials and simulated resulting potentials when applied electrode voltages are filtered, for (a) curvatures and (b) gradients.

calibration models to our solutions, the voltage solutions typically do not lead to successful rotations at the desired time scales (20 - $30\,\mu s$). Voltage solutions that take our models into account lead to successful rotations that, despite inducing a high level of motional heating, can be used as a starting point for further optimization. Our method of minimizing rotation-induced heating is outlined in the following section.

5.4.3.3 Automated optimization

So far we have determined the potentials required for ion crystal rotations, simulated and calibrated the trap's electrodes (see Section 2.3.2) to determine the static voltages required to produce those potentials, and calculated the adjusted voltage waveform required to overcome distortions due to the system's filters. Upon applying the resulting voltage sequence to two trapped 40 Ca $^{+}$ ions, the results are on the one hand encouraging — rotations are successful — but on the other hand disappointing — the ion crystal has acquired tens of phonons. To avoid this heating, further optimization of the voltage sequences must be performed. The methods of the previous sections have provided us with a usable starting point for in-experiment optimization, with a measurable objective, namely minimization of the mean phonon number. Our method for optimizing the rotation sequences is outlined in this section.

The optimization procedure is schematically shown in Figure 5.14, and is based on the following four steps: (I) An initial rotation sequence is simulated, using the calibrated multipole and filter models, and described using a set of control parameters. (II) Using this sequence as a starting point, a neural network-based learning algorithm is used to optimize rotation sequences to minimize induced motional heating. (III) Experimental data, where heating is measured for various rotation sequences, is used to train the network, from which an optimized rotation is found. (IV) This optimized rotation is used in further experiments. These steps are further detailed below.

Step I: To start off with, we require a set of control parameters that can be tuned that change the voltage sequence. This set represents the optimization parameters — the tunable parameters that, when set to the right value, produce the desired outcome, namely the lowest achievable increase in phonon number after a rotation. In the experiment, the parameters that can be changed are the set of time-dependent electrode voltages. The voltage sequences consist of about 50 steps (based on our target sequence duration of 20 - 30 µs and AWG sample rate of 2 MHz), with voltages applied to 21 electrodes. In total, 1050 parameters can be changed. Finding an optimum in a 1050-dimensional parameter space could be tedious, thus some simplification of the parameter space is desired. First, instead of describing the sequence in terms of electrode voltages, we translate applied voltages to the expansion of multipole potentials. These potentials give a better intuition of what forces are being applied to ions, and thus allow for easier manipulation in the optimization. Second, instead of defining the potentials at each time-step, we parameterize the sequence: noting that the waveforms




Figure 5.14: Schematic overview of rotation sequence optimization. (I) An initial guess for sequence parameters is made by simulating ideal rotations and applying our filter model and multipole calibrations. (II) Sequence parameters are optimized using a machine learning neural network, which is trained using (III) experimentally determined rotation induced heating. (IV) An optimized rotation sequence is used in further experiments. The symbols indicate parameters relevant to that step, of which the description can be found the main text.

of the sequence of multipole potentials resemble sine-curves (see Figure 5.11), we choose a Fourier component parametrization. The sequence of applied potentials, discretized into N steps, is given by

$$\vec{v}_{n,\text{set}} = \vec{a}_0 + \sum_{p=1}^{p_{\text{max}}} \vec{a}_p \sin\left(\frac{p\pi n}{N}\right). \tag{5.30}$$

Here, $\vec{v}_{n,\text{set}}$ is a vector of size 8, containing coefficients of spherical harmonic potentials of degrees l=1 and l=2, at each sequence step n. \vec{a}_p is a set of weights that describe the amplitude of the Fourier components. An increasingly high chosen order of Fourier components improves the approximation of the filter-distorted multipole waveforms. For our optimization approach, we choose $p_{\text{max}}=2$ to limit the number of optimization parameters \vec{a}_p .

In addition to varying the total rotation time (by adjusting N), if $p_{\text{max}} = 2$, there are 24 weights that parametrize the sequence ($\vec{a}_{\{0,1,2\}}$, each with 8 terms), giving a total of 25 parametrization coefficients. It is not required to include the sample rate of the sequence (2 MHz) as an optimizable sequence parameter, as this frequency is heavily filtered in our set-up. For notational simplicity, we cast all parametrization coefficients into a single vector, \vec{a} .

The intended potential sequence $\vec{v}_{n,\text{des}}$ shown in Figure 5.11 differs from the set potential $\vec{v}_{n,\text{set}}$, due to static deviations described by the potential calibration model and dynamic deviations described by our filter model (Section 5.4.3.2). We numerically find a set of coefficients \vec{a}_{init} that, under the influence of these discrepancies, result in the intended potentials $\vec{v}_{n,\text{des}}$.

Steps II - III: With a first guess for control parameters, we can start adjusting their values to reduce the sequence induced heating. The control parameters do not independently influence the heating, so

manually tuning each parameter separately is not viable. We therefore use an automated calibration routine that can find optima of multi-parameter systems. There exist many possible optimization candidates [208] (hill climbing, direct stochastic, simulated annealing, genetic algorithms, etc), each with advantages and disadvantages. We use a machine-learning approach, based on a single hidden layer neural network [209] (using the MATLAB® Neural Network Toolbox™), to find the set of waveforms that has the least influence on the motional state of the ion crystal. We choose this method for its relative computational simplicity, its ability to explore a high dimensional parameter space, and ability to find a global optimum, not just local ones. Furthermore, neural networks have the advantage that one does not require a predetermined model that describes the relation between input (sequence parameters) and output (motional heating)¹¹.




Figure 5.15: Schematic of a neural network for machine learning. The value l_j of a node is given by the sum of weighted features of the previous layer and a bias, with a sigmoid function applied.

A neural network¹² is a collection of nodes and edges, schematically shown in Figure 5.15. The nodes i all hold scalar values. Each node value is the sum of the values associated with the edges j entering each node i from the left, plus a bias b_j . The edge values are a function of the value of the node that it originates from, and a weight factor w_j . A commonly used function for the edge values is a sigmoid function¹³, which produces a value between -1 and 1, with a smooth transition between the two. The collection of intermediate nodes (i.e., not the input or output nodes) are referred to as $hidden\ layers$. This type of nomenclature underscores why a neural network is often referred to as a 'black-box' function, highlighting that the internal mechanism is either irrelevant or too complicated to be described intuitively. The important functionality of a neural network is that it maps a set of input values (features) to a set of output values (features), dependent on the network's weights features0 and features1. The goal of our network in particular is to make a prediction of how much heating will be induced by the rotation sequence features1 based on the potential sequence parameters. features2 and features3 and features4 are the inputs and output of the network, denoted as

$$\bar{n}_{\text{pred}} = \mathcal{N}(\vec{a}, \vec{w}),$$
(5.31)

with the network function \mathcal{N} , and the weights and biases grouped into the vector \vec{w} .

Choosing a correct network size (i.e. the amount and dimension of the hidden layers) is not an exact science, and often a product of trial and error and rules-of-thumb [210]. For the measurements presented here, we have found a single hidden layer consisting of 15 nodes produce trainable networks that are able to produce statistically significant correlation between network predictions and measured outcomes.

In order for the network \mathcal{N} to make accurate predictions, it must be 'trained'. Training is the process of feeding the network realistic examples of inputs and outputs. The network then adjusts its weights and biases \vec{w} to attempt to improve its ability to predict an output accurately. If a trained network

¹¹ Note 30 in Appendix e

¹² Note 31 in Appendix e

¹³ $\sigma(x) = \frac{2}{1 + \exp(-2x)} - 1$

makes accurate predictions on heating rate, that network can be used to find a set of input parameters that minimize the predicted output, and therefore minimize the experimental mean phonon occupation \bar{n} induced by a rotation. We focus on minimizing heating on the axial modes, but the described method is applicable to any set of motional modes.

Training and optimizing comprise the following general steps, which will be outlined in more detail further below. The simulated sequence parametrized by \vec{a}_{init} is used as an initial point for the optimization procedure. (i) A large amount of sequences with parameters that marginally deviate from \vec{a}_{init} are generated. (ii) Electrode voltages that realize these sequences are determined using the method in Section 2.3.1. (iii) The voltage sequences are applied in the experiment to rotate an ion crystal, from which (iv) heating is measured. (v) The heating results are used as targets to train the neural network, with associated sequence parameters \vec{a} as features. (vi) The trained network is used to find a new parameter set \vec{a}_{new} that, according to the network, should result in minimal heating. (vii) This new parameter set is used as a new central point around which to create a set of sequences, as in (i). Steps (i) - (vii) constitute a single iteration of the optimization. These steps are repeated several times, until the measured heating rate converges. This procedure is outlined in more detail below.

(i) In our experiment, for the initial training iteration, K = 500 rotation sequences (indexed k) are generated by randomly selecting a set of sequence parameters $\{\vec{a}^{(k)}\}$ within given bounds around the initial parameter set \vec{a}_{init} . We have chosen bounds by experimentally finding limits on parameters that do not result in expulsion of ions from the trap and result in rotations. Bounds that are within a safe limit of these requirements are typically ± 2 - 5% of the respective coefficients of \vec{a}_{init} . We do not adjust the zeroth-order Fourier components of the parametrized sequences (\vec{a}_0 in Eq. 5.30), which govern the static potential before and after a rotation, to avoid possible changes in mode frequencies before a sequence. The fourth term of \vec{a}_0 , which corresponds to the $Y_{2,-2} = xy$ spherical harmonic potential, has been set to $16 \,\mathrm{V}\,\mathrm{mm}^{-2}$. This tilts the radial motional axes such that the x-mode, which would otherwise be parallel to the trap surface, has more overlap with the Doppler cooling beam, and are therefore more efficiently cooled. This additional potential ideally leaves the axial mode frequencies unaffected. The number of sequence steps n is included in \vec{a} as an optimizable parameter. (ii) The voltages $\vec{\alpha}$ required to realize this set of rotation sequences are determined with $\vec{a}_n = \mathbf{V}\vec{v}_{n,\mathrm{set}}$ (see Section 2.3.1).

(iii) These voltage sequences are applied to an ion crystal in our experimental setup. Rotation sequences are prefaced with the usual pulses of Doppler cooling, optical pumping, and sideband cooling (see Section 2.2.8) which prepares ions in the $S_{1/2}(m=-1/2)$ level, near the motional ground state for the axial common and stretch modes. (iv) At the end of the sequence, the resulting mean phonon occupation is measured. The method of measuring this value depends on the severity of rotation induced heating. The initial set of sequences produce 10 - 100 phonons. In this regime, phonon numbers can be estimated by monitoring the decay of Rabi oscillations when resonantly driving a carrier transition. When rotations have been optimized enough to produce only a few motional quanta, the phonon number is more accurately estimated by monitoring the excitation of motional sidebands [211]. In contrast to Section 2.3.4, when doing sideband thermometry with more than one ion (with only a global 729 nm beam, as in our experiment at the time), one cannot extract a mean phonon number simply by evaluating the ratio of the blue and red sideband excitation. Instead, we fit the measured excitation curves to an analytic model with the mean phonon number $\bar{n}^{(k)}$ as a fitting parameter. The analytic model is derived from the two-ion Hamiltonian given by the James-Cummings model (see Appendix b). Note, when doing sideband thermometry, we only probe the final phonon number of the axial common mode.

(iv) The neural network \mathcal{N} is trained, with the set of sequence parameters $\{\vec{a}^{(k)}\}$ as features, and the set of measured heating rates $\{\bar{n}^{(k)}\}$ as targets. Training is the method of adjusting the network's weights and biases \vec{w} to minimize the cost function J given by

$$J = \sum_{k} \left(\mathcal{N}(\vec{a}^{(k)}, \vec{w}) - \bar{n}^{(k)} \right)^{2}$$
 (5.32)

Minimizing this cost function in a neural network is commonly achieved by a process called backpropagation [212], which is a computationally cheap version of optimizing each individual weight through gradient descent. This procedure is built in to the Neural Network Toolbox.

(vi) The trained network is now capable of making estimates of rotation induced heating as a function of sequence parameters \vec{a} , though only within the bounds initially set for generating $\{\vec{a}^{(k)}\}$. We can thus use the network to find a parameter set \vec{a}_{new} that minimizes¹⁴ the predicted heating \bar{n}_{pred} , as $\min_{\vec{a}} \mathcal{N}(\vec{a}, \vec{w})$, where \vec{a} is constrained to the same bounds as $\{\vec{a}^{(k)}\}$.

The network \mathcal{N} typically has multiple local minima. Therefore, the computed value of $\min_{\vec{a}} \mathcal{N}(\vec{a}, \vec{w})$ depends on the supplied initial values of \vec{a} . We use parameters corresponding to the best results from the scan (typically the lowest 10% mean phonon numbers \vec{n}) as starting points for the minimization. For our experimental data, this selection results in a unanimous consensus for the value of \vec{a}_{new} that minimizes \mathcal{N} . This parameter set is used as a central point around which new random sequences are generated. We thus start again from step (i), with a new central value \vec{a}_{new} .

Steps II and III are iteratively repeated, until the measured heating rate no longer decreases. In our experiment, we have seen that after four iterations, the optimization routine generally converges. Figure 5.16 displays the measured average phonon number during the four iterations of the optimization. The first two iterations make use of the Rabi oscillation decay method of determining phonon number, while the last two use sideband thermometry. Several optimization iterations bring down the average rotation-induced phonon number from tens to a few quanta.

In our setup, a heating rate measurement for a single parameter setting \vec{a} takes a little over a second, and in four optimization iterations, two-thousand of such measurements are taken. In each iteration (except the last), a neural network is trained, a new optimal value of \vec{a} is determined, and a new set of sequence parameters is generated, which takes a few minutes of computation time. Overall, the duration of the optimization procedure is just short of an hour. Due to drifts in the stray electric field (caused, for example, by changes in surface charges on the trap electrodes), which typically occur in the span of several hours or days, rotation sequences may require re-optimization. Although not experimentally verified, we presume that re-optimization could be achieved using the latest working experimental parameter set as an initial guess for \vec{a}_{init} , instead of using the value as calculated from models, and would require less iterations to reach convergence.




Figure 5.16: Histograms of measured rotation induced heating during four iterations of the automated sequence optimization. Each iteration contains about 500 measurements.

Step IV: The final iteration produces another series of randomly generated rotations, with a mean heating rate of a few phonons. From this set, the result with the lowest measured heating is selected. This optimized rotation sequence is used in further characterization experiments, discussed in the following section.

The optimized rotation sequence is $25\,\mu s$ long. Figures 5.17(a) and (b) display the calculated potentials and mode frequencies that the voltage sequence is predicted to produce. The potential calibration and the effect of electrical filters are included in this estimate.

The waveforms of the calculated mode frequencies deviate from the intended behavior, being that mode frequencies should remain unchanged throughout a rotation sequence. This deviation suggests either that the *calculated* mode frequencies do not accurately represent that actual mode frequencies

¹⁴ Minimization is done with Matlab's fmincon, using an 'interior-point' algorithm [213].




Figure 5.17: (a) Simulated potentials and (b) mode frequencies for of the sequence selected from the final iteration of the rotation optimization, measured to have the lowest rotation induced heating.

in the experiment, or that the optimization procedure has found a low-mode-heating solution for rotation sequences that does not necessitate constant mode frequencies, and that phonon exchange due to mode-mixing is not prominent. The fact that the initial rotation sequence, parametrized by $\vec{a}_{\rm init}$, required further optimization despite it having been calculated to produce constant mode frequencies during the rotation, suggests that our mode frequency calculation is inaccurate. This inaccuracy is likely due to an incomplete model of the electronic transfer function of the trap electrodes.

5.4.4 Experimental results

5.4.4.1 Single species rotations: Ca - Ca

In this section, results of the Ca - Ca ion crystal rotations are presented, using the sequence calculated and optimized in the previous sections. When at rest, (i.e., when at the position before and after a rotation) ions have the common mode frequencies $\omega_{\{x,y,z\}}/(2\pi)=\{2.8,3.5,1.25\}$ MHz, where z is the axial (RF-free) direction, and x and y the radial directions parallel and perpendicular to the trap surface. Rotations occur in the xz plane. Both ions are sideband-cooled near the motional ground states of the axial common and stretch modes (0.017(16) and $0.011^{(+4)}_{(-5)}$ phonons respectively) and prepared in the $S_{1/2}(m=-1/2)$ level prior to the rotation.

As introduced in section 5.4.1, the goal of ion crystal rotations is to emulate a **SWAP** gate, in which we require a high rotation success rate, low rotation induced heating, and a minimal loss of qubit coherence. These aspects will be discussed in turn below.

-Success rate-

The success rate describes how often a rotation sequence physically swaps the position of two ions. To measure this value, we must be able to distinguish the two ions. This is achieved by preparing the two ions in two different states, $|0\rangle$ and $|1\rangle$, applying the rotation, and monitoring if the states have switched position. In our experiment, these states are encoded in the $|0\rangle \equiv |4S_{1/2}(m=-1/2)\rangle$ and $|1\rangle \equiv |3D_{5/2}(m=-1/2)\rangle$ levels. State detection is performed by monitoring state-dependent fluorescence ($4S_{1/2} \leftrightarrow 4P_{1/2}$, 397 nm) with a CCD camera, which allows us to measure the two ions' state separately.

An obvious success rate measurement would be to apply a π -pulse to one of the two ions to prepare, for example, the $|01\rangle$ state, apply a rotation, and use state-detection to check if $|10\rangle$ is measured. At the time of measuring, however, our setup had not yet incorporated single ion addressing. Limited to global operations, our sequence is as follows: we first apply a global π -pulse on the ions initialized in $|00\rangle$, defined as

$$R_x\left(\frac{\pi}{2}\right) = \exp(-i\pi/4(\sigma_x^{(1)} + \sigma_x^{(2)})),$$
 (5.33)

Table 5.1: Truth table for state detection before and after a number of rotations. Values are the percentages of detecting a final state, for a given input state. For each number of rotations, 250 trials were performed, of which the initial states $|01\rangle$ and $|10\rangle$ were post-selected

		Final state			
Number	Initial	00>	$ $ $ 01\rangle$	$ 10\rangle$	$ 11\rangle$
of			' '	' '	
Rotations	state	(%)	(%)	(%)	(%)
0	$ 01\rangle$	6 ± 3	94 ± 3	0 ± 1	0 ± 1
	$ 10\rangle$	10 ± 5	0 ± 2	90 ± 5	0 ± 2
1	$ 01\rangle$	2 ± 2	0 ± 2	98 ± 2	0 ± 2
	$ 10\rangle$	5 ± 3	95 ± 3	0 ± 2	0 ± 2
2	$ 01\rangle$	9 ± 4	91 ± 4	0 ± 2	0 ± 2
	$ 10\rangle$	10 ± 4	0 ± 2	90 ± 4	0 ± 2
3	$ 01\rangle$	12 ± 4	0 ± 1	88 ± 4	0 ± 1
	$ 10\rangle$	8 ± 4	92 ± 4	0 ± 1	0 ± 1
4	$ 01\rangle$	6 ± 3	94 ± 3	0 ± 1	0 ± 1
	$ 10\rangle$	6 ± 3	0 ± 1	94 ± 3	0 ± 1
5	$ 01\rangle$	12 ± 4	0 ± 1	88 ± 4	0 ± 1
	$ 10\rangle$	3 ± 3	97 ± 3	0 ± 2	0 ± 2

with Pauli operators $\sigma_x^{(i)}$ for ions 1 and 2 [11]. This operation prepares the ions in the state $1/2(|00\rangle+i|01\rangle+i|10\rangle-|11\rangle)$. State detection is performed, followed by a number of crystal rotations, and finalized with another state detection. Measured data is post-selected to only include measurements that were in $|01\rangle$ or $|10\rangle$ before a rotation. A successful rotation is indicated by a final measurement of $|10\rangle$ or $|01\rangle$, respectively, for odd numbers of rotations, and no change from the initial state for even numbers.

Table 5.1 displays the resulting truth table. The data for each number of rotations consists of 250 repetitions of the experiment. For all numbers of rotations, including no rotations, a wait time of 200 µs is used between detection events, which is enough time to allow all rotations to finish. After one rotation, 94(3)% of initial states of $|01\rangle$ are measured to end up in $|10\rangle$, and 90(5)% vice versa. The deviations from 100% are attributed to limitations caused by the CCD camera's readout noise in ion state detection. At the time of measuring, the imaging system was not set up for efficient state readout using the CCD camera. Readout errors are caused by spontaneous decay to the S state during the 50 ms detection time (corresponding to a 4% decay probability after the first detection, and 8% after the second detection). As can be seen in Table 5.1, 0% of cases in which states are initially in $|01\rangle$ or $|10\rangle$ remain unchanged after one rotation. Therefore, despite detection errors, we can conclude that the success probability of the rotation is $100\%_{(-1\%)}$. The same conclusion can be drawn for multiple rotations in succession, as the success probability does not depend on the number of rotations. We can therefore say with confidence that the errors in Table 5.1 are due to detection errors, and not to rotation failures.

-Heating-

The changes in phonon number of the axial modes of the crystal are measured using sideband thermometry, as was the case in the final iterations of the rotation optimization procedure. The mean phonon number is measured for the axial common and stretch modes, after a number of rotations. Measuring the phonon number does not require ion dependent state readout, since the crystals motional state is shared by the two ions. State readout can therefore be done using the PMT, which does not resolve the state of the two ions independently, but has a notably lower detection time compared with the CCD camera. Readout errors due to spontaneous decay are less than 0.5% when using the PMT.

Before a rotation, the axial modes are sideband-cooled to an average phonon number of the common mode $\bar{n}_{com}=0.017(16)$ and the stretch mode $\bar{n}_{str}=0.011^{(+4)}_{(-5)}$. A number of rotations sequences, zero to six, are applied to the ion crystal, and the ion temperature is subsequently probed by resonantly exciting the blue and red sideband transitions of both modes, for varying probing times. The resulting excitation profile is compared to a model for two-ion sideband excitation for thermally distributed phonon levels¹⁵, by means of a least-squares fit. When fitting the experimental data to the sideband excitation model, the only free parameter is the mean phonon number. The coupling strength to the motional sidebands is measured independently by resonantly exciting a blue sideband without having applied any rotations and determining the Rabi frequency of the excitation. The mean phonon numbers of the axial common and stretch modes, \bar{n}_{com} and \bar{n}_{str} , after a number of rotations is shown in Figure 5.18(a). From a linear fit of mean phonon number as function of the number of applied rotations, the axial heating rates are determined to be $\Delta \bar{n}_{com} = 0.6^{(+3)}_{(-2)}$ phonons per rotation on the common mode, and $\Delta \bar{n}_{str} = 3.9(5)$ on the stretch mode. Stationary heating rates have been determined to be several phonons per second in our trap [106] and are therefore negligible in the 25 µs rotation duration.




Figure 5.18: Rotation induced heating: (a) Measured phonon numbers of the axial radial and stretch modes, after a number of rotations. (b) Rabi oscillations on an optical carrier transition, after a set of rotations.

To complement the rotation induced heating data, we monitor how the rotations affect the performance of on-resonant carrier excitation pulses, i.e. Rabi oscillations (see Section 2.2.3), which gives an indication of the fidelity of single-qubit gates after several rotations have been carried out. Excitation of a carrier transition using a light field whose wavevector is aligned with the ion crystal's axial direction results in an excited state population $p_{|1\rangle}^{(i)}$ of each ion i is given by:

$$p_{|1\rangle}^{(i)} = \sum_{n_{\text{com}}=0}^{\infty} \sum_{n_{\text{str}}=0}^{\infty} P_{\bar{n}_{\text{com}}}(n_{\text{com}}) P_{\bar{n}_{\text{str}}}(n_{\text{str}}) \sin^2\left(\frac{\Omega_{n_{\text{com}},n_{\text{str}}}t}{2}\right)$$
(5.34)

where $P_{\bar{n}}(n)$ is the occupation probability to have n phonons, given a mean phonon number \bar{n} . The occupation probability for a thermal distribution is given by

$$P_{\bar{n}}(n) = \frac{1}{\bar{n}+1} \left(\frac{\bar{n}}{\bar{n}+1}\right)^n. \tag{5.35}$$

 $\Omega_{n_{\text{com}},n_{\text{str}}}$ is a coupling constant (in the second order Lamb-Dicke approximation):

$$\Omega_{n_{\text{com}},n_{\text{str}}} = \Omega_{0,0} \left(1 - \eta_{\text{com}}^2 n_{\text{com}} - \eta_{\text{str}}^2 n_{\text{str}} \right), \tag{5.36}$$

with $\Omega_{0,0}$ the ground state carrier coupling constant, and η_{com} and η_{str} the mode-frequency dependent Lamb-Dicke factors, 0.061 and 0.047 in our particular setup.

Figure 5.18(b) shows Rabi oscillations, probed after a number of rotations. The displayed excitation is the mean of the excited state population of the two ions, $p_{|1\rangle}^{(1)} + p_{|1\rangle}^{(2)}$. The decay of the oscillation at a higher number of rotations is attributed to, and consistent with, the phonon numbers shown in Figure 5.18(a).

Heating in the stretch mode is attributed to phonon exchange during moments of degeneracy with the $X_{\rm str}$ and $Y_{\rm str}$ modes, as can be seen in Figure 5.17(b). While a local parameter optimum has been found with our optimization procedure, the resulting evolution of mode frequencies in Figure 5.17(b) suggests that a more stable solution, in which mode frequencies do not cross, might exist.

The efforts in reducing heating caused by ion crystal rotations stem from the requirement to minimize loss of qubit coherence, imperative when considering rotations for use as coherent operations in quantum computational sequences as a SWAP gate. The minimization routine results in rotation sequences that induce less than 1 phonon on the common axial mode. The following section discusses how qubit coherence is maintained after applying an optimized rotation sequence.

-Coherence-

Ramsey measurements, described in Section 2.3.5, are used to infer how well qubit coherence is preserved. For the two-ion coherence measurement, we first apply a global $\pi/2$ pulse, identical to the one used for determining swap success rate (Eq. 5.33), followed by a variable number of crystal rotations, and finalized by another global $\pi/2$ pulse with variable phase, which applies the operation

$$R_{\phi^{(1)},\phi^{(2)}}\left(\frac{\pi}{2}\right) = \exp(-i\pi/4(\sigma_{\phi^{(1)}}^{(1)} + \sigma_{\phi^{(2)}}^{(2)})),\tag{5.37}$$

with $\sigma_{\phi^{(i)}}^{(i)} = \cos(\phi^{(i)})\sigma_x^{(i)} + \sin(\phi^{(i)})\sigma_y^{(i)}$. Here, $\phi^{(i)}$ is a phase shift experienced by ion i between the initial and analysis pulse. This phase consists of a controlled phase change in the light field, ϕ_{var} , a phase evolution due to experimental drifts and fluctuations, ϕ_{drift} (for example, magnetic field and laser frequency), and a phase difference due to changes in the ions' position $\phi_{\text{pos}}^{(i)}$ caused by the rotation:

$$\phi^{(i)} = \phi_{\text{var}} + \phi_{\text{drift}} + \phi_{\text{pos}}^{(i)} \tag{5.38}$$

Scanning ϕ_{var} produces Ramsey fringes. The fringe contrast is an indication of qubit coherence, and can decrease if ϕ_{drift} is fluctuating. An additional source of contrast loss is fluctuations in coupling strength of the transition on which the Ramsey pulses are being applied. The mechanism in which motional heating leads to a decay in Rabi oscillations, as seen in the previous section, is therefore also responsible for loss of contrast in Ramsey measurements.

The Ramsey contrast, measured after various numbers of ion crystal rotations, is shown in Figure 5.19. For every number of rotations, including zero, the total wait time between the first and second Ramsey pulse is kept at 300 µs, so that ions experience the same dephasing caused by $\phi_{\rm drift}$, independent of the amount of rotations. The time of 300 µs ensures a sufficient settling time for filtered voltages and thus ion positions, also for six rotations. The contrast of 79% at zero rotations is attributed to CCD detection errors. Regardless of this initial contrast, we can infer loss of coherence from the decay in contrast for multiple rotations. Complementarily, we repeat the experiment with a PMT, which reduces detection errors to 0.2%. The PMT does not differentiate between the $|01\rangle$ and $|10\rangle$ states, thus a contrast measurement with the PMT does not account for non-identical phase differences for the two ions, and will in general result in loss of contrast if $\phi^{(1)} \neq \phi^{(2)}$. Since swapping the position of two ions generates an equal but opposite phase shift on both ions ($\phi_{pos}^{(1)} = -\phi_{pos}^{(2)}$), and therefore causes




Figure 5.19: Ramsey contrast after a set of rotations, measured using the PMT and CCD camera. The dotted lines are fit to the data as an exponential decay, from which a 0.2(2)% coherence loss per rotation is inferred.

a loss in contrast, we omit data for odd numbers of rotations 16 . The contrast of $98.8^{(+0.7)}_{(-1.0)}\%$ for zero rotations is consistent with previously measured Ramsey contrasts of our system, at $300\,\mu s$. After applying 6 rotations, a contrast of $98.4^{(+0.9)}_{(-1.2)}\%$ is maintained. From a weighted exponential decay fit, using both the CCD and PMT data, we infer a coherence loss of 0.2(2)% per rotation.

5.4.4.2 *Mixed species rotations: Ca - Sr*

Scalable quantum computation envisions making use more than one ion species. A second species can be used, for example, for sympathetic cooling or as an ancilla for state readout. In the conceptual usage of the QCCD architecture, ions transported along the trap chip may need to pass by an ion of a different species. Once again, an efficient solution is to use ion crystal rotations, but now of mixed-species crystals. As described in Section 5.4.3.1, the success of mixed-species rotations is more sensitive to changes in applied potentials, as fields unequally displace ions of different mass in the radial directions. In this section, we present findings of mixed-species ion crystal rotations. The goals and intentions of these rotations in terms of success rate, heating, and coherence, are the same as for single species. Unfortunately, as will be discussed in this section, we were unable to meet these goals, due to limitations of the setup. We conclude by suggesting setup improvements that will improve the quality of mixed-species rotations.

Applying the protocol described in Section 5.4.2 (the same as for Ca - Ca crystals, though with different masses in equation 2.20), suitable sequences of potentials that provide successful rotations are found. Determining the success rate of a rotation does not require detection of the internal states of the ions since both species fluoresce at different wavelengths. It is thus sufficient to monitor the position of the Ca ion. We opted to detect Calcium fluorescence with a PMT to obtain the data, since in our setup at the time our PMT's detection rate was much higher than that of the CCD camera. The two possible positions of the Calcium ion are resolved by placing a knife edge in an intermediate image plane of the detection optics, resulting in a different count rate depending on the position of the Ca ion. A threshold in PMT count rate thus distinguishes ion positions. The two possible ion positions produce count rates that are 3σ from the threshold, resulting in less than 0.2% detection error between the two states. The position is determined both before and after a rotation sequence.

The success of mixed-species rotations is inherently more sensitive to stray fields, notably in the radial direction in the plane of rotation. To characterize this sensitivity, an offset field can be intentionally applied during a rotation. The offset field is applied along the *x*-axis, the trap's radial direction parallel to the trap surface, which lies in the plane that the rotation takes place.

Figure 5.20 shows simulated rotation results, under the influence of different stray field strengths. Ion positions are calculated as function of applied potentials to determine whether a rotation is successful or not. Ions are initialized in the Sr - Ca configuration (top) or the Ca - Sr configuration (bottom). The left and right panels indicate a successful rotation in only one direction: regardless of the




Figure 5.20: Simulated Ca - Sr rotation trajectories. A radial offset field of $\pm 0.1 \, \mathrm{V} \, \mathrm{mm}^{-1}$, perpendicular to the rotation axis, disrupts rotations such that they work in only one direction.

initial configuration, the final configuration is Ca - Sr if a field is applied in the negative x-direction. Likewise, Sr - Ca is always the final configuration if a positive field is applied. If no offset field is applied, rotations are successful from both starting configurations.




Figure 5.21: Success of Ca - Sr crystal rotations under influence of a uniform external field E_x along the x-axis. The data points show the experimentally measured outcomes for a Strontium ion after a rotation. The bar above the plot indicates outcomes predicted by ion trajectory simulations.

This dependency of the mixed-species rotation on an offset field can be experimentally confirmed by adding voltages to the trap electrodes (identical to the voltage sets used for micromotion compensation) throughout the rotation sequence. The ion configuration is measured directly before and after a rotation sequence is applied. Figure 5.21 displays the measured outcome of rotations, as a function of the applied offset field, for each of three outcomes: 1) Sr starts on the left and stays on the left, 2) Sr starts on the right and stays on the right, and 3) both ions have switched positions after a rotation sequence. Above the plot, bars indicate outcomes as predicted by simulations, such as those shown in Figure 5.20, for varying offset fields. The measurement results are consistent with the simulated predictions.

The measurements in Figure 5.21 indicate a 100% success rate in bi-directional rotations (inferred from 100 repetitions of the sequence¹⁷) when the offset field is set to $-0.01\,\mathrm{V}\,\mathrm{mm}^{-1}$. A deviation of $\pm 0.013\,\mathrm{V}\,\mathrm{mm}^{-1}$ from this field reduces the success rate to 85(5)%. There is thus a window of less than

25 V m⁻¹ in which the field must be compensated to ensure that the rotations are successful in both directions.

The simulation is more forgiving, suggesting an acceptable range of about $50 \,\mathrm{V}\,\mathrm{m}^{-1}$. The deviation stems from the discrepancy between the simulated potentials and the potentials realized in the experiment. Regardless, $50 \,\mathrm{V}\,\mathrm{m}^{-1}$ is still a narrow window, from which we can conclude that careful calibration is imperative for successful mixed-species rotations. For comparison, for the same range of offset fields shown in Figure 5.21, a Ca - Ca rotation is always successful, according to simulations.

Similar as for single-species rotations, the rotation-induced heating is to be minimized. Unfortunately, after the rotation sequence, when exciting a carrier transition, there was no contrast in the Rabi oscillation. This behavior suggests a combined increase of phonon number of more than 80 quanta on the axial common and stretch modes. We have attempted to use the same machine-learning based optimization routine to reduce this value. However, even with conservative bounds on rotation parameters, the routine frequently led to expulsion of an ion (almost always Sr) from the trap. Due to the inconvenience of persistent reloading of ions, and the lack of optimizable contrast in measurements, we did not further pursue improving the rotation induced heating. We attribute our filter set-up to be a probable cause of rotation induced heating, both because it hinders us from exactly controlling motional frequencies during a sequence, as well as inducing excessive micromotion. Anharmonicity in the trap potential, especially as ions are further displaced from the RF null, may also contribute to undesired exchange of phonons between modes [66]. Anharmonicity may also lead to errors in determining correct potential sequences, as our method only takes into account harmonic potentials.

5.4.5 Summary and outlook

We have discussed a framework to generate the potentials required to rotate single- and mixed-species ion crystals confined in a linear trap. The spherical harmonic expansion provides a basis of potentials that are intuitive for controlling rotations. The experimental setup is characterized such that the applied voltage sequence corresponds to the desired potential sequence, taking into account the distortion of voltages due to the setup's electrical filters, and calibrations of the applied multipole potentials, described in Section 2.3.2. Sequences are further optimized by machine-learning based parameter optimization.

The machine-learning based algorithm was found to be a useful method for optimizing of the potential sequence to rotate a Ca - Ca crystal. The benefits of using machine learning are that relatively little needs to be known about the model that describes your system, and that a relatively simple neural network can model complicated functions. Also, little needs to be known about optimization procedures, since most of the heavy lifting is done internally by the chosen program that hosts the neural network (in the case of this section, Matlab was used, but the ability to train and use machine learning neural networks are ubiquitous across other data analysis tools, i.e. Python, Mathematica, R, etc.). Neural networks also have the benefit of being robust to noisy data¹⁸.

That being said, improvements can be made to the optimization protocol, to reduce both the final rotation induced heating and the experimental time required to find the optimum. The optimization results presented in this section required about an hour of data acquisition and analysis to produce the final rotation sequence used in Section 5.4.4. We expect that re-optimization on subsequent days, to compensate for drifts in stray electric fields, would likely take tens of minutes. In the scope of quantum computation, the duration of calibration and optimization of various experimental parameters must be kept to a minimum [214–216], as such operations can be seen as down-time for the computation. While the duration of our optimization procedure is not unacceptable, it could certainly do with some improvement. For example, in our algorithm, search bounds were set conservatively, to avoid accidentally searching in parameter regions that do not lead to rotations¹⁹. This restriction on search bounds negatively affects the total optimization duration and may limit the parameter space in which optima can be found. It is, for example, possible that in the results presented in 5.4.3.3 a local optimum was found, but not a global one. Search bounds could be set to be dynamic throughout the

¹⁸ In fact, neural networks often flourish over other models when dealing with noisy data. To give a classic example: neural networks are used for written text recognition and are able to decipher my handwriting.

¹⁹ The 'do-nothing' sequence would certainly optimize the amount of heating induced, but poorly affects the rotation success rate.

optimization: initially, a broader search range would aid in ensuring that a global optimum is found, whereas later on in the optimization, narrower search bounds will allow the algorithm to quickly hone in on an optimum. Instead of neural networks, alternative optimization algorithms could be employed, such as the DIRECT algorithm [217], which strikes a compromise between exploring a broad, bounded parameter space and focusing on converging on likely parameter candidates. Additional checks would be required to ensure that the sequence still produces a rotation. Additionally, care must be taken to avoid parameter regions in which the voltage sequences expel ions from the trap.

In the final stages of optimization, we focused on minimizing heating on the axial *common* mode and were able to reduce this value below a phonon per rotation of a Ca - Ca crystal. However, the *stretch* mode phonon occupation increases by several quanta per rotation. Upon analysis of mode frequencies of our optimized rotation sequence, one sees that the $X_{\rm str}$ and $Z_{\rm str}$ modes cross several times. The exchange of phonons between these modes is a dominant source of heating and is reduced by cooling the $X_{\rm str}$ mode before a rotation sequence. Ideally, one would want to cool all modes close to their ground states. Here, electromagnetically induced transparency (EIT) cooling [84] and polarization gradient cooling [85, 218] would be suitable approaches. Additionally, the optimization algorithm could be adjusted to attempt to reduce heating on both the common and stretch mode, which likely results in sequences in which relevant motional modes do not cross.

Parameter optimization was required in our setup because of a discrepancy between the intended potential sequence and the actual potentials that the time-dependent voltages on the trap electrodes produce. The largest source of discrepancy is expected to be the electrical filter model. Including a primitive filter model was absolutely necessary to produce voltage sequences that allow rotations to be successful, but the model is expected to be imperfect. In this model, resistances, inductances, and capacitances of all electrode connections are assumed to be identical, and tolerances in individual components are ignored. Self-capacitance, self-inductance and mutual capacitance between channels are neglected. The slew-rate of the voltage source is modeled as an additional resistance, which oversimplifies the device's effective transfer function. It is therefore probable that voltages applied to the trap electrodes deviate from the intended ones.

Unfortunately, it is notoriously difficult to directly measure the filter transfer function of the trap electrodes in-situ. One could conceive measurements where trapped ions are used as tools to measure the filter function (for example, one could measure time-dependent spectra of motional sidebands upon applying a step-function in electrode voltages). Improved filter models would result in predicted initial rotation sequences with less motional heating. This benefits mixed-species rotations, as voltage sequences could be generated within a regime where rotation induced heating is low enough such that sequences can be optimized, without the risk of ion expulsion.

Despite experimental obstacles such as limited knowledge of the filter transfer function, rotations of two Ca - Ca qubits in a qubit superposition are achieved with nearly no loss of qubit coherence. Furthermore, the rotation outperforms an alternative exchange of qubit information that can be realized with a quantum information protocol that leaves the physical ions in place, a so-called SWAP gate. One possible way to achieve a SWAP gate is to apply three alternating CNOT gates, as shown in Figure 5.22. With the usual trapped-ion experimental toolbox, a CNOT operation can be generated using a MS-gate and a set of local and global Pauli operations, as shown in Figure 5.22. A SWAP operation cannot be attained with less than 3 MS gates²⁰. Given the currently attainable entangling fidelity of $\sim 99.9\%$ [26, 219], a SWAP gate infidelity of $1 - (0.999)^3 = 0.3\%$ can be inferred, in the most optimistic scenario where fidelities of the Pauli operations are assumed to be unity. The Ca - Ca rotations in this section produce an identical operation with a fidelity of 0.2(2)%, which is on-par with state-of-the-art gate fidelities. The physical rotations are, however, achieved on time-scales of tens of microseconds, about an order of magnitude lower than the SWAP gate.

In this section, methods for achieving rotations have been studied within the context of quantum information processing. Ion crystal rotations are a fundamental component of the proposed QCCD architecture. The presented methods of generating and optimizing sequences are applied to ion crystal rotations on planar traps, but can easily be generalized to other architectures and other physical operations, further expanding the toolbox of physical ion manipulation for quantum information processing.

²⁰ The total number of local Pauli operations can be reduced when concatenating the three CNOT gates.




Figure 5.22: A SWAP gate requires three entangling gates (MS $_{\chi}(\pi/4)$) and a set of local Pauli operations.

CONTROL OF QUBITS WITH DISSIPATION THROUGH ENGINEERED RESONANCE

After all these implements and texts designed by intellects / So vexed to find, evidently there's still so much that hides.

— The Shins, Saint Simon

The commonly used universal set of gate operations involved in trapped ion quantum computation (QC) and quantum information processing (QIP), introduced in Section 2.2.7, are *unitary* operations that control the quantum state of a collection of qubits. Using the Bloch sphere picture presented in Section 2.2.1, if a Bloch vector points to the surface of the Bloch sphere, it will remain to do so after application of the unitary operation. Additionally, the unitary operation applied to the full set of points that lie on the surface of the Bloch sphere, maps to the same set with a one-to-one relation (i.e. a unitary operation is bijective). As such, unitary operations are reversible, and they preserve the *purity* and *coherence* of a quantum state (see Section 2.2.5).

In the context of QC, purity and process fidelity go hand-in-hand: a loss of purity leads to a loss in fidelity and thus leads to computational errors. It therefore makes sense that one would attempt to keep a qubit's purity as high as possible throughout a computational sequence. Indeed, one of the outstanding challenges in high-fidelity quantum computation is minimizing sources of noise that lead to qubit decoherence, thus a loss of purity.

However, a new paradigm of quantum algorithms is emerging, which stands in stark contrast with the archetypical idea that quantum algorithms should have unitary evolution. Such novel algorithms combine unitary processes with non-unitary operations [220], which are operations which may be irreversible and may produce mixed states. For example, they are used to generate low-temperature thermal states [221, 222], or for simulation of quantum systems that interact with an open environment [223, 224]. Such algorithms can be realized with controlled non-unitary gates, whose operation currently cannot be simply implemented using the standard set of qubit manipulation tools (such as the ones discussed in Section 2.2.7).

One of the most imperative applications of non-unitary operations in quantum computation is quantum error correction (QEC) [45, 225, 226]. In QEC, several *physical* qubits (individual ions) together encode a single *logical* qubit. The encoding of logical qubits corresponds to a specific subset of the full Hilbert space that the physical qubits span. The encoding is chosen such that if an error occurs (for example, one qubit's state changes in phase or in population), the logical qubit leaves this subset. The goal of QEC is to identify whether the state is in the logical subset, and, if not, apply appropriate gate operations to return it to the correct state. The QEC operation recovers the original error-free qubit state by mapping the erroneous state within a broad Hilbert space and compressing it to a smaller subset, which makes the operation inherently non-unitary and irreversible.

Up until now, the experimental tools for manipulating qubits do not allow us to realize such non-unitary operations for QEC directly in the quantum system. Instead, in many of the proposed and demonstrated QEC schemes [49, 227–232], the irreversible part of the error correction operation is outsourced to a classical computer. In such schemes, information about whether or not an error has occured, an error syndrome, is mapped to an ancilla qubit, is measured through state detection of that ancilla and forwarded to a classical computer. The computer then assesses this information, and decides which, if any, correction sequence is required to fix possible errors. This method of error correction, shown on the left side of Figure 6.1, is time consuming, requires multiple entangling operations and has yet to demonstrate fault-tolerance, a regime where QEC realistically solves more problems than that it introduces.

QEC still has a long road ahead before a qubit can be fully considered to be protected from errors, with many different strategies being investigated [225]. One avenue of investigation is to see if the




Figure 6.1: Conceptual diagrams of quantum error correction methods, showing the conventional syndromereadout approach, which relies on classical feedback, and the dissipative state transfer approach discussed in this chapter.

classic feedback mechanism can be replaced by a fully automatized conditional operation that exists entirely within the quantum processor [233]. The conditional decision making process is then not carried out by an external classical computer but is made by the quantum system, severely lightening the classical overhead. This implies that the current set of gates used within quantum circuits is to be extended to include non-unitary operations.

In this chapter, we introduce a novel experimental technique that prospectively makes this extension of the quantum computational toolbox to include controlled non-unitary operations. We refer to the experimental technique presented in this chapter as *dissipative state transfer* through *engineered resonance*. The term refers to the concatenation of two distinct methods:

- **Engineered resonance** is the application of a unitary operation that acts conditionally on the initial state of multiple qubits.
- **Dissipation** is a non-unitary operation that makes the state transfer due to engineered resonance irreversible.

Combining these two methods opens up new possibilities of qubit control, including the potential of realizing the goal of integrating QEC fully in the quantum processor. The focus of this chapter is to introduce the methods of controlling dissipation through engineered resonance and demonstrate their experimental feasibility.

In Section 6.1, the primary mechanism of engineered resonance is introduced. Section 6.2 describes how this mechanism can be used to realize quantum error correction.

For the purpose of QEC, practical application of dissipation through engineered resonance still faces many challenges and obstacles. The goal of the rest of this chapter, Section 6.3, is an experimental demonstration of the viability, usability, and assessment of challenges of these tools. They are used to execute another type of non-unitary operation: boolean logic gates on a quantum platform. Here, we merge the unbelievably successful paradigm of classical computation (whose functionality heavily relies on irreversible gates) with that of quantum computation (whose functionality heavily relies on reversible gates). Although future directions of the implementation of such a classical-quantum-hybrid system are as of yet unclear, our primary goal is to use such a system as a demonstrator for the novel tools presented in this chapter.

6.1 Dressed state splitting for dissipation through engineered resonance

The non-unitary operations discussed throughout this chapter are conditional operations: from a set of input states, only a specific selection is to undergo state transfer, while others are to be left unchanged. We achieve this conditional transfer by using *dressed states*. In this section, dressed states are introduced, followed by a description of how they can be used to enable dissipative state transfer through engineered resonance.

6.1.1 Dressed state splitting of one ion

The general concept of the engineered resonance scheme is to generate a set of states known as dressed states, whose energy levels can be precisely controlled. This control of energy levels allows us to adjust which light-field frequency is required for resonant population transfer between these states. As we see later in this section, we can use these engineered resonances to conditionally allow chosen transitions to occur, while suppressing others.

We first consider an ion with three distinct electronic states, denoted $|0\rangle$, $|f\rangle$, $|1\rangle$. A light field is applied that couples the states $|0\rangle$ and $|f\rangle$, with coupling strength Ω_f and detuning Δ_Ω . This light field is denoted as the *probe* beam. Simultaneously, a second light field couples the states $|1\rangle$ and $|f\rangle$, with coupling strength g_f and detuning Δ_g (the *pump* beam). These states, coupling strengths, and detunings are schematically shown in Figure 6.2(a).




Figure 6.2: (a) Diagram of three-level scheme (b) Excitation from the initial state $|0\rangle$ in a three-level system after time $t=\pi/\Omega_f$, with probe power Ω_f , and pump power $g_f=10\Omega_f$. The plot on the right is a cross section of the image plot, at $\Delta_g=0$, which displays dressed-state splitting.

The most general state of the ion $|\phi(t)\rangle$ is given by

$$\psi(t) = c_0(t) |0\rangle + c_f(t) |f\rangle + c_1(t) |1\rangle \tag{6.1}$$

with time-dependent complex weights obeying $|c_0(t)|^2 + |c_f(t)|^2 + |c_1(t)|^2 = 1$. Given a known initial state $|\psi(0)\rangle$, the time-dependent solution of the state $|\psi(t)\rangle$ can be obtained with $i\hbar\partial |\psi(t)\rangle / \partial t = H |\psi(t)\rangle$, with the Hamiltonian

$$H = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_f & 0 \\ \Omega_f & 2\Delta_{\Omega} & g_f \\ 0 & g_f & 2\Delta_{\Omega} - 2\Delta_g \end{pmatrix}. \tag{6.2}$$

Figure 6.2(b) shows an example of excitation when this Hamiltonian acts on the initial state $|0\rangle$, given by the population $1-P_{|0\rangle}=1-|c_0|^2$, for various detunings Δ_Ω and Δ_g . In this example, we have used $g_f=10\Omega_f$, and $t=\pi/\Omega_f$. There is no simple analytic expression for $|\psi(t)\rangle$, but a few limiting cases can be observed. Naturally, "switching off" the pump beam, $g_f=0$, reproduces the Rabi cycles as function of detuning Δ_Ω of Eq. 2.45, for starting state $|\psi(0)\rangle=|0\rangle$. The same behavior would be observed even if the pump was not off but had a large detuning $|\Delta_g|\gg g_f$: The detuning at which the probe beam excites the ion asymptotically approaches $\Delta_\Omega=0$ with increasing detuning Δ_g . This asymptote is denoted by the horizontal dashed line in 6.2(b).

Another interesting case occurs when the pump and probe detunings are equal, $\Delta_{\Omega} = \Delta_g$. If $|\Delta_g| \gg g_f$, excitation to $|f\rangle$ is almost fully suppressed, and instead state population cycles between $|0\rangle$ and $|1\rangle$ with an effective Rabi frequency of $\Omega_f g_f/(2\Delta_{\Omega})$. The state $|f\rangle$ mediates population transfer, but is not populated itself. This type of excitation is known as a Raman transition, and is commonly

used for operations on a ground-state qubit (i.e., whose qubit is encoded in the Zeeman sub-levels of the atomic ground state) using optical transitions. This particular resonance at $\Delta_{\Omega} = \Delta_g$, valid when $|\Delta_{\Omega}| \gg g_f$, is shown as an asymptote by the diagonal dashed line in Figure 6.2(b).

If the two mentioned resonance conditions are simultaneously met, $\Delta_g = \Delta_\Omega = 0$, an avoided crossing is observed. If we assume that the pump is much stronger than the probe, $g_f \gg \Omega_f$, and is on resonance, $\Delta_g = 0$, excitation occurs at detunings of $\Delta_\Omega = \pm g_f/2$, as indicated in the right plot in Figure 6.2(b). The original condition that excitation would occur at zero detuning of the probe beam, Δ_Ω , has apparently *split* into two resonance conditions. At these particular detunings, an ion's state oscillates between $|0\rangle$ and $1/\sqrt{2}(|f\rangle \mp |1\rangle)$, with an effective Rabi frequency of $\Omega_{\rm eff} = \Omega_f/\sqrt{2}$. The excited states $1/\sqrt{2}(|f\rangle \mp |1\rangle)$ are known as dressed states, and their changes in energy with respect to the bare states is known as dressed state splitting.

The above may feel like a very hand-waving and perhaps somewhat unsatisfying explanation of dressed state splitting. A more formal description is provided in Appendix c.

6.1.2 *Two ion dressed state splitting*

In the previous section, the resonance condition for the detuning of the probe beam for excitation between two levels has been altered due to strong coupling to a third level. This effect plays a fundamental role in resonance engineering, by which we mean that an appropriately detuned light field may or may not excite an ion's electronic state, depending on the state of one or more other ions. The ion whose state is to be conditionally changed is referred to as the target qubit. Information about the state of multiple qubits must be shared between ions to be able to apply a conditional operation on the target qubit. In the scheme presented below, one of the ions' shared motional mode is used for this exchange of qubit information.

To expand the single-ion dressed state excitation of the previous section, we now introduce a second ion, with the same three electronic levels as the first. Both ions are prepared in the motional ground state $|0\rangle_n$. The probe beam is only applied to the target qubit, chosen to be the first of the two. The pump beam is applied to both ions simultaneously. Crucially, this beam is now set to be a red sideband transition, such that it couples the states $|1\rangle |n+1\rangle_n \leftrightarrow |f\rangle |n\rangle_n$ of both ions, with an interaction strength $g_f\sqrt{n+1}$.

We can now analyze what happens to the ions in each of the four possible qubit states, $|00\rangle|0\rangle_n$, $|01\rangle|0\rangle_n$, $|10\rangle|0\rangle_n$, $|11\rangle|0\rangle_n$, though from here on occasionally omitting the motional occupation notation $|n\rangle_n$ when the ions are in the motional ground state, $n=0^1$. The left ket is the shorthand notation for the tensor product of the first and second ion, $|ij\rangle=|i\rangle_1\otimes|j\rangle_2$. Since the initial states are all in the motional ground state, they initially do not couple with the red sideband pump beam.




Figure 6.3: Level schemes for dressed state splitting from starting states (a) $|00\rangle$ and (b) $|01\rangle$, with the probe applied only to ion 1. (c) Excitation from initial state $|01\rangle$ with settings identical to those in Figure 6.2. The interpretation of the dashed lines is also identical. At $\Delta_g=0$, resonance occurs at $\Delta_\Omega=0,\pm\sqrt{2}g_f/2$, shown in the plot on the right.

We once again analyze what occurs to each of the initial states. The probe beam only couples $|0\rangle$ and $|f\rangle$ of the *first* ion, so the states $|10\rangle$ and $|11\rangle$ remain unchanged. Figure 6.3(a) and (b) schematically show level diagrams with the possible transitions that can occur from the other two starting states $|00\rangle$ and $|01\rangle$. Let us first consider the initial state $|00\rangle$, in Figure 6.3(a). The dynamics that emerge are identical to the single ion case in the previous section, where the basis $\{|0\rangle$, $|f\rangle$, $|1\rangle$ is replaced with $\{|00\rangle|0\rangle_n$, $|f0\rangle|0\rangle_n$, $|10\rangle|1\rangle_n$. The Hamiltonian that describes the dynamics of this basis is identical to the one in Eq. 6.2. Therefore, similar to the single ion case, we have

$$g_f \gg \Omega_f$$
, $\Delta_g = 0$, $\Delta_\Omega = \frac{g_f}{2}$: $|00\rangle |0\rangle_n \leftrightarrow \frac{1}{\sqrt{2}} (|f0\rangle |0\rangle_n - |10\rangle |1\rangle_n) \equiv |p_{00}\rangle$. (6.3)

where $|p_{00}\rangle$ is one of the dressed states generated by applying the pump beam. Note that the Rabi cycling on the transition $|00\rangle\,|0\rangle_n \leftrightarrow |p_{00}\rangle$ is valid by approximation when $g_f\gg\Omega_f$, and is exact in the limit $g_f\to\infty$.

Starting from the state $|01\rangle |0\rangle_n$, four different states can be reached, as shown in Figure 6.3(b). Using the basis $\{|01\rangle |0\rangle_n$, $|f1\rangle |0\rangle_n$, $|11\rangle |1\rangle_n$, $|1f\rangle |0\rangle_n\}$, the Hamiltonian that describes the time evolution is given by

$$H = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_f & 0 & 0 \\ \Omega_f & 2\Delta_{\Omega} & g_f & 0 \\ 0 & g_f & 2\Delta_{\Omega} - 2\Delta_g & g_f \\ 0 & 0 & g_f & 2\Delta_{\Omega} \end{pmatrix}.$$
(6.4)

Excitation from the initial state $|01\rangle\,|0\rangle_n$ is shown in Figure 6.3(c) for various detunings, using $g_f=10\Omega_f$, at time $t=\pi/\Omega$. The plot on the right shows a cross section of the image plot at $\Delta_g=0$. There are two notable differences from the excitation curve for the $|00\rangle$ state (see Figure 6.2): First, $\Delta_\Omega=0$, corresponding to the horizontal line, always results in excitation, even at $\Delta_g=0$. Second, and more importantly, the dressed state splitting at $\Delta_g=0$ results in resonance at $\Delta_\Omega=\pm g_f/\sqrt{2}$, which is a factor $\sqrt{2}$ further detuned than was the case for $|00\rangle$ as an initial state. This enhanced splitting is a result from the additional pump transition. The two pump transitions constructively interfere and produce an enhancement of the dressed state splitting². Under the conditions $\Delta_\Omega=\sqrt{2}g_f/2$ and $g_f\gg\Omega_f$, Rabi cycling occurs between the states

$$|01\rangle |0\rangle_n \leftrightarrow \frac{1}{2} \left(-\sqrt{2} |11\rangle |1\rangle_n + |f1\rangle |0\rangle_n + |1f\rangle |0\rangle_n \right) \equiv |p_{01}\rangle, \tag{6.5}$$

with an effective Rabi frequency of $\Omega_f/2$ (see Appendix c.2). As before, the indicated Rabi cycling transition is an approximation and is exact when $g_f \to \infty$.

The crucial insight is that the condition for resonant excitation is different for the starting states $|00\rangle$ and $|01\rangle$. We thus have a method of state transfer of one ion, conditioned on the state of the second by selecting the appropriate detuning of the probe beam. We refer to this conditional state transfer as engineered resonance.

6.1.3 Dissipation

The methods presented in the previous section enable resonant Rabi cycling between an initial state and a *manifold* of states, the dressed states $|p_{00}\rangle$ or $|p_{01}\rangle$. The goal, however, is non-unitary transfer of population to a single chosen bare state. This implies that an irreversible process must occur, for which spontaneous decay is a straightforward choice. Sideband cooling (Section 2.2.8) is well-suited for this purpose: it uses spontaneous decay to dissipatively remove motional excitation, and it acts

² An intuitive but possibly inaccurate way of describing this is that the second pump transition splits the already-split state, forming a total of four dressed states, two of which with higher detuning. The other two interfere with one another to form a single dressed state at zero detuning, which could promiscuously be referred to as un-dressing.

only on states that have non-zero motional mode occupation, which is the case for one of the states in both excited dressed states, $|p_{00}\rangle$ and $|p_{01}\rangle$.

One might consider directly sideband-cooling the ions that encode the qubits (see Section 2.2.8). However, qubit coherence cannot be maintained if they undergo spontaneous decay, the relevance of which becomes clear later in this chapter, in Section 6.2, where its application for QEC is discussed. The solution then, is to introduce a second ion species that shares the motional motion, and can thus be sideband cooled without affecting the state of the qubit species. This approach, known as *sympathetic* cooling [41, 42, 130] is a common method to combat heating without affecting qubits during a quantum computation sequence, and is introduced here as a method of dissipation.

Sideband cooling transfers the state $|10\rangle |1\rangle_n$ and $|11\rangle |1\rangle_n$ to $|10\rangle |0\rangle_n$ and $|11\rangle |0\rangle_n$. Neither of these final states are influenced by the probe or pump beams, so any further dynamics cease once either state is reached. If sideband cooling is applied simultaneously with the engineered resonance Rabi cycling described previously, we ideally have the conditional non-unitary population transfer:

$$|00\rangle \rightarrow |10\rangle \quad \text{iff} \quad \Delta_{\Omega} = \pm \frac{g_f}{2}$$
 (6.6)

$$|00\rangle \rightarrow |10\rangle$$
 iff $\Delta_{\Omega} = \pm \frac{g_f}{2}$ (6.6)
 $|01\rangle \rightarrow |11\rangle$ iff $\Delta_{\Omega} = \pm \frac{g_f}{\sqrt{2}}$.

Experimental limitations may inhibit the success of these transfers. In particular, we see that the peaks in Figures 6.2 and 6.3 have a finite width. This means that in reality the non-unitary population transfers occur at a broader range of probe detuning Δ_{Ω} than indicated in Eqs. 6.6 and 6.7, and may be broad enough that the two conditions overlap. Therefore, if a detuning of $\Delta_{\Omega} = g_f/2$ is chosen with the intention of exciting $|00\rangle$ as in Eq. 6.6, the state $|01\rangle$ may still be off-resonantly excited, which we refer to as cross-talk. This cross-talk is reduced by increasing the ratio of coupling strengths of the pump and probe beams, g_f/Ω_f : Loosely speaking (to be discussed less loosely and more quantitatively in Section 6.3.3), the pump coupling g_f dictates the separation between resonance peaks (which should be maximized to reduce cross-talk), and the probe coupling Ω_f dictates the width of the peaks (which should be minimized). Experimental constraints limit how far the ratio g_f/Ω_f can be pushed: finite laser power sets an upper limit on g_f , and setting a low Ω_f lowers the effective Rabi oscillation frequency, thus extending the sequence duration. At longer sequence times, errors due to limited coherence times and motional heating increase.

6.2 DISSIPATION THROUGH ENGINEERED RESONANCE FOR QUANTUM ERROR CORRECTION

This previous section has introduced the mechanisms of engineered resonance, the concept that population transfer on one ion can be achieved, conditioned on the state of a second ion. Dissipation ensures that this process is unidirectional. This forms the basis for conditional classical boolean logic gates, presented in Section 6.3. First though, we will discuss the prospective application of dissipation through engineered resonance.

The fact that the engineered resonance mechanisms are *conditional* opens the door to autonomous quantum error correction (QEC). The core concept of QEC is to use a set of physical qubits to encode a logical qubit redundantly. The occurrence of an error, for example a bit-flip of one of the physical qubits, results in the system leaving the logical subspace, which is designed to be a detectable event. As a basic illustrative example, using the classical repetition code [234–236], we could define a logical qubit as

$$|\phi\rangle_L = c_0 |0\rangle_L + c_1 |1\rangle_L = c_0 |000\rangle + c_1 |111\rangle.$$
 (6.8)

If a single spin flip were to occur on the first ion, the correct logical state can be uniquely reconstructed from the error state $|\phi\rangle_E = c_0 |100\rangle + c_1 |011\rangle$. Naively, one would do a majority-vote check to determine if one of the three physical qubits is an outlier, and, if so, perform a bit-flip on that physical qubit to correct it. Of course, attempting to determine if an error occurred on a logical qubit by directly detecting its state is problematic, since this action collapses the state. The prevailing strategy to circumvent this obstacle is by first mapping information about the parity of the qubit states, known as an error syndrome, onto an ancilla qubit and measuring that syndrome. As was discussed in the beginning of this Chapter, highlighted in Figure 6.1, this is a time-consuming and error-prone method, requiring multiple entangling gates, and classical analysis and decision-making. Dissipation through engineered resonance has the prospect of simplifying this procedure, as it avoids the necessity for state readout of ancilla qubits altogether, instead using population transfer conditioned on the state of multiple ions to uni-directionally return a faulty state to the logical subspace.

As of yet, there are no experimental demonstrations of using engineered dissipation for QEC, owing to the fact that there are multiple technical challenges to overcome, as will be discussed in Section 6.3.3. Therefore, instead of jumping into a fully-fledged error correction scheme, it is sensible to simplify the experimental methods and focus on demonstrating the core concepts of dissipative QEC. This type of proof-of-principle demonstration is discussed in the following section, where we introduce a minimal-instance scheme for dissipative QEC. Afterwards, in Section 6.2.2, we build up from the concepts used in the minimal instance scheme, and introduce the full scheme for QEC, as proposed by F. Reiter et al. [233].

6.2.1 Minimal instance dissipative error correction

A logical qubit is encoded using two physical qubits as

$$|\phi\rangle_I = c_0 |01\rangle + c_1 |10\rangle. \tag{6.9}$$

In this minimal instance scheme we make the assumption that the second qubit is immune to errors, and only the first qubit is susceptible to possible erroneous spin flips. This admittedly unreasonable assumption is why this scheme is considered to be a proof-of-principle and not a full QEC protocol.

The occurrence of a spin-flip error shifts population to the state $|\phi\rangle_E$, given by

$$|\phi\rangle_E = c_0 |11\rangle + c_1 |00\rangle. \tag{6.10}$$

leaving the logical qubit in a mixed state consisting of the logical and error qubit, $|\phi\rangle_L$ and $|\phi\rangle_E$. We denote the basis states with negative parity, $|01\rangle$ and $|10\rangle$, as odd states, and those with positive parity as even. In this example, a logical qubit consists of odd states. The even states are faulty and require correction. The process that enables the non-unitary transfer $|\phi\rangle_E \to |\phi\rangle_L$ is discussed below.




Figure 6.4: Level scheme for minimal instance dissipative error correction

The two-ion dressed-state splitting model of Section 6.1.2 is applied here. The level scheme and transitions are shown in Figure 6.4. As before, a weak probe beam couples the $|0\rangle$ and $|f\rangle$ of only the first ion, $|0\rangle_1 \leftrightarrow |f\rangle_1$, with coupling strength Ω_f . Also as before, a strong pump beam couples the red sideband of the $|1\rangle \leftrightarrow |f\rangle$ transition of both ions and has a coupling strength g_f when the transition involves the motional ground state, $|1\rangle |1\rangle_n \leftrightarrow |f\rangle |0\rangle_n$. The pump beams are both set to be on resonance with their respective transitions.

Extending the dressed-state splitting model of Section 6.1.2, two additional light-fields are introduced, along with an additional auxiliary level $|e\rangle$. A second probe beam with coupling strength Ω_e is applied

to the first ion with detuning Δ_{Ω}^e , which couples $|1\rangle_1 \leftrightarrow |e\rangle_1$ (The detuning of the original probe beam, with coupling Ω_f , is now relabeled as Δ_{Ω}^f). A second pump beam is applied on resonance with the red sideband of the $|0\rangle$ and $|e\rangle$ transition of both ions, with coupling strength g_e when $|e\rangle$ is in the motional ground state.

It is once again instructive to analyze which excitation can occur for each of the initial states, $|00\rangle\,|0\rangle_n,\,|01\rangle\,|0\rangle_n,\,|10\rangle\,|0\rangle_n,\,$ and $|11\rangle\,|0\rangle_n$. The pump beams are red sideband transitions and cannot excite these states as they are in the motional ground state. The probe beams are only applied to the first ion: The probe with coupling Ω_f acts on $|00\rangle\,|0\rangle_n$ and $|01\rangle\,|0\rangle_n$, while the probe with coupling Ω_e acts on $|10\rangle\,|0\rangle_n$ and $|11\rangle\,|0\rangle_n$. The level schemes depicting these four cases are shown in Figure 6.5(a), with the initial states highlighted yellow. In each case, after excitation from the initial state by the probe beam, further transitions are made possible because of interaction with the pump beams. These transitions and the respective states are also shown in Figure 6.5(a), where it can be seen that the possible transitions of each initial state form a four-level system. For example, the initial state $|10\rangle\,|0\rangle_n$ is excited by one of the probe beams to $|e0\rangle\,|0\rangle_n$, which can undergo a red sideband transition by a pump beam to $|00\rangle\,|1\rangle_n$ and from there to $|0e\rangle\,|0\rangle_n$. One might note that the probe would also enable excitation from $|00\rangle\,|1\rangle_n$ to $|e0\rangle\,|1\rangle_n$, but we see later that the probe's detuning Δ_Ω^e is chosen such that it is off resonance from the $|1\rangle_1\,|n\rangle_n\leftrightarrow|e\rangle_1\,|n\rangle_n$ transition with the special exception of n=0.




Figure 6.5: (a) Level schemes for minimal instance dissipative error correction, showing coupled levels for all four possible starting states (highlighted). (b) Excitation from initial states, as function of probe detuning and second pump coupling strength g_e , using $g_f/10 = \Omega_f = \Omega_e$. Choosing a probe detuning of g_f (vertical dashed lines), and a pump power of $g_e = \sqrt{3}g_f$ (horizontal dashed lines) enables resonant population transfer from initial states $|00\rangle$ and $|11\rangle$, while transfer from the other initial states is suppressed. (c) Excitation profiles for all starting states, corresponding to the horizontal dashed lines in (b).

We note the similarity of these level schemes of Figure 6.5(a) to that in Figure 6.3(b), where we discussed two-ion dressed state spliting. In fact, the dynamics of initial states $|01\rangle$ and $|10\rangle$ are identical to those in Figure 6.3(b). Therefore, using the same treatment as in Section 6.1.2, we see that resonant

excitation from these initial states occurs when $\Delta_{\Omega}^f = g_f/\sqrt{2}$ and $\Delta_{\Omega}^e = g_e/\sqrt{2}$, for $|01\rangle$ and $|10\rangle$, respectively. On the other hand, for the initial states $|00\rangle$ and $|11\rangle$, both g_f and g_e contribute to dressing the excited manifold. The resulting dressed state splitting resonantly excites the initial states when at detunings $\Delta_{\Omega}^f = \Delta_{\Omega}^e = \sqrt{g_f^2 + g_e^2}/2$ (see Appendix c.3 for more details).

Consequently, if the coupling strengths of the two pump beams are chosen to be unequal, $g_f \neq g_e$, the resonance conditions for the even starting states, $|00\rangle$ and $|11\rangle$, are unequal to those of the odd starting states. This is exemplified in Figure 6.5(b), where we show how the resonance condition for excitation from the four initial states depends on the ratio of coupling strengths, g_e/g_f . As an example for this figure, we have set $g_f = 10\Omega_f$. Upon choosing $g_e = \sqrt{3}g_f$, denoted by the horizontal dashed line, and $\Delta_{\Omega}^e = \Delta_{\Omega}^f = \sqrt{g_f^2 + g_e^2}/2$, denoted by the vertical dashed line, we see that resonant transfer is enabled from the even states and suppressed from the odd states. For the example in Figure 6.5(b), the ratio between pump powers, $g_e/g_f = \sqrt{3}$ is chosen arbitrarily but is a reasonable experimental choice because it allows for resonance conditions to be spectrally separated while optimizing available beam power.

The various resonance conditions are summarized and emphasized in Figure 6.5(c), which shows excitation profiles for all initial states. The same parameters as in Figure 6.5(b) are used, though with a fixed value of $g_e = \sqrt{3}g_f = 10\sqrt{3}\Omega_f$, corresponding to the horizontal dashed line.

In the limit that the probe beams are much weaker than the pump beams, Ω_f , $\Omega_e \ll g_f$, g_e , if the probe detunings are set to $\Delta_{\Omega}^f = \Delta_{\Omega}^e = \sqrt{g_f^2 + g_e^2}/2$, Rabi cycling occurs between the states

$$|00\rangle |0\rangle_n \leftrightarrow a_{f0} |f0\rangle |0\rangle_n + a_{1e} |1e\rangle |0\rangle_n + a_{10} |10\rangle |1\rangle_n \equiv |\psi_{10}\rangle,$$

$$|11\rangle |0\rangle_n \leftrightarrow a_{e1} |e1\rangle |0\rangle_n + a_{0f} |0f\rangle |0\rangle_n + a_{01} |01\rangle |1\rangle_n \equiv |\psi_{01}\rangle$$
(6.11)

where we have denoted the dressed states that are accessible from $|00\rangle |0\rangle_n$ and $|11\rangle |0\rangle_n$ as $|\psi_{10}\rangle$ and $|\psi_{01}\rangle$, respectively. The state coefficients are given by

$$a_{f0} = a_{0f} = \frac{g_f}{\sqrt{2}\sqrt{g_f^2 + g_e^2}} \tag{6.12}$$

$$a_{e1} = a_{1e} = \frac{g_e}{\sqrt{2}\sqrt{g_f^2 + g_e^2}} \tag{6.13}$$

$$a_{01} = a_{10} = \frac{1}{\sqrt{2}} \tag{6.14}$$

These Rabi cycles oscillate with an with an effective Rabi frequency of

$$\Omega_{\rm eff} = \frac{\Omega_i g_i}{\sqrt{2} \sqrt{g_f^2 + g_e^2}} \tag{6.15}$$

with i = f for initial state $|00\rangle$ and i = e for $|11\rangle$. Excitation from the odd states, $|01\rangle$ and $|10\rangle$ are suppressed. See Appendix c.3 for a derivation of dressed states, resonant detunings, and the effective Rabi frequency.

As in Section 6.1.3, this transfer process is made non-unitary by introducing dissipation through sideband cooling of a spectator ion, which only acts on the part of the states in Eq. 6.11 that contain an occupied oscillator mode, $|1\rangle_n$. Population is then trapped in that state, while removing the motional excitation, $|1\rangle_n \to |0\rangle_n$. The full engineered resonance dissipation scheme combines the resonant transfer between the even states and their respective dressed states and the dissipation from the dressed state towards the target state:

$$\begin{aligned}
|00\rangle |0\rangle_n \leftrightarrow |\psi_{10}\rangle, & |\psi_{10}\rangle \to |10\rangle |0\rangle_n, \\
|11\rangle |0\rangle_n \leftrightarrow |\psi_{01}\rangle, & |\psi_{01}\rangle \to |01\rangle |0\rangle_n,
\end{aligned} (6.16)$$

as required.

Until now, we have shown that the scheme corrects both $|00\rangle$ and $|11\rangle$ independently, but we must consider how well the process preserves coherence of a possible superposition of those states, such that

$$c_0 |00\rangle + c_1 |11\rangle \rightarrow c_0 |10\rangle + c_1 |01\rangle$$
, (6.17)

where the amplitude and phase of the superposition state is preserved after the correction operation. A prerequisite for maintaining coherence is deduced by considering the following: During the dissipative state transfer, a photon is emitted from the sideband-cooled ion. This photon is potentially measurable, and can therefore be treated as a way to detect if an error has been corrected. However, one should not be able to infer from the temporal information of the detected photon which of the two transfer processes in Eq. 6.16 had occurred. In other words, a condition to maintain qubit coherence is that the two transfer processes must happen simultaneously, making them indistinguishable from each other. This condition is met by setting their effective Rabi frequencies, $\Omega_{\rm eff}^{(f)}$ and $\Omega_{\rm eff}^{(e)}$ in Eq. 6.15, to be equal. As such, the probe powers should be chosen such that $\Omega_f g_f = \Omega_e g_e$, resulting in the coherent error correction operation described by Eq. 6.17.

6.2.2 Autonomous error correction with engineered resonance dissipation

The error correction protocol of the previous section works under the assumption that one of the two physical qubits is immune to errors. This is not a realistic assumption in actual experimental conditions, as all qubits are likely equally susceptible to errors. In this section, the models of the previous sections are extended to an engineered resonance dissipation error correction scheme, based on the theoretical proposal in [233], which is capable of correcting a single bit-flip error occurring on any of the physical qubits.

The logical qubit now consists of three physical qubits, whose logical (error-free) state is given by

$$|\phi\rangle_L = c_0 |0\rangle_L + c_1 |1\rangle_L = c_0 |000\rangle + c_1 |111\rangle.$$
 (6.18)

Based on the schemes presented before (and using the same notation conventions), we could pose to once again apply two weak (Ω_f and Ω_e) beams as probes on each of the three ions, and two strong (g_f and g_e) red-sideband beams to create a state-dependent dressed state splitting. One could then carefully choose the detuning of the probe with coupling Ω_f such that erroneous states $|100\rangle$, $|010\rangle$, and $|001\rangle$ undergo resonant population transfer due to dressed state splitting, to a dressed state that contains $|000\rangle$ $|1\rangle_n$, which can be dissipatively transferred to the logical subspace. Likewise, other states could be corrected with the appropriate detuning of the second probe beam.

While this is a viable approach in the theoretical limit that $g_f/\Omega_f \to \infty$, and likewise for g_e and Ω_e , in practice, the dressed state splitting for various starting states poses resonance conditions on the probe detuning Δ_{Ω} that are not well-resolved. As a result, the beam parameters for the desired transfer, $|100\rangle \to |000\rangle$, would erroneously allow the transfers $|101\rangle \to |001\rangle$ and $|101\rangle \to |100\rangle$ to occur. We thus pursue an improved approach, covered below.

The approach is conceptually similar as before: state transfer from certain states are enabled through a detuned weak probe, through engineered resonance due to dressed state splitting. In contrast to before, additional motional modes are introduced, which helps to avoid spectral overlap of undesired transfer processes. A schematic outline of the beams and level scheme are shown in Figure 6.6(a), where, as an example, the erroneous state $|011\rangle$ is used as an initial level. A beam with coupling strength Ω_f couples the transition $|0\rangle\leftrightarrow|f\rangle$ of all three ions and is detuned by Δ_Ω . This beam is once again referred to as the probe. A pump beam couples to the red-sideband of the same transition, $|f\rangle\leftrightarrow|0\rangle|1\rangle_a$ with coupling strength g_f . The subscript a is used as a label for a chosen motional mode (as more modes are required, which will be introduced later). This beam configuration provides a similar type of conditional state transfer as in previous examples: the number of physical qubits in an initial state $|0\rangle$ dictates the frequency splitting of the dressed state. Figure 6.6(b) shows the excitation under these beam conditions from various starting states, where we have used $g_f = 10\Omega_f$.




Figure 6.6: Autonomous error correction protocol, for a three-ion logical qubit. (a) Level scheme for correcting the erroneous $|011\rangle$ state. See main text for description of coupling strengths Ω_f , g_f and Ω_{c_1} , detuning Δ_{Ω} , and decay rate Γ_{c_1} . (b) Dressed state splitting, showing excitation from various starting states as function of probe detuning Δ_{Ω} , using the values $g_f = 10\Omega_f = 5\Omega_{c_1} = \Gamma_{c_1}$. The intended population transfer from $|011\rangle$ occurs at $\Delta_{\Omega} = g_f/2$. (c) Excitation of various initial states, using $\Delta_{\Omega} = g_f/2$. The dashed line shows population in $|111\rangle$ for initial state $|011\rangle$

The levels that couple from $|011\rangle$ reflect the same structure as the single ion case in Figure 6.2 and have the resonance condition $\Delta_{\Omega}=\pm g_f/2$. Identical excitation is expected from $|101\rangle$ and $|110\rangle$, since both the pump and probe are global beams. Starting states with two zeros have resonance conditions $\Delta_{\Omega}=\pm g_f/\sqrt{2}$, and the starting state $|000\rangle$ has $\Delta_{\Omega}=\pm \sqrt{3}g_f/2$.

Unlike previously, we cannot correct the error by dissipatively removing the additional phonon in mode a, since this would bring us back to the initial erroneous state. Therefore, another weak probe, which we refer to as the transfer probe, can be used to couple the states $|f\rangle$ and $|1\rangle$, on a red sideband corresponding a motional mode c, which is a different one than the one involved in generating the dressed state splitting (a). Assuming the mode c is prepared in the ground state, the weak probe couples the transition $|f\rangle |0\rangle_c \leftrightarrow |1\rangle |1\rangle_c$, the latter of which can be dissipatively cooled to $|1\rangle |0\rangle_c$. Sticking with the example of the initial erroneous state $|011\rangle$, the following transitions occur to correct this error:

$$|011\rangle |0\rangle_{a} |0\rangle_{c} \leftrightarrow \frac{1}{\sqrt{2}} (|f11\rangle |0\rangle_{a} |0\rangle_{c} + |011\rangle |1\rangle_{a} |0\rangle_{c}) \tag{6.19}$$

$$|f11\rangle |0\rangle_a |0\rangle_c \leftrightarrow |111\rangle |0\rangle_a |1\rangle_c \tag{6.20}$$

$$|111\rangle |0\rangle_a |1\rangle_c \rightarrow |111\rangle |0\rangle_a |0\rangle_c. \tag{6.21}$$

However, the attentive reader might notice that from the state $|111\rangle |0\rangle_a |1\rangle_c$ in Eq. 6.19, the transfer probe allows further transitions to $|1f1\rangle |0\rangle_a |1\rangle_c$ and $|11f\rangle |0\rangle_a |1\rangle_c$, which opens up a whole forest of additional possible levels to couple to. To avoid this, three transfer probe beams are used, each addressing one of the three ions, and each on a red sideband of a unique motional mode. Denoting individual ions with the index i, and the motional modes with c_i , these probes couple the states $|f\rangle_i |0\rangle_{c_i} \leftrightarrow |1\rangle_i |1\rangle_{c_i}$, with coupling strength Ω_{c_i} . The transitions of the transfer probe beam, is shown for the example of the initial state $|011\rangle$ in 6.6(a), with coupling strength Ω_{c_1} . Setting $\Omega_{c_i} \ll g_f$ ensures that the dressed state splitting of the pump beam is negligibly affected by this additional beam. Note that each of the three modes c_i must be sympathetically cooled to dissipatively correct errors on any of the three ions. Sympathetically cooling the c_i modes through a second species dissipatively transfers population to the error-free $|111\rangle$ state, with a rate denoted by Γ_{c_1} in Figure 6.6(a). Sideband cooling the c_2 and c_3 modes (with rates c_2 and c_3) produces the same result if the error had occurred on the second or third qubit.

Figure 6.6(c) shows the development of several starting states, using $g_f = \Gamma_{c_1} = 10\Omega_f = 5\Omega_{c_1}$. As intended, only the initial state $|011\rangle$ undergoes population transfer. The dashed line shows the corresponding population in state $|111\rangle$, indicating successful state transfer.

The above set of beams correct erroneous states that contain a single $|0\rangle$ (i.e., the states $|011\rangle$, $|101\rangle$, and $|110\rangle$). The full dissipative error correction scheme is completed by setting up a similar set of correction pulses to enable transfer from erroneous states that contain a single $|1\rangle$. This second set uses a different auxiliary level $|e\rangle$ and mode b, so that this process does not interfere with the first.

All added up, including the sympathetic cooling beams, the full dissipative error correction (DEC) scheme requires 11 beams to be on simultaneously. Most beams have stringent requirements on the precision at which the frequencies and powers are set and need to take into account a large list of AC stark shifts. At least 5 separate motional modes are to be cooled to the ground states, and should not heat during the correction sequence. All considered, the presented scheme for dissipative error correction is an ambitious one, not yet within reach considering constraints of our experimental setup. A sensible approach for experiments is to take a step back and explore simpler proof-of-principle schemes, such as the one described in Section 6.2.1. However, experimental limitations also prohibit us from demonstrating the minimal instance error correction scheme (these limitations are discussed in Section 6.3.3). We thus take yet another step back, and aim to demonstrate the core methods involved in dissipative state transfer through engineered resonance by applying them in a more basic type of non-unitary operation. This operation, classical Boolean logic gates, is discussed in the following section.

6.3 BOOLEAN LOGIC GATES WITH ENGINEERED RESONANCE DISSIPATION

This section details an experimental demonstration of dissipation through engineered resonance as a tool for non-unitary operations on qubits. Other than quantum error correction, one of the applications of this tool is *classical* Boolean gate operations on ionic qubits. With such gate operations, we refer to the well-known logic gates, in which Boolean functions determine the output of one or several binary inputs. Boolean gates are proposed to find their way into quantum computation through hybrid classic-quantum algorithms in quantum machine learning [237, 238], quantum optimization [239], and simulation [223, 224, 240, 241]. In order to have a functionally complete set of gates (i.e., a universal gate set), at least one of the gate operations is required to be irreversible [242]. In previous proposals, the non-unitary action is achieved through auxiliary qubits or by preparation of mixed-state inputs. Such non-unitary action thus requires enlarging the system size, which comes at the cost of needing to account for larger degrees of freedom and therefore higher sensitivity to noise.

Boolean gates are irreversible conditional operations, making the tool-set of dissipation through engineered resonance a suitable candidate. In this section we aim to apply these tools in order to achieve the OR and NOR operations. We choose these operations because they represent two sets of universal logic gates³. These operations therefore serve as a basis for any logical gate operation.

Figure 6.7 shows the truth tables for the desired OR and NOR actions. As before, we use the notation $|ij\rangle$ to denote the electronic levels i and j of ions 1 and 2, given by the tensor $|i\rangle_1 \otimes |j\rangle_2$. The output of the logical gate is mapped onto the first qubit, referred to as the target qubit, and the state of the second qubit always remains unaltered.




Figure 6.7: Truth tables for the NOR and OR operations. The logical action of the gate is mapped onto the first qubit, highlighted blue. Purple arrows indicate cases where a bit-flip is required from input to output.

³ The OR gate on its own is not universal, but the OR and NOT gates together are. The NOT operation is trivial: for the Boolean states represented by $|0\rangle$ and $|1\rangle$, the NOT action is simply a π -pulse between those states.




Figure 6.8: Level diagram of relevant states and transitions of ⁴⁰Ca⁺and ⁸⁸Sr⁺, used for classical gates with dissipation through engineered resonance. The table on the right links coupling strengths and beam nomenclature in the main text.

6.3.1 Experimental overview

The experiments have been carried out on the 'Golden Gate' surface trap (see Section 3.2.2). Two 40 Ca⁺ions and one 88 Sr⁺ion are co-trapped. The 40 Ca⁺ions are used as information carriers, qubits, upon which the gate is performed. The 88 Sr⁺ion enables the dissipative process through sideband cooling, as described in Section 6.1.3.

The ions are stored in the Ca-Sr-Ca configuration. This configuration places the ⁴⁰Ca⁺ions maximally far apart, reducing crosstalk in single-ion addressing. Collisions with background-gas particles occasionally alter this configuration. We therefore periodically apply a sequence of voltages to the trap electrodes that deterministically place the ions in the desired order [63, 243].⁴

An overview of the relevant electronic levels and transitions is shown in Figure 6.8 and is discussed below. The notation for coupling strengths is the same as used in Section 6.1. The logical states $|0\rangle$ and $|1\rangle$ are encoded in the calcium ions' $4S_{1/2}(m=-1/2)$ and $4S_{1/2}(m=1/2)$ levels. Transfer of population between these states is enabled through optical transitions between various $3D_{5/2}$ and $4P_{3/2}$ levels. In particular, $3D_{5/2}(m=1/2)$ is used as an auxiliary level, denoted $|f\rangle$ for engineered resonance. A probe beam couples $4S_{1/2}(m=-1/2) \leftrightarrow 3D_{5/2}(m=1/2)$ of only the first ion, with coupling strength Ω_f , and a pump beam couples the axial common-mode red sideband of $4S_{1/2}(m=1/2) \leftrightarrow 3D_{5/2}(m=1/2)$ of both ions with coupling strength g_f . In principle any of the $3D_{5/2}$ levels that couple to both $4S_{1/2}$ ground states (m = -3/2, -1/2, 1/2, 3/2) could be chosen as an auxiliary level $|f\rangle$ for the pump and probe beams. For our experiment's beam and magnetic field geometry (see Section 3.1.2.2), coupling is maximized for transitions between $4S_{1/2}$ and $3D_{5/2}$ where the change in magnetic quantum numbers is zero, $\Delta m = 0$. The choice to use $3D_{5/2}(m = 1/2)$ as an auxiliary level to mediate the engineered resonance ensures that g_f is maximized under the constraint of finite beam power. As was shown in Section 6.1, a high pump power g_f improves the performance of the operation since dressed state splitting is enhanced and therefore reduces crosstalk between the intended and unintended state transfer mechanisms.

Comparing the OR and NOR operations, as shown in Figure 6.7, we note that the OR operation has one input state that requires the target qubit to be flipped, whereas the NOR operation has three. For the latter, for both $|10\rangle$ and $|11\rangle$ the target qubit is to be flipped independently of the state of the second qubit and therefore does not require engineered resonance methods for state transfer. We can separate the NOR operation into two sub-routines, one that uses engineered resonance to transfer $|00\rangle$ to $|10\rangle$, and another that uses a different method of dissipation to transfer $|1\rangle_1$ to $|0\rangle_1$ (the subscript indicates the ion number). To achieve this, before the engineered resonance step is carried out, population in $|1\rangle_1$ is temporarily transferred and stored in another auxiliary level, $3D_{5/2}(m=-3/2)\equiv |e\rangle$, using an addressed beam that couples these levels with coupling strength Ω_e . At the end of a sequence, $|e\rangle_1$ is dissipatively transferred to $|0\rangle_1$ using 854 nm light, on resonance

with the $3D_{5/2}\leftrightarrow 4P_{3/2}$ transition (denoted by Γ_e in Figure 6.8), from which spontaneous decay brings the population down to the $4S_{1/2}$ ground states⁵. The polarization of the 854 nm beam is set to suppress $\Delta m=+1$ in the $D\to P$ transition, which enhances the likelihood of transferring to $4P_{3/2}(m=-3/2)$, and (m=-1/2). Spontaneous decay from this state is predominantly towards the $4S_{1/2}$ level, where the state $4S_{1/2}(m=-1/2)$, $|0\rangle_1$, is guaranteed. State population in $4P_{3/2}$ has a 0.7% chance of decaying to $3D_{3/2}$, and an approximately 6% chance of decaying to $3D_{5/2}$. For the former, 866 nm light brings population to $4P_{1/2}$, from where it decays to $4S_{1/2}$, though not necessarily in the m=-1/2 sublevel. For the latter, the 854 nm beam remains in effect.

Confining potentials are set so that the axial common-mode (com) frequency is $\omega_{\rm com}/(2\pi)=550\,\rm kHz$. This value is chosen as a trade-off between ensuring a low motional mode heating rates (which increase with lower motional frequency, shown in Section 3.2.6) and a sufficiently high coupling to the motional mode g_f through laser interaction, since coupling scales inversely with the square root of the motional frequency, $g_f \propto \omega_{\rm com}^{-1/2}$. At the chosen frequency, the axial common mode heats at a rate of 106(20) phonons per second.

At the start of an experimental sequence, the ion crystal's axial common and stretch modes are cooled to the motional ground state, using resolved sideband cooling. Both ion species are used for cooling the axial common mode, which is cooled to ≈ 0.14 phonons. Only 40 Ca $^+$ is used for cooling the axial stretch mode, as the centrally placed 88 Sr $^+$ ion does not couple to this mode. Cooling the higher frequency axial mode (the "Egyptian" mode) has not been observed to lead to improved performance of the engineered dissipation sequence, and is omitted to reduce the overall sequence duration. After sideband cooling, the 40 Ca $^+$ and 88 Sr $^+$ ions are prepared in the $S_{1/2}(m=-1/2)$ state, which for the two 40 Ca $^+$ ions corresponds to $|00\rangle$.

The pulse sequences for the OR and NOR gates are shown in Figure 6.10. Preceding the engineered resonance and dissipation pulses, a set of π -pulses applied (globally and addressed) to the transitions $|0\rangle\leftrightarrow|f\rangle$ and $|1\rangle\leftrightarrow|f\rangle$ prepares the ions in one of each possible starting state, $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$:

$$|00\rangle \rightarrow |00\rangle$$
 (6.22)

$$|00\rangle \xrightarrow{R_{1,2}^{0 \leftrightarrow f}(\pi)} |ff\rangle \xrightarrow{R_{1}^{0 \leftrightarrow f}(\pi)} |0f\rangle \xrightarrow{R_{1,2}^{1 \leftrightarrow f}(\pi)} |01\rangle \tag{6.23}$$

$$|00\rangle \xrightarrow{R_1^{0\leftrightarrow f}(\pi)} |f0\rangle \xrightarrow{R_{1,2}^{1\leftrightarrow f}(\pi)} |10\rangle \tag{6.24}$$

$$|00\rangle \xrightarrow{R_{1,2}^{0 \leftrightarrow f}(\pi)} |ff\rangle \xrightarrow{R_{1,2}^{1 \leftrightarrow f}(\pi)} |11\rangle \tag{6.25}$$

where $R_k^{i\leftrightarrow j}(\pi)$ is a π -pulse on the transition of the states $|i\rangle$ and $|j\rangle$, applied to ion(s) k.




Figure 6.9: Excitation spectra for initial states $|00\rangle$ and $|01\rangle$, showing dressed state splitting. The solid lines are simulated data.

We can now calibrate the detuning of the probe beam, by applying a 400 μ s pulse while the pump is switched on, and monitoring the excitation from the initial state for various frequencies. Figure 6.9 shows measured excitation spectra for the initial states $|00\rangle$ and $|01\rangle$. The dressed state splitting is apparent (see for reference the spectra shown in Figures 6.2 and 6.3). Importantly, for the two initial

⁵ Similar to the optical pumping technique for state preparation, as described in Section 2.2.8

states, resonant excitation occurs at distinct detunings, at approximately $\Delta_{\Omega}=4.1\,\mathrm{kHz}\approx g_f/2$ for $|00\rangle$ and $\Delta_{\Omega}=5.8\,\mathrm{kHz}\approx g_f/\sqrt{2}$ for $|01\rangle$. This distinction allows us to transfer population from a chosen initial state, ideally leaving the other unaffected.

In both the OR and NOR gates, the procedure for enabling dissipation through engineered resonance is split into two steps, firstly a transfer pulse using engineered resonance, and secondly a dissipation pulse. These steps are indicated with dashed boxes in Figure 6.10. Though initially conceived as a single-step continuous procedure as discussed in Section 6.1, splitting it into two steps reduces the overall error of the gate, as will be described later in this chapter, in Section 6.3.3.




Figure 6.10: Pulse sequence diagram for the OR and NOR operations. States, transitions, and coupling strengths are as denoted in Figure 6.8.

For the first step in the OR operation, the pump and probe beams, with coupling strengths g_f and Ω_f , and detuning $\Delta_\Omega = g_f/\sqrt{2}$ are applied simultaneously. This results in the interaction described by Eq. 6.5, Rabi cycling between $|01\rangle$ and $|p_{01}\rangle$, with the dressed state

$$|p_{01}\rangle \equiv \frac{1}{2} \left(-\sqrt{2} |11\rangle |1\rangle_n + |f1\rangle |0\rangle_n + |1f\rangle |0\rangle_n \right). \tag{6.26}$$

Since $\Delta_{\Omega} \neq g_f/2$, the state $|00\rangle$ ideally remains unchanged. The pulse is applied for a duration $t = 2\pi/\Omega_f$, resulting in a full population transfer from $|01\rangle$ to $|p_{01}\rangle$.

The pump beam is on resonance with a red sideband, but also couples to the carrier transition in the form of an AC stark shift. This continuously alters the relative phase of the $|1\rangle$ and $|f\rangle$ levels (corresponding the $Z(\theta) = \exp(-i\theta\sigma_z)$ operation, as in Section 2.2.7, Eq. 2.68). This in itself can be corrected for by adjusting the frequency pump beam to compensate this phase drift. However, similar to the decay seen on a carrier Rabi excitation when ions are not perfectly ground-state cooled, the phase coherence when applying a $Z(\theta)$ operation also decays. As will be discussed in more detail in Section 6.3.3, this decay impairs the intended *full* population transfer in the engineered resonance pulse. To mitigate this source of error, an additional beam with similar power as the pump beam is applied near (but not on) resonance with the *blue* sideband, with the aim of canceling out the AC stark shift.

The second step in the OR gate is to transfer $|11\rangle\,|1\rangle_n$, one of the sub-states in $|p_{01}\rangle$, to $|11\rangle\,|0\rangle_n$. This is done by sympathetically cooling the strontium ion. Simultaneously, population that has been depleted from $|11\rangle\,|1\rangle_n$ is repopulated from the remaining sub-states, $|1f\rangle\,|0\rangle_n$ and $|f1\rangle\,|0\rangle_n$, by maintaining coupling with the pump beam.

The NOR operation follows the same two steps as the OR, though with a probe detuning of $\Delta_{\Omega} = g_f/2$ and a pulse duration of $t = \sqrt{2}\pi/\Omega_f$. This enables the transfer of $|00\rangle$ to $|p_{00}\rangle$, and suppresses transfer from $|01\rangle$. Additionally, as previously discussed, the sequence is preceded by the

transfer $|1\rangle_1 \rightarrow |e\rangle_1$, and is succeeded by dissipation to $|0\rangle_1$ through spontaneous decay over the $4P_{3/2}$ level, which achieves the transfers $|10\rangle \rightarrow |00\rangle$ and $|11\rangle \rightarrow |01\rangle$.

At the end of the sequence, state population is read-out with state-dependent fluorescence detection, using a CCD camera, which distinguishes excitation in the S and D levels. Since the logical information, $|0\rangle$ and $|1\rangle$, is carried in the two S levels, a global π -pulse is applied to transfer $|1\rangle$ to $|e\rangle$ (denoted as 'analysis' in Figure 6.10), such that the two logical states can be resolved.

6.3.2 Results

In this section, we display and analyze the results after implementing the sequences described in the previous section in our experimental setup. Figure 6.11 shows the development of the two qubit populations while the state-dependent transfer pulse is being applied for the NOR operation, corresponding to the purple dashed box in Figure 6.10, though noting that for this measurement the preceding transfer pulse has been omitted so that we can focus on solely the functionality of the engineered resonance state transfer. Populations are denoted by P_{ij} , with i and j the states of the first and second ion. The development of the state populations is shown for each of the four initial states. The state read-out does not differentiate between the states $|1\rangle$ and $|f\rangle$, so we show their joint populations as $P_{10} + P_{f0}$ when only the first ion is in $|1\rangle$ or $|f\rangle$, $P_{01} + P_{0f}$ if only the second, and $P_{11} + P_{1f} + P_{f1} + P_{ff}$ if both. Also, these measurements do not ascertain the phonon numbers, which are presumed to remain near the ground state for all initial states except for $|00\rangle$, since a change in motional state must be accompanied by a change in the electronic state.




Figure 6.11: Engineered resonance population transfer in the NOR gate, shown for each possible initial state. The lines indicate simulated results, which include the measured initial phonon number and heating rate as simulation parameters. The dashed line denotes the time that the pulse is applied within the full NOR sequence.

The intended behavior, Rabi cycling from $|00\rangle$ and no transfer from the other initial states, is apparent in Figure 6.11. The solid lines denote simulated data, where we have used $\Omega_f/(2\pi)=1.15\,\mathrm{kHz}$, $g_f/(2\pi)=8\,\mathrm{kHz}$, an initial mode occupation of $\bar{n}=0.14$, and a heating rate of $\dot{n}=106$ phonons per second. These parameters were experimentally obtained: Ω_f and g_f are determined by estimating the Rabi oscillation frequency of their respective bare transitions (see Section 2.2.3). Mode occupation and

heating rate are determined through sideband spectroscopy (see Section 2.3.4). The dark red dotted line marks the duration of the pulse when applied within the full NOR sequence, at $t\approx 600\,\mu s$. At this time, 82(3)% of population is measured to be depleted from the initial state $|00\rangle$, and only 16(3)% is depleted from $|01\rangle$. For comparison, simulations predict 85% excitation from $|00\rangle$, and 11% from $|01\rangle$. The deviation from a full population transfer from $|00\rangle$ is attributed to the non-zero initial phonon number and heating rate, corroborated by the simulated results. Unintended population transfer from $|01\rangle$ is attributed to overlap in the dressed state resonance condition. These error sources are discussed in more detail in Section 6.3.3.

For the initial state $|00\rangle$, after a pulse duration of $t=600\,\mu s$, population has mostly transferred to $|10\rangle$ and $|f0\rangle$, presumed to be distributed evenly between the two, with mode occupation n=1 for the former and n=0 for the latter. This combination of states needs to be dissipatively transferred to $|10\rangle\,|0\rangle_n$. The dressed state $|p_{00}\rangle$ is maintained by continuing to apply the pump beam with g_f , while depleting the additional phonon in $|10\rangle\,|1\rangle_n$ with sympathetic sideband cooling (see the blue dashed box in Figure 6.10).

We experimentally verify the transfer from the dressed state $|p_{00}\rangle$ to the target state $|10\rangle |0\rangle_n$ by monitoring the ions' electronic and motional state population as a function of pulse duration of the probe and sideband cooling beams. Figure 6.12 shows the development of the populations of states $|f0\rangle$ and $|10\rangle$, denoted P_{f0} and P_{10} . The two states are differentiated by running the measurement twice, once with the analysis pulse (see Figure 6.10) which provides us with the populations P_0 and $P_1 + P_f$ for each ion, and once without, which provides us with $P_0 + P_1$ and P_f . Population in the state $|10\rangle$, P_{10} , is inferred from the difference between the first and second measurement. Additionally, the phonon number can be measured by applying a red and blue sideband pulse on one of strontium's $S_{1/2} \leftrightarrow D_{5/2}$ transitions at the end of a sequence. As was described in Section 2.3.4, the occupation in the motional ground state can be inferred from the relative difference in the excitation of the two pulses. In Figure 6.12, the result of this measurement for the population in the motional ground state $P_{n=0}$ is shown alongside the development of the electronic states. We see at time t=0 that P_{f0} and P_{10} are approximately equal, and that the ground state occupation is approximately 50%. Although these measurements do not indicate how the motional occupation is attributed to each of the states |10\rangle and $|f0\rangle$, the measured occupation is consistent with that of the dressed state $|p_{00}\rangle$, where $|10\rangle$ has n=1and $|f0\rangle$ has n=0, and are equally distributed. As the dissipation pulse is applied, population is depleted from $|f0\rangle$. After 1 ms of applying the dissipation pulse, we determine a state population of $P_{10} = 80\%$ and a ground state occupation of $P_{n=0} = 90\%$, from which we conclude that the majority of population resides in the state $|10\rangle |0\rangle_n$, as required of the NOR gate.

The solid lines in Figure 6.12 are simulated results, using the same parameters as in Figure 6.11, including the non-ideal starting values at the end of the engineered resonance step. A dissipation rate of $\Gamma_f = 4.5(6)$ kHz is determined by a least-squares fit between the simulated and measured results. The deviation from a 100% transfer to the intended final state, seen both in simulations and measured results, is attributed to imperfect motional ground state preparation and the motional heating rate.




Figure 6.12: State population and ground state occupation as a function of pulse duration of the dissipation sequence. Markers are measured results, and solid lines are simulated results.

We are now ready to put all the steps together: following the full pulse sequence of Figure 6.10, we measure the output states for all possible input states, for both the OR and NOR sequences. In these sequences, the dissipation step has a 1 ms duration. At longer times, unintended population transfer from initial states that should remain unaltered reduce the overall gate performance.




Figure 6.13: Measured population truth tables of the OR and NOR gates, with the intended output states marked with red dashed lines. Measured populations are derived from 50 repetitions of the experimental sequence. The OR and NOR gates have an average population fidelity of 87(5)% and 81(5)%.

Figure 6.13 shows the truth tables for the full NOR and OR sequences, which display the measured population outcome for each of the four possible initial states, determined from 50 repetitions of the experimental sequence for each input state. Both tables exhibit the intended behavior: for all input states, the majority of the population is transferred to the desired state, marked in the figure with dashed boxes.

To quantify the performance of the gate operation, we calculate a statistical overlap of the gate's output state and the desired state. The output from one specific input state i, has a measure of performance F_i given by $F_i = \sum_j P_{i \to j}^{(\text{meas})} P_{i \to j}^{(\text{des})}$, where $P_{i \to j}^{(\text{meas})}$ and $P_{i \to j}^{(\text{des})}$ are the measured and desired population of output state j. We take the overall statistical overlap of the operation to be the average performance for each of the input states as $\bar{F} = 1/4 \sum_i F_i$. In other words, we take an average of the values in the red boxes in Figure 6.13. With this metric, we determine a gate performance of 87(5)% for the OR gate, and 81(5)% for the NOR gate.

6.3.3 Sources of error

In the following, we describe the main mechanisms that lead to imperfect gate performance. Errors are further quantified in Appendix d.

-Mode occupation-

The primary sources of error, which have previously been alluded to multiple times now, are the non-zero phonon number and heating rate. These manifest as errors in several different ways.

The resonance condition is a function of the pump beam's coupling strength, g_f . The pump beam is a red sideband pulse, which has a coupling strength that is dependent on the motional mode occupation, $g_f \propto \sqrt{n}$ (with n the mode occupation of the lower-energy electronic state, see Eq. 2.64). The resonance condition is not met for any population that is not in the motional ground state, and the intended excitation will not occur for that population. The initial phonon occupation of $\bar{n}=0.14$ suggests that the motional ground state is approximately 86% populated, and that therefore only 86% of population is excited to the dressed state through engineered resonance. This makes up the majority of the error in the transfer of $|01\rangle \rightarrow |11\rangle$ for the OR gate, and the transfer of $|00\rangle \rightarrow |10\rangle$ for the NOR gate.

Similarly, heating *during* the engineered resonance pulse reduces the success of the population transfer. The heating rate of $\dot{n} = 106$ phonons per second corresponds to an increase of 0.06 phonons

during the 600 µs pulse. Simulations suggest that this increase in phonons contributes 3% error in the $|00\rangle\leftrightarrow|p_{00}\rangle$ transfer of the NOR gate, and 2% in the $|01\rangle\leftrightarrow|p_{01}\rangle$ transfer of the OR gate.

Mode heating also limits the final achievable phonon number during the dissipation step. We can approximate the steady-state phonon number from the differential equation $\dot{P}_{n=0}(t)=-\dot{n}P_{n=0}(t)+P_{n=1}(t)\Gamma_f$. If the final mean phonon number is small, $\bar{n}\ll 1$, which is the case when the cooling rate is much higher than the heating rate, $\Gamma_f\gg\dot{n}$, the phonon occupation for $n\geq 2$ is negligible, such that $P_{n=0}+P_{n=1}\approx 1$. Solving for the steady-state solution, $\dot{P}_{n=0}(t)=0$, gives a final phonon number of $\dot{n}/(\dot{n}+\Gamma_f)=0.02$. ⁶ A state with mode occupation above zero will continue to be transferred to other states by the pump pulse and therefore does not settle to the correct final state.

The non-zero mode occupation during the sequence also disrupts the success of π -pulses, such as those used in state initialization, the transfer pulse to $|e\rangle$, and the analysis pulse at the end of the sequence (see Figure 2.10 in Section 2.3.4). Errors due to imperfect π -pulses are estimated to be on the order of 1-2%.

-Overlap in resonance condition-

As discussed in Section 6.1.2, the resonance condition for population transfer, $\Delta_{\Omega} = \pm g_f/2$ for the NOR gate, and $\Delta_{\Omega} = \pm g_f/\sqrt{2}$ for the OR gate, are in reality conditions with a finite bandwidth, as can be seen in the peaks in the excitation plots of Figures 6.2(b) and 6.3(b). The dressed state excitation spectrum is a modification of the Rabi excitation spectrum, described in Section 2.2.3, generically given by the excitation population $P_{\rm ex}$:

$$P_{\rm ex} = \frac{\Omega^2}{\Omega^2 + \Delta^2} \sin^2 \left(\frac{\sqrt{\Omega^2 + \Delta^2} t}{2} \right),\tag{6.27}$$

for Rabi frequency Ω and detuning Δ . Excitation from the initial state $|00\rangle$ is on resonance at $\Delta_{\Omega} = \pm g_f/2$ and has an effective Rabi frequency of $\Omega_f/\sqrt{2}$. Substitution into Eq. 6.27 gives the excitation spectrum

$$1 - P_{00} = \frac{\Omega_f^2}{\Omega_f^2 + 2(\Delta_\Omega \pm \frac{g_f}{2})^2} \sin^2 \left(\sqrt{\frac{\Omega_f^2}{2} + \left(\Delta_\Omega \pm \frac{g_f}{2} \right)^2} \frac{t}{2} \right). \tag{6.28}$$

Similarly, excitation from the initial state $|01\rangle$ is on resonance at $\Delta_{\Omega} = \pm g_f/\sqrt{2}$ and has an effective Rabi frequency of $\Omega_f/2$ and has an excitation spectrum given by

$$1 - P_{01} = \frac{\Omega_f^2}{\Omega_f^2 + 4(\Delta_{\Omega} \pm \frac{g_f}{\sqrt{2}})^2} \sin^2 \left(\sqrt{\frac{\Omega_f^2}{4} + \left(\Delta_{\Omega} \pm \frac{g_f}{\sqrt{2}} \right)^2} \frac{t}{2} \right). \tag{6.29}$$

There also exists resonant excitation from the initial state $|01\rangle$ when $\Delta_{\Omega}=0$, corresponding to the central peak in the dressed state splitting shown in Figure 6.3(b). Excitation from $|01\rangle$ at this transition is given, as function of detuning Δ_{Ω} , by:

$$1 - P_{01} = \frac{\Omega_f^2}{\Omega^2 + 2\Delta_\Omega^2} \sin^2\left(\sqrt{\frac{\Omega^2}{2} + \Delta_\Omega^2} \frac{t}{2}\right). \tag{6.30}$$

If a pulse is applied to the state $|00\rangle$ with a detuning $\Delta_{\Omega}=g/2$ and a gate duration of $t=\sqrt{2}\pi/\Omega_f$, as is intended for the NOR gate, this same pulse produces an excitation of about 5% if the initial state were $|01\rangle$ (about 4% from off-resonant excitation on the central peak, and the rest from the other dressed state). For the OR gate, pulse settings engineered to excite the $|01\rangle$ state result in off-resonant excitation of 17% from $|00\rangle$ as an initial state, which is a substantial error. This error is worse than the

⁶ In fact, the sideband cooling step before the gate sequence should be able to reach a similar, if not better, minimum phonon number. We have not figured out why our measured initial phonon number, $\bar{n} = 0.14$, is higher.

NOR gate, because a longer pulse duration is required, and the effective coupling of the undesired excitation is higher.

An obvious experimental improvement, in theory at least, to reduce this off-resonant crosstalk is to improve the ratio between the pump and probe coupling strengths, g_f and Ω_f , which narrows the bandwidth of the resonance conditions. However, as previously mentioned, g_f is limited by experimental hardware (i.e. beam power) and a low probe power Ω_f implies a high sequence duration, where other sources of error such as motional heating and decoherence grow.

Another method to reduce error induced by off-resonant excitation is with a clever choice of beam parameters: For example, in the OR gate, a detuning of $\Delta_{\Omega}=g_f/\sqrt{2}$ and a pulse duration of $t=2\pi/\Omega_f$ are used, which enables population transfer from $|01\rangle$. Unintended excitation from $|00\rangle$ can be calculated using Eq. 6.28. This excitation vanishes when the argument in the sine term is an integer multiple of π :

$$\sqrt{\frac{\Omega_f^2}{4} + \left(\Delta_{\Omega} \pm \frac{g_f}{\sqrt{2}}\right)^2} \frac{t}{2} = n\pi \tag{6.31}$$

with n a positive integer. Substituting the values for t and Δ_{Ω} , we find that off-resonant excitation is avoided by setting the ratio to be $g_f/\Omega_f = (2+\sqrt{2})\sqrt{2n^2-1}$. A ratio of ~ 9 allows the undesired excitation to undergo a full 4π rotation, returning the excited population to zero.

Similarly, in the NOR gate, undesired excitation occurs from the initial state $|01\rangle$, described by both Eqs. 6.29 and 6.30. Choosing the pump/probe ratio to cancel out undesired excitation results in two unique conditions: $g_f/\Omega_f=(1+\sqrt{2})\sqrt{8n^2-1}$ and $\sqrt{8m^2-2}$, which cannot be met simultaneously. A reasonable choice is $g_f/\Omega_f\approx 5.8$, which results in a 1.5% error.

-(Off-)resonant carrier excitation-

The pump beam is a sideband pulse with power g_f . With respect to the carrier transition, this beam has power g_f/η , with in our case a Lamb-Dicke parameter of $\eta=0.056$, and is detuned from the carrier transition by $\Delta_{\rm rsb}=550\,{\rm kHz}$. Off-resonant Rabi excitation is described in Section 2.2.3. Using the effective Rabi frequency g_f/η and detuning $\Delta_{\rm rsb}$, gives the excitation probability

$$\frac{g^2}{g^2 + \eta^2 \Delta_{\text{rsb}}^2} \sin^2\left(\sqrt{\frac{g^2}{\eta^2} + \Delta_{\text{rsb}}^2} \frac{t}{2}\right) \tag{6.32}$$

which for our parameters is a high-frequency oscillation on the carrier with a mean excitation of 3%. While in theory the mean excitation remains at this value, in practice dephasing of the coupled states (see next subsection, AC Stark shift) can lead to a steady increase of this off-resonant carrier excitation.

Another error of a more technical nature is resonant carrier excitation: while the laser used to drive the sideband transition has a sufficiently narrow linewidth (about $10\,\mathrm{Hz}$), the spectral profile has a non-negligible noise-floor extending several MHz from the peak center. We have determined the relative coupling strength of the noise floor to be about 10^{-4} . Resonant coupling with the carrier transition is given by $10^{-4}g_f/\eta=2\pi\cdot14\,\mathrm{Hz}$, which corresponds to an excitation of about 1% after 2 ms.

-AC Stark shift-

In addition to carrier excitation, the pump beam induces an AC Stark shift that shifts the frequency of $|f\rangle$ by $g_f^2/(4\eta^2\Delta_{\rm rbs})\approx 9\,{\rm kHz}$. The pump coupling strength g_f is dependent on both beam intensity and the ions' motional mode occupation. Beam intensity fluctuations and a distribution of phonon numbers broaden this frequency shift. This broadening leads to a dephasing of the $|f\rangle$ and $|1\rangle$ states, which affects the population transfer of the engineered resonance pulse. If left untreated, in our experimental conditions the engineered resonance population transfer would be about 90% successful, where the source of error is predominantly intensity noise. We attempt to cancel out the AC stark shift by introducing an additional beam near resonance with the blue sideband of the $|1\rangle\leftrightarrow|f\rangle$ transition, thus reducing the dephasing error. We have not not precisely characterized the errors due to AC stark shifts with or without the correction beam but have qualitatively noticed an improvement. The

addition of the correction beam does, however, introduce another source of errors due to (off-)resonant carrier excitation on the order of 3 - 4%, as discussed in the previous paragraph.

-Spontaneous decay-

In addition to dissipation through sideband cooling, a second channel of dissipation is required for the NOR gate, which transfers the state of the first ion from $|1\rangle_1$ to $|0\rangle_1$. This transfer involves spontaneous decay in Calcium's $4P_{3/2} \to 4S_{1/2}$ dipole transition. Ideally, spontaneous decay only occurs to $4S_{1/2}(m=-1/2)$, which can be ensured to be the case if only $4P_{3/2}(m=-3/2)$ is populated. The polarization of the repumper that populates the $4P_{3/2}$ levels (at 854 nm) is set to maximize the occupation of this particular state. We have measured a 97% success rate in the transfer from $|1\rangle_1$ to $|0\rangle_1$.

This success rate can likely be improved but requires changes to our optical setup: Currently, the 854 nm shares a beam path with multiple other beams, making it difficult to independently set its polarization.

-State preparation and measurement-

Optical pumping (see Section 2.2.8) places ions in the $|00\rangle$ state, with a > 99% success rate. Producing the other initial states, $|01\rangle$, $|10\rangle$, and $|11\rangle$, makes use of a set of global and addressed π -pulses, described by Eqs. 6.22 - 6.25. The transfer success of these pulses is limited by the non-zero mode occupation in various motional modes, which leads to a decay in Rabi oscillations, as described in Section 2.3.4. Even though the axial modes are cooled near the ground state, global operations suffer from some decay, likely due to coupling with uncooled radial motional modes in anharmonic trap potentials [66]. Similarly, addressed (radial) Rabi-oscillations decay due to the uncooled radial modes. This results in state preparation errors of about 2% for all initial states other than $|00\rangle$. State readout also requires the application of a global π -pulse, and thus also introduces a similar error.

Radial modes are not sideband-cooled in order to avoid lengthy sequences that heighten the risk of ion loss (the mechanism of ion loss in periods where Doppler cooling beams are off is discussed in Chapter 4). Similar experiments in deeper traps, where ion loss is less likely, would benefit from radial mode cooling. Additionally, improved cooling techniques such as electromagnetically induced transparancy (EIT) cooling or polarization gradient cooling (PGC) would help cool multiple modes rapidly.

6.3.4 Conclusions

In this section we have used the tools of engineered resonance and dissipation to produce classic gates with ionic qubits. This opens up the way to exciting new possibilities in hybrid quantum-classical algorithms, such as quantum machine learning, quantum optimization, and quantum simulation. In many of such algorithms, nonunitary operations are used in conjunction with unitary quantum gates. The tools presented in this section demonstrate a promising ability to implement conditional dissipative routines within quantum algorithms. The gates OR (together with the trivial NOT) and NOR, have been chosen not just because of their experimental simplicity, but because of their logical universality. These gate elements are thus widely applicable for many kinds of non-unitary qubit transformations.

However, it must be noted that little has been demonstrated about the quantum nature of the procedure. For example, could we pose that applying the NOR gate to a not-entangled superposition $1/\sqrt{2}(|00\rangle+|01\rangle)$ would result in the formation of an entangled Bell state, as

$$NOR\left[\frac{1}{\sqrt{2}}(|00\rangle|0\rangle_n + |01\rangle|0\rangle_n)\right] \to \frac{1}{\sqrt{2}}(|10\rangle|0\rangle_n + |01\rangle|0\rangle_n)? \tag{6.33}$$

Although this was not the purpose of the NOR gate, it is interesting to wonder to what extent a classical boolean operation can mix with a quantum system to produce quantum gates. Unfortunately, the NOR gate does *not* produce the operation described by Eq. 6.33. While applying the NOR gate to the superposition state would result in the desired populations, the state is mixed as a result of the dissipation process. In order to maintain qubit coherence between two states, one of the conditions is that *both* states are indistinguishably subjected to the same dissipation process.

An obvious application where such an indistinguishable dissipation process is desirable is quantum error correction, of which the theory was described in Section 6.2. While some experimental efforts have been made in our setup to realize these concepts, the sources of error presented in the previous section are ironically more dominant in the presented error correction schemes. If we take the logical qubit to be $|\phi\rangle_L = c_0 |00\rangle + c_1 |11\rangle$, then a correction of the faulty state $|\phi\rangle_E$ could be achieved by a simultaneous application of an OR and an AND gate. Assuming that both operations have comparable error rates, the gate would endure a 75% fidelity under our experimental conditions. However, we have noted in Section 6.3.3 that phase between qubit states decays while the pump beam is applied. In the 2 ms gate time, the phase coherence is fully decayed. The minimal instance error correction scheme (Section 6.2.1) would suffer even more from such coherence decay mechanisms, as more beams are required. The limitations, however, are only technical in nature: dephasing and imperfect population transfer are both primarily limited by the crystal's initial phonon numbers and heating rates. The ongoing search for improved trap designs and electronics aims to minimize heating rates for high-fidelity quantum computation. Once such improvements are in place, a logical follow-up to the measurements presented in this section is to attempt similar gates while maintaining qubit coherence. This paves the way to applying such processes to fully-fledged dissipative quantum error correction.

7

SUMMARY AND OUTLOOK

It's now safe to turn off your computer

— Windows 95

This thesis presented several advances in the field of quantum computation. The employed setup has been designed to demonstrate scalability, using a cryogenic segmented surface ion trap, with capabilities of trapping and manipulating multiple ion species.

7.1 RECAP

Chapter 3 provides an overview of the experimental apparatus, which has undergone multiple developmental changes throughout the years. The setup has hosted a large number of unique surface traps, two of which have been instrumental to producing the results presented in Chapters 4 - 6, and those presented in other works [106, 111, 112, 125].

Chapter 4 discusses RF heating, a phenomenon that has been known to affect trapped ions in RF potentials since the early days of Paul traps but has until now hardly been considered in terms of its influence on small numbers of trapped ions, such as in the context of quantum computation. RF heating affects ions when they melt, and reduces the rate at which they can be recrystallized. Efficient recrystallization, and therefore a good understanding of the RF heating process, is increasingly important with decreasing trap size and with increasing numbers of trap sites, as will be the case for scalable quantum computation.

Chapter 5 introduces the fundamentals of physical manipulation of trapped ions through tailored trap-electrode voltage sequences. The ability to transport, split, and rotate ion chains are fundamental concepts of the quantum processor array used in the QCCD trap architecture, and are therefore paramount in scalable quantum computation. We provide generalized guidelines on how to generate a set of voltage sequences to realize transport and splitting operations. Furthermore, an in-depth study of ion crystal rotations is presented, showing experimental results of the calibration and optimization procedures, and finally experimental results of optimized rotation sequences. We show that ion crystal rotations are achievable while maintaining coherence and low motional excitation.

Chapter 6 introduces a novel toolset that combines classical operations with quantum systems, consisting of engineered resonance conditioned on many-body observables of a system and dissipation techniques. Applications of such hybrid quantum-classical operations are envisioned, for example, in simulations of quantum systems that interact with a classic environment. One notable application, which is discussed in Chapter 6, is to use the conditional state transfer approach for quantum error correction. This presents a new paradigm of error correction. Instead of outsourcing the decision making process of whether to correct an error or not to a classical computer, the conditional decision is integrated into the qubit system. As a proof-of-principle of the mechanism of dissipation through engineered resonance, the chapter shows how these tools can be used to perform classical gates with qubits. Technical issues such as motional heating rates and coherence times limit the success of these processes and furthermore prevent our setup from using the presented tools for QEC. However, with the continuing development of improved traps, many of these obstacles will be overcome.

7.2 NEXT STEPS FOR THE EXPERIMENTAL APPARATUS

The experimental apparatus has been designed to demonstrate the prospects of scalability of quantum computation. One may note, in that regard, that the focus of this thesis is geared more towards examining specific tools (ion manipulation, recrystallization, and classic control) and not as much toward actually applying quantum circuits in our system. My vision for our experimental setup is

to go beyond examining individual tools and implement full-fledged quantum algorithms (as was initially proposed when the project was conceptualized). The setup has several unique features that are still rare in the ion trapping community, making it ideal for exploring new territory in trapped-ion quantum computation: we can co-trap two species, have coherent control of both of them, and can transport them in-sequence. With that prospect in mind: Below, we suggest improvements to the cryogenic apparatus and experimental setup.

-High-NA lens-

The high-NA lens described in Section 3.1.1.2 was placed in the upgraded design of the apparatus as a way to show that readout error caused by spontaneous decay of the qubit during state readout can be lowered by reducing the total required detection time. Additionally, quantum circuits that require mid-sequence state readout [244–246], for example those that include quantum error correction [50, 231, 235], benefit from quick state measurements to reduce the overall sequence duration and limit measurement-induced heating of the ion string.

With the high-NA lens in our setup, we are able to distinguish between qubit states in $50\,\mu s$ with a certainty of $\approx 99.999\%$. During this time-window, there is a less-than 0.001% chance of decay from the excited state. Notably, it seems that one can safely place a lens with such high NA uncomfortably close to the trap surface without adverse effects¹ and without notably compromising the fluorescence collection efficiency.

That being said and done, while a great demonstration of state readout, the inclusion of the lens has complicated the design and usage of the setup. The inner heat shield, which was cramped enough as it was, underwent a major overhaul to accommodate the nanopositioners and wiring. The lens itself had to have slots cut out to ensure optical paths at 45° with respect to the trap surface are not obstructed. Even with those cuts, the beams coming in at 45° are now still limited in how tightly they can be focused down without scattering from the lens. Also, the fact that the mesh grid needed to be included was a consequence of having an in-cryo high-NA lens.

In future versions of the apparatus, this lens should be replaced, where the focus² is shifted from having immensely high readout fidelity to having a lens (or set of lenses) that provides good multi-ion and multi-species addressing and readout. A suitable optical solution should be chosen such that the in-vacuum nanopositioner will not be a requirement. The lens should be custom-made for low-aberration ion addressing and well-resolved state detection. The design of the lens goes hand-in-hand with the design of the ion addressing unit, discussed below.

-Ion addressing unit-

Our setup has two beams per ion species for coherent qubit manipulation³, one along the axial direction that enables collective operations on all ions simultaneously, and a tightly focused beam that addresses individual ions. In our current setup, the beam path of both of these beams is fixed. Single ion addressing can therefore currently only be applied to one ion at a time, and in a fixed place along the trap. The latter restriction can be overcome by in-sequence shuttling of the ion chain, which physically places the ion on which the qubit operation is to be performed in-line with the beam. However, it is desirable to be able to address any arbitrary combination of ions from the chain.

The experimental apparatus therefore will require an ion addressing unit that is capable of producing multiple focused beams. There exist many strategies to achieve this [247, 248]. To date, the experimental hardware that has proven to have the best combination of optical control and precision is a crossed acousto-optic deflector (AOD) setup [248]. Similar designs are planned for our experimental setup. The design of the addressing unit must be coordinated alongside the design of the in-vacuum high-NA lens, discussed previously.

-Axial micromotion-

It has been shown in Section 3.2.5 that we can minimize micromotion in the radial direction (along the direction of the addressing beam) to a micromotion index of $\beta_r < 10^{-3}$. The *axial* direction is measured to have a micromotion index of $\beta = 0.25$. This is counter-intuitive, since the axial direction in theory should not have any RF field, and therefore no micromotion component. We suspect that the

¹ Although this statement is speculative; we do not have data on the trap's performance without the lens in place for comparison.

² No pun intended

³ Not counting the Raman beams used in [112], which have not been used in the works described in this thesis

source of axial RF field is a short unshielded wire, off the side of the trap carrier, that transports the RF signal from the resonator to the trap PCB. Indeed, basic electric field simulations indicate that the voltage that this wire carries is more than enough to produce the RF field amplitude of $0.76\,\mathrm{V\,mm^{-1}}$ measured at the position of the ion. A portion of this RF field is shielded by the clamp that holds the trap carrier itself and by the surface of the trap, but apparently enough is transmitted to result in a notable amount of axial micromotion.

For low numbers of same-species ions, axial micromotion does not significantly affect ionic qubits. However, in larger ion crystals, where trap anharmonicity becomes increasingly noticeable, micromotion can play a dominant role in ion heating [249]. This is especially troublesome for mixed-species ion chains, where individual ions experience unequal amplitudes due to the RF field, thus providing a strong channel of mode mixing. Axial micromotion is a technical issue that should not be present in our system, and can be avoided by redesigning the trap PCB, shielding the RF wiring, or both.

-Control software and hardware upgrade-

Our experimental system is hanging on to an outdated under-supported piece of control software and hardware. Control and communication with external hardware such as the CCD camera and the electrode voltage supply are clumsy and limited in functionality.

We suggest a major overhaul of the experimental control system, by replacing the current version with hardware and software that is designed for use in quantum physics experiments. In particular, the Sinara hardware, together with ARTIQ software, make up a well-maintained control system geared towards quantum physics experiments, and has been implemented successfully in several experimental setups in our institute. One of the main improvements over our current software, and an absolute requirement for scalable quantum computation, is the ability to run sequences with conditional branching.

-Improved ground-state cooling-

Many of the sources of error discussed throughout this thesis stem from mode heating and motional decoherence. At least in part, these issues originate from mode mixing [66], in which cooled modes are able to exchange phonons with uncooled ones. Such mixing is especially prominent in anharmonic trap potentials, a typical characteristic of surface traps. It is therefore desirable to implement improved cooling techniques that can cool multiple modes simultaneously.

Techniques such as polarization gradient cooling (PGC) [85] and electromagnetically induced transparancy (EIT) cooling [84] have both been demonstrated in our group to be reliable methods to efficiently cool numerous motional modes simultaneously. Either of these techniques would require restructuring of the beam paths on our optical tables, but would be well worth the effort.

-Improved trap design-

One of the lessons learned from having had various types of traps in our apparatus is that currently simple trap designs outperform the complicated ones. However, the road towards scalability in trapped-ion quantum computation will inevitably require more complicated trapping architectures, including junctions and electrical connectivity to a large number of electrodes.

In order to determine how to proceed in quantum computation with ion traps, one must assess which challenges in scalability are to be tackled in the near future in this particular setup. Experience tells that tackling too many challenges simultaneously is counter-productive. That is to say, if the focus of the project is to investigate ways to advance ion trap architectures, one need not work on having a system with optimal ion addressing, detection, state-of-the-art shuttling capabilities, and exceptional vibration and magnetic field isolation.

My vision for the experimental setup, therefore, is to make an incremental upgrade to the ion trap, while maintaining proven designs and fabrication techniques. The trap layout will be modified to increase the trap depth, to reduce the risk of ion loss, which is especially prominent in mixed-species operation. Additionally, the electrode structure will be tailored towards generating well-isolated trap sites, with dedicated regions for physical ion manipulation such as chain splitting and rotations. Electrodes can be further tailored to minimize potential anharmonicity around trapping sites. The optical access that the current trap provides should be maintained.

7.3 CONCLUSION

As the introduction of this thesis started with a cliché, it only makes sense to end with one: the works presented in this thesis are stepping stones towards the development of scalable digital quantum computers with trapped ions. We have investigated and presented improvements in experimental techniques that make up fundamental components of quantum computers. Specifically, our project demonstrates the advantages, but also the obstacles, in control of mixed-species ions in a segmented surface trap in a cryogenic environment. While the road towards reliable quantum computers that outperform classical ones is a long one, and is full of challenges, I hope to think that the work in this thesis reassure that we are still on the right path.



In section 4.6.1.2, we provide an analytic approximation for the average collision rate for two melted ions with three dimensional harmonic motion. This average rate is a parameter for a distribution function from which a random collision time is generated. In this section, we provide a 'derivation' for the average collision rate¹. We then verify that the collision time model as used in the simplified ion cloud energy simulation accurately portrays the collision dynamics. The contents of this section are adapted from [139].

A.1 COLLISION RATE DERIVATION

In Section 4.6.1.2, the motion of two ions is described as sinusoidal, with different amplitudes $a_{i,k}$ and frequencies ω_k in all three dimensions. In each dimension separately, ions are within collision range (such that the distance between ions d_k is below a given range r_c) twice per oscillation period, $2\pi/\omega_k$. This collision condition in one dimension is thus represented by a pulse wave $B_k(t)$, with period $T_k = \pi/\omega_k$ and pulse duration Δt_k , as in Eq. 4.49. We define a two-ion collision as an event where the collision condition is satisfied in all three dimensions simultaneously, given by the pulse wave $B_{3D}(t) = B_x B_y B_z$. As this three-dimensional pulse wave is aperiodic (the periods T_k are not rationally related), $B_{3D}(t)$ does not have a fixed pulse period². However, an average pulse period can still be defined, given by the average time between pulses. The collision rate \bar{f}_{coll} is then the average number of pulses in B_{3D} per time. In this section, we derive the average collision rate in B_{3D} as function of pulse parameters T_k and Δt_k .

The probability that the 3D collision condition, $B_{3D}(t) = 1$, is met at any moment in time t is given by the product of the probabilities that $B_k = 1$ for k = x, y, z, $P_{3D} = \prod_k P_k$, with the 1D probabilities $P_k = \Delta t_k / T_k$. Intuitively, the collision rate is given by the product of the momentary collision probability P_{3D} , and the effective rate at which P_{3D} is resampled.

We derive f_{coll} with a geometric argument, depicted in Figure a.1 (shown in two dimensions, for clarity). The pulse wave $B_{2D}(t) = B_x(t)B_y(t)$ in Figure a.1(a) is a function of time t. The individual pulse waves B_x and B_y can, however, graphically be separated into two time dimensions, t_x and t_y , depicted in Figure a.1(b) as two time axes. The vertical and horizontal shaded regions correspond to regions where $B_x = 1$ and $B_y = 1$, respectively. The locations where the vertical and horizontal bars meet are places that satisfy the collision condition, $B_{2D} = 1$. These are graphically represented by the yellow boxes, with sidelengths Δt_x and Δt_y . 'Real' time t parametrically follows the diagonal line, $t_x = t_y$. Whenever the real time line crosses a yellow box, a collision occurs (starred regions).

We consider the parallelogram unit cell U (red dashed line), whose height is given by the lowest value of $\{T_i\}$ (in this example, the lowest value is T_y . We see later that this choice is made without loss of generality), and a base width given by the remaining value (T_x) . The parallelogram angle follows the real time line t, at 45° . Each unit cell contains exactly one collision box. The collision box is projected at 45° through the unit cell, denoted by the green shaded area. Graphically, a collision occurs if the time line passes through this area, as the real time line then has overlap with the collision box. In the example unit cell of Figure a.1(b), no collision occurs. Since T_x and T_y can be assumed to be irrationally related, the location where the time line enters a unit cell is uniformly distributed. Therefore, the probability of a collision occurring in a unit cell is given by the ratio of the green shaded area to the area of the unit cell. This is identical to the ratio of the lengths of the base of the green area to the base the unit cell. The base of the green area is the projection of the collision condition box along the axis of the time line onto the base of the unit cell box (indicated by the thick green line). Since the projection is along 45° onto the x-axis, the length of the projection is given by $\Delta t_x + \tan (45^{\circ}) \Delta t_y = \Delta t_x + \Delta t_y$. The

¹ Where 'derivation' is in quotes because the lack of mathematical rigorousness.

² Nor do the pulses have a fixed duration




Figure a.1: Graphical representation of collision condition, for deriving the 3D collision frequency. (a) In dimensions x and y, the collision conditions are represented by pulse waves B_x and B_y . In 2D, a collision is represented by the non-periodic pulse wave $B_{2D} = B_x B_y$. (b) $B_x(t)$ and $B_y(t)$ are displayed as two orthogonal temporal dimensions, such that their respective collision conditions are vertical and horizontal bars. "Real" time is the diagonal line $t_x = t_y$. Collisions occur where real time crosses vertical and horizontal bars simultaneously, denoted by the yellow boxes. This is equivalent to the real time line crossing the green shaded area, the 45° projection of collision boxes in unit cells U. Therefore, in the example unit cell in (b), no collision occurs. The probability of a collision occurring in a unit cell is thus given by the ratio of the projection of the green region onto the base of the unit cell, and the base of the unit cell itself. This probability, schematically shown in (c) 2D and (d) 3D, is multiplied by the frequency that the time line enters new unit cells to give the collision frequency.

probability that a collision occurs within a unit cell is therefore $P_{U,2D} = (\Delta t_x + \Delta t_y)/T_x$ (see Figure a.1(c)).

Extending this concept into three dimensions (see Figure a.1(d)), the base of the unit cell is now two dimensional, with an area of $A_U = T_x T_y$, assuming the shortest time in $\{T_i\}$ is T_z . The area of the projection of the collision box is given by $A_{\text{coll}} = \sum_k \sum_{l>k} \Delta t_k \Delta t_l$. The probability of a collision occurring within a unit cell is $P_U = A_{\text{coll}}/A_U$. Note that with this geometric argument P_U can exceed 1, and should be numerically capped off at this value. For typical experimental values of Δt_k and T_k it is generally the case that $P_U \ll 1$.

After passing through N unit cells, on average $n_{\text{coll}} = P_U N$ collisions have occurred. The time line enters a new unit cell at intervals T_z , so $N = t/T_z$. The collision rate is thus $\bar{f}_{\text{coll}} = n_{\text{coll}}/t = P_U/T_z$. Rewriting gives

$$\bar{f}_{\text{coll}} = \prod_{i} \left(\frac{\Delta t_i}{T_i} \right) \sum_{i} \frac{1}{\Delta t_i}, \tag{a.1}$$

which conforms with the intuition that the collision rate is given by the product of P_{3D} and an effective resample rate.

A.2 COLLISION RATE VERIFICATION

In the simplified RF heating model presented in Section 4.6.2, we update the ions' secular energy at intervals given by an estimated time t_{coll} between ion-ion collisions. To generate this time for a

given set of ion and trap parameters, we first find an *average* collision time \bar{t}_{coll} , as derived in the previous section. We then randomly sample from a truncated exponential distribution characterized by this average rate. In this section, we analyze these methods through comparison with numerical simulations

The average collision time $\bar{t}_{\rm coll}$ is a function of the ions' secular motion amplitudes $a_{i,k}$ (for ion i in direction k), and their relative motional phase ϕ_k . We generate a random set of these parameters $(a_{i,k} \in (0,10^{-4}) \text{ m}, \phi_k \in (0,2\pi))$ and simulate the ions' time-dependent positions under the influence of the RF potential with the full ion dynamics simulation, though neglecting Coulomb interaction to ensure that the parameter set $a_{i,k}$ and ϕ_k remains unaltered for the duration of the simulation. We then query how often the collision condition is met, i.e. $|r_{1,k}-r_{2,k}| < r_c$, simultaneously for all $k=\{x,y,z\}$. The number of collisions occurring in the duration of the simulation gives us the simulated collision rate. We compare this collision rate with the analytically obtained collision rate given by Eq. 4.50. We repeat this method 20 times with randomly generated parameters, and plot the results in Figure a.2, which compares the simulated and analytically obtained average collision rate (triangles). We additionally repeat this process another 500 times, though replace the ion dynamics simulation with an analytic expression for ion position, as given by Eq. 4.66, which greatly reduces the computation time per repetition (circles). Results in Figure a.2 show good agreement between simulations and the analytic model. The larger deviation at large collision times is attributed to larger statistical errors, as only a few collisions occur during the millisecond simulation time.

After each collision event, a new collision time is to be determined. For this, one could simply use the average time \bar{t}_{coll} , since this suitably approximates the time over which RF heating occurs after thousands of collisions. However, we attempt to design our time selection procedure to be more physically accurate, without compromising the computation time of the simplified RF heating simulation. We do this by drawing a new time randomly from a distribution function, which is parametrized by the analytically obtained average collision time.

Designing a physically accurate description of the distribution function is mathematically cumbersome, given the aperiodic nature of the 3D collision criterion. We thus aim to find a simplified approximation of the distribution function. In our work, we have chosen an exponential distribution based on the notion that the aperiodic collision conditions are uncorrelated events. This is a simplification, since after a collision occurs, a subsequent collision is more likely to occur close to a multiple of the three motional periods. Using the same randomized parameters as the results presented in Fig. a.2(a), we obtain histogram bin counts of times between collisions, which are normalized to the numerically obtained *average* collision time. The resulting data is shown in Figure a.2(b). For comparison, an exponential distribution function is plotted. There is a good agreement at longer (> \bar{t}_{coll}) collision times, but qualitatively less overlap at lower times. Regardless, this method represents a computationally inexpensive improvement over simply selecting the average time \bar{t}_{coll} as a new collision time. A further improvement is made by assuming that the minimum time between collisions is required to be at least one half oscillation period, T_k . This restriction is imposed by rejecting and resampling the randomly generated collision times, effectively truncating the exponential distribution function. This truncated distribution is also shown in Figure a.2(b).

Ultimately, the main performance check is comparing the simplified simulation to the full ion dynamics simulation, as in Figure 4.8(a). We can also compare the performance of the simplified simulation using various strategies for generating new collision times, using 1) the analytically obtained average time directly without sampling from a distribution, 2) sampling from an exponential distribution characterized by the average time, 3) sampling from a truncated exponential distribution, and 4) simulating ion positions after a collision to determine the time of the next collision. The energy dynamics of the simplified simulation with these strategies are shown in Figure a.2(c), where each shaded region represents standard deviation around the mean ($\mu \pm \sigma$) of 50 simulation runs for each strategy. The change in energy is similar for all methods. However, the simulated ion position approach has significantly more computational overhead. In our work we elect to use the truncated exponential distribution function. While sampling from this distribution doesn't lead to a notable improvement of the simulation results compared to using the average time, it depicts a more accurate physical representation, at nearly no extra computational cost.




Figure a.2: Numerical analysis of collision time model. (a) Correlation plot of analytically and simulated obtained average collision times \bar{t}_{coll} for a randomized set of ion motion parameters $a_{i,k}$ and ϕ_k . The simulated data is derived from numerical ion dynamics simulations (triangles) and analytically approximated ion positions (circles). (b) Distribution of collision times, normalized to the average collision rate. For comparison, an (truncated) exponential distribution function is shown in blue (red). Various collision time strategies are used in the simplified rf heating simulation, of which the results are compared in (c).

TWO-ION HEATING MEASUREMENTS

This appendix provides an analytic expression that describes excitations of motional sidebands of a two-ion crystal. The time-dependent excitations provide an accurate way to calculate mean phonon numbers of individual modes, since their evolutions are sensitive to the mean phonon number. This analytic model has been used to fit experimental data of motional sideband excitations after a series of ion crystal rotations, discussed in Section 5.4. The content of this appendix is adapted from [88].

Measuring mode occupation numbers \bar{n} for low phonon numbers ($\bar{n} \lesssim 10$) is done by addressing motional sidebands at frequency $\nu = \omega_0 \pm \omega_k$, with ω_0 the frequency of a $|0\rangle \leftrightarrow |1\rangle$ carrier transition, and ω_k the motional mode frequency of mode k. The Jaynes-Cummings Hamiltonian

$$\hat{H}_{I_{\pm}} = \frac{1}{2}\hbar\Omega_{n\mp1,n}\sum_{i}\left(\hat{a}\sigma_{i}^{\pm} + \hat{a}^{\dagger}\sigma_{i}^{\mp}\right),\tag{b.1}$$

describes the interaction of this field with a red (blue) motional sideband, where the motional state $|n\rangle$ is occupied by n phonons. Here, $\Omega_{n-1,n}=\eta\sqrt{n}\Omega_{0,0}$ and $\Omega_{n+1,n}=\eta\sqrt{n+1}\Omega_{0,0}$ are the blue and red sideband coupling strengths, $\sigma_i^+=|1\rangle$ $\langle 0|$ and $\sigma_i^-=|0\rangle$ $\langle 1|$ the atomic transition operators acting on ion i, \hat{a} and \hat{a}^\dagger annihilation and creation operators. η is the mode-frequency dependent Lamb-Dicke parameter, in our case $\eta=0.061$ and 0.047 for the axial common and stretch modes. H_{I_+} and H_{I_-} represent interaction Hamiltonians on the blue and red sidebands, respectively.

The generalized ion state $|\phi(t)\rangle$ for two ions is given by

$$|\phi(t)\rangle = \sum_{n=0}^{\infty} \alpha_{00}^{(n)} |n\rangle |00\rangle + \alpha_{01}^{(n)} |n\rangle |01\rangle + \alpha_{10}^{(n)} |n\rangle |10\rangle + \alpha_{11}^{(n)} |n\rangle |11\rangle$$
(b.2)

with time-dependent coefficients $a_n^{(j)}$ obeying $\sum_n \sum_j |a_j^{(n)}|^2 = 1$, with $j \in \{00,01,10,11\}$. Given a symmetric starting state $(a_{01}^{(n)}\Big|_{t=0} = a_{10}^{(n)}\Big|_{t=0})$ and solely global operations, we can assume that $a_{01}^{(n)} = a_{10}^{(n)}$ for all t. Applying a beam on resonance with a red motional sideband transition to a pair of ions in the optical ground state with phonon number n, $|\phi(0)\rangle = |n\rangle |00\rangle$, the system will remain in the subspace

$$|\phi(t)\rangle = \alpha_{00} |n\rangle |00\rangle + \alpha_{01} |n-1\rangle |10\rangle + \alpha_{10} |n-1\rangle |01\rangle + \alpha_{11} |n-2\rangle |11\rangle,$$
(b.3)

dropping the superscript (n) for convenience. The time-dependent Schrödinger equation $i\hbar \frac{\partial}{\partial t} |\phi(t)\rangle = \hat{H}_{I_-} |\phi(t)\rangle$ is then:

$$\begin{split} i\hbar(\dot{\alpha}_{00}\left|n\right\rangle\left|00\right\rangle + \dot{\alpha}_{01}\left|n-1\right\rangle\left|10\right\rangle + \\ \dot{\alpha}_{10}\left|n-1\right\rangle\left|01\right\rangle + \dot{\alpha}_{11}\left|n-2\right\rangle\left|11\right\rangle) = \\ \frac{1}{2}\hbar\Omega_{n-1,n}\left[\alpha_{00}\left(\sqrt{n}\left|n-1\right\rangle\left|10\right\rangle + \sqrt{n}\left|n-1\right\rangle\left|01\right\rangle\right) + \\ 2\alpha_{01}\left(\sqrt{n-1}\left|n-2\right\rangle\left|11\right\rangle + \sqrt{n}\left|n\right\rangle\left|00\right\rangle\right) + \\ \alpha_{11}\left(\sqrt{n-1}\left|n-1\right\rangle\left|01\right\rangle + \sqrt{n-1}\left|n-1\right\rangle\left|10\right\rangle\right) \right] \end{split} \tag{b.4}$$

Combining identical states gives a set of coupled first order differential equations:

$$i\dot{\alpha}_{00} = \Omega_{n-1,n}\sqrt{n}\alpha_{01}$$

$$2i\dot{\alpha}_{01} = \Omega_{n-1,n}\left(\sqrt{n}\alpha_{00} + \sqrt{n-1}\alpha_{11}\right)$$

$$i\dot{\alpha}_{11} = \Omega_{n-1,n}\sqrt{n-1}\alpha_{01}$$
(b.5)

For initial conditions $\alpha_{00} = 1$, $\alpha_{01} = \alpha_{11} = 0$, the solutions are given by:

$$\alpha_{00}(t) = \frac{n - 1 + n\cos\left(\sqrt{n - \frac{1}{2}}\Omega_{-}t\right)}{2n - 1}$$

$$\alpha_{01}(t) = -i\frac{\sqrt{2n^{2} - n}\sin\left(\sqrt{n - \frac{1}{2}}\Omega_{-}t\right)}{\sqrt{2}(2n - 1)},$$

$$\alpha_{11}(t) = -\frac{2\sqrt{n^{2} - n}\sin^{2}\left(\frac{1}{2}\sqrt{n - \frac{1}{2}}\Omega_{-}t\right)}{2n - 1}$$
(b.6)

using the shorthand notation $\Omega_{\pm} = \Omega_{n\pm 1,n}$. A similar method is used to find the system equations for applying a beam on resonance with the blue motional sideband, which results in

$$\alpha_{00}(t) = \frac{2+n+(1+n)\cos\left(\sqrt{n+\frac{3}{2}}\Omega_{+}t\right)}{2n+3}$$

$$\alpha_{01}(t) = -i\frac{\sqrt{n+1}\sin\left(\sqrt{n+\frac{3}{2}}\Omega_{+}t\right)}{\sqrt{2}(\sqrt{2n+3})}$$

$$\alpha_{11}(t) = -\frac{2\sqrt{n+1}\sqrt{n+2}\sin^{2}\left(\frac{1}{2}\sqrt{n+\frac{3}{2}}\Omega_{+}t\right)}{2n+3}$$
(b.7)

In general, the excitation of multiple ions can be collectively expressed by a global excitation *E*, given by:

$$E = \frac{1}{N} \sum_{i=0}^{N} i p_i \tag{b.8}$$

where N is the total number of ions, and p_i is the population of i ions in the excited state. The collective excitation for two ions with n phonons is thus given by $E_n(t) = |\alpha_{11}(t)|^2 + |\alpha_{01}(t)|^2$, using $p_1 = |\alpha_{01}|^2 + |\alpha_{10}|^2$, and $\alpha_{01} = \alpha_{10}$. A thermally distributed mode has a phonon number n occupation probability given by

$$P_{\bar{n}}(n) = \frac{1}{\bar{n}+1} \left(\frac{\bar{n}}{\bar{n}+1}\right)^n, \tag{b.9}$$

for mean phonon number \bar{n} . The excitation for a thermally distributed mode is then

$$E_{\bar{n}}(t) = \sum_{n=0}^{\infty} P_{\bar{n}}(n) E_n(t)$$
 (b.10)

 $\Omega_{n\pm1,1}$ is calculated with Equation b.1, with $\Omega_{0,0}$ determined from a ground state-cooled carrier Rabi oscillation. $E_{\bar{n}}(t)$ is fit to experimental data of pulses on resonance with red and blue motional sidebands to determine the mean phonon number.

C

Chapter 6 discusses population transfer dynamics of multi-level systems, in which strong coupling between levels forms dressed states. In each of the examples discussed in Chapter 6, a specific initial qubit state is driven with a weak probe. The probe is detuned by Δ_{Ω} from the bare transition frequency, in accordance with a dressed state produced by coupling other transitions with additional beams. This appendix shows how the formation of dressed states can be derived, how appropriate detunings of the probe beam are obtained, and what the effective Rabi frequency between the initial states and dressed states are.

C.1 DRESSED STATE DYNAMICS OF ONE ION

Section 6.1.1 describes dressed state dynamics of one ion, with three levels (see Figure 6.2). Using the basis $\{|1\rangle, |f\rangle, |0\rangle\}$.

The dynamics following the dressed state splitting of one ion (see Section 6.1.1 and Figure 6.2) are derived from the Hamiltonian

$$H = \begin{pmatrix} 0 & \Omega_f & 0 \\ \Omega_f & 2\Delta_{\Omega} & g \\ 0 & g & 2\Delta_{\Omega} - 2\Delta_g \end{pmatrix}, \tag{c.1}$$

omitting the prefactor $\hbar/2$. Ω_f is the coupling strength of the weak probe, which couples $|0\rangle$ and $|f\rangle$ and g is the strength of the pump, which couples $|1\rangle$ and $|f\rangle$. The probe and pump have detunings Δ_{Ω} and Δ_{g} .

Analyzing the eigenvectors and eigenvalues of H reveals information about dressed states and their associated resonance frequencies. In particular, we can establish which basis of states is formed by setting the probe power to zero, $\Omega_f = 0$, which is then given by the eigenvectors of H, ignoring the normalization factors,

$$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} \Delta_g - \sqrt{\Delta_g^2 + g^2} \\ g \end{pmatrix}, \quad \begin{pmatrix} 0 \\ \Delta_g + \sqrt{\Delta_g^2 + g^2} \\ g \end{pmatrix}, \quad (c.2)$$

with associated eigenvalues

0,
$$2\Delta_{\Omega} - \Delta_g - \sqrt{\Delta_g^2 + g^2}$$
, $2\Delta_{\Omega} - \Delta_g + \sqrt{\Delta_g^2 + g^2}$. (c.3)

The interpretation is that applying g produces a new set of basis states, though includes the initial state $|0\rangle$. Resonant excitation from the initial state occurs when its eigenvalue is degenerate with the eigenvalue of one of the other basis states:

$$0 = 2\Delta_{\Omega} - \Delta_{g} \pm \sqrt{\Delta_{g}^{2} + g^{2}} \tag{c.4}$$

which is precisely the condition that produces the avoided-crossing excitation curves of Figure 6.2(b). We see, for example, that if the pump detuning is large, $\Delta_g\gg g$, there exist two cases of resonance. $\Delta_\Omega=0$ gives resonant Rabi cycling between $|0\rangle$ and $|f\rangle$, which is just a run-of-the-mill two-level Rabi oscillation. $\Delta_\Omega=\Delta_g$ produces resonant Rabi cycling between $|0\rangle$ and $|1\rangle$, which is known as a Raman transition.

For the engineered resonance scheme, we set the pump detuning to zero, $\Delta_g = 0$. The eigenvectors reduce to

$$\begin{pmatrix} 1\\0\\0 \end{pmatrix}, \quad \frac{1}{\sqrt{2}} \begin{pmatrix} 0\\-1\\1 \end{pmatrix}, \quad \frac{1}{\sqrt{2}} \begin{pmatrix} 0\\1\\1 \end{pmatrix}, \tag{c.5}$$

and the resonance condition reduces to $\Delta_{\Omega}=\pm g/2$. A weak pump beam with this detuning then produces Rabi cycling between $|0\rangle$ and $1/\sqrt{2}(|f\rangle \mp |1\rangle)$.

This set of eigenvectors represents a new basis of states, $\{|0\rangle, |-\rangle, |+\rangle\}$, with $|\pm\rangle = 1/\sqrt{2}(|1\rangle \pm |f\rangle)$. To find the effective coupling strength between these states, we can express the Hamiltonian in terms of this new basis, using $\tilde{H} = R^{-1}HR$, with the matrix $R = (|0\rangle, |-\rangle, |+\rangle)$.

Using $\Delta_{\Omega} = g/2$ and $\Delta_g = 0$, the Hamiltonian now reads

$$\tilde{H} = \begin{pmatrix} 0 & -\Omega_f / \sqrt{2} & \Omega_f / \sqrt{2} \\ -\Omega_f / \sqrt{2} & 0 & 0 \\ \Omega_f / \sqrt{2} & 0 & 2g \end{pmatrix}.$$
 (c.6)

For a large pump coupling strength, $g \gg \Omega_f$, transfer to the state $|+\rangle$ is far off-resonance, thus will not be populated. Rabi cycling occurs between $|0\rangle$ and $|-\rangle$, with a coupling strength $\Omega/\sqrt{2}$.

In Section 2.2.3, we described what happens to the excited state population P_{exc} of Rabi oscillations with Rabi frequency Ω if the light-field is detuned from resonance by Δ :

$$P_{\rm exc} = \frac{\Omega^2}{\Omega^2 + \Delta^2} \sin^2\left(\frac{\sqrt{\Omega^2 + \Delta^2}t}{2}\right) \tag{c.7}$$

Substituting $\Omega = \Omega_f / \sqrt{2}$ and $\Delta = \Delta_f \pm g/2$, we find an excitation spectrum given by the sum of the two plus/minus versions of

$$\frac{\Omega_f^2}{\Omega_f^2 + 2(\Delta_{\Omega} \pm g/2)^2} \sin^2 \left(\frac{\sqrt{\Omega_f^2/2 + (\Delta_{\Omega} \pm g/2)^2} t}{2} \right).$$
 (c.8)

C.2 DRESSED STATE DYNAMICS OF TWO IONS

The same analysis can now be used for the extension to two ions, as presented in Section 6.1.2. The starting state $|00\rangle$ follows identical dynamics as the three-level system in the previous section, just with a different state notation. From the starting state $|01\rangle$, a total of four states are reachable, making up the basis $|01\rangle$, $|f1\rangle$, $|11\rangle$, and $|1f\rangle$ with the Hamiltonian

$$H = \begin{pmatrix} 0 & \Omega_f & 0 & 0 \\ \Omega_f & 2\Delta_{\Omega} & g & 0 \\ 0 & g & 2\Delta_{\Omega} - 2\Delta_g & g \\ 0 & 0 & g & 2\Delta_{\Omega} \end{pmatrix}.$$
 (c.9)

As before, we analyze the eigenvectors and eigenvalues of the system under the condition that the probe is off, thus exposing the dressed states and their energies. We will only consider the case where the pump detuning is zero, $\Delta_g = 0$. The eigenvectors are then

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ -1 \\ 0 \\ 1 \end{pmatrix}, \frac{1}{2} \begin{pmatrix} 0 \\ 1 \\ -\sqrt{2} \\ 1 \end{pmatrix}, \frac{1}{2} \begin{pmatrix} 0 \\ 1 \\ \sqrt{2} \\ 1 \end{pmatrix}, \tag{c.10}$$

with associated eigenvalues

$$\{0, 2\Delta_{\Omega}, 2\Delta_{\Omega} - \sqrt{2}g, 2\Delta + \sqrt{2}g\}.$$
 (c.11)

Resonance from the initial state $|01\rangle$ occurs when its energy is degenerate with another one of the dressed states. We thus have three possibilities for resonant excitation from $|01\rangle$, occurring at $\Delta_{\Omega}=0$ and $\Delta_{\Omega}=\pm g/\sqrt{2}$, which can also been seen in the right plot of Figure 6.3(c).

We denote the normalized eigenvectors as $\{|0\rangle, |f^+\rangle, |-\rangle, |+\rangle\}$. Casting these vectors into the matrix R, the Hamiltonian can be expressed in terms of the basis of dressed states as

$$\tilde{H} = R^{-1}HR = \begin{pmatrix} 0 & -\Omega_f/\sqrt{2} & \Omega_f/2 & \Omega_f/2 \\ -\Omega_f/\sqrt{2} & \sqrt{2}g & 0 & 0 \\ \Omega_f/2 & 0 & 0 & 0 \\ \Omega_f/2 & 0 & 0 & 2\sqrt{2}g \end{pmatrix},$$
(c.12)

where we have set the detuning $\Delta_{\Omega}=g/\sqrt{2}$. For $g\gg\Omega_f$, we have resonant Rabi oscillations between $|0\rangle$ and $|-\rangle$, with a Rabi frequency of $\Omega_f/2$. The excitation profile is found as before using Eq. c.7, with the substitutions $\Omega=\{\Omega_f/\sqrt{2},\Omega_f/2,\Omega_f/2\}$ and $\Delta=\Delta_{\Omega}+\{0,+g/\sqrt{2},-g/\sqrt{2}\}$ given by the sum of

$$\frac{\Omega_f^2}{\Omega_f^2 + 4(\Delta_\Omega \pm g/\sqrt{2})^2} \sin^2\left(\frac{\sqrt{\Omega_f^2 + 4(\Delta_\Omega \pm g/\sqrt{2})^2}t}{4}\right) \tag{c.13}$$

and

$$\frac{\Omega_f^2}{\Omega_f^2 + 2\Delta_\Omega^2} \sin^2\left(\frac{\sqrt{\Omega_f^2 + \Delta_\Omega^2/2}t}{2}\right). \tag{c.14}$$

C.3 DRESSED STATE DYNAMICS OF MINIMAL INSTANCE ERROR CORRECTION

We now analyze the dynamics of the minimal instance error correction scheme described in Section 6.2.1, whose levels schemes are shown in Figure 6.5. We follow the same procedure as the previous sections.

Starting states $|01\rangle$ and $|10\rangle$ follow the same dynamics as in the previous section, requiring detunings $\Delta_{\Omega}^f = g_f/\sqrt{2}$ and $\Delta_{\Omega}^e = g_e/\sqrt{2}$. Rabi cycling occurs at a frequencies $\Omega_f/\sqrt{2}$ and $\Omega_e/\sqrt{2}$.

¹ Though not explicitly shown here, the excitation peak at $\Delta_{\Omega}=0$ has a Rabi frequency of $\Omega_f/\sqrt{2}$, which can be obtained from \tilde{H} with $\Delta_{\Omega}=0$.

The dynamics of the starting state $|00\rangle$ are derived from the Hamiltonian:

$$H = \begin{pmatrix} 0 & \Omega_f & 0 & 0 \\ \Omega_f & 2\Delta_{\Omega}^f & g_f & 0 \\ 0 & g_f & 2\Delta_{\Omega}^f & g_e \\ 0 & 0 & g_e & 2\Delta_{\Omega}^f \end{pmatrix}, \tag{c.15}$$

using the basis $\{|00\rangle|0\rangle_n$, $|f0\rangle|0\rangle_n$, $|10\rangle|1\rangle_n$, $|1e\rangle|0\rangle_n$, and assuming that g_f and g_e are both on resonance. Non-normalized eigenvectors are

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -g_e/g_f \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ g_f/g_e \\ -\sqrt{g_f^2 + g_e^2}/g_e \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ g_f/g_e \\ \sqrt{g_f^2 + g_e^2}/g_e \\ 1 \end{pmatrix}, \tag{c.16}$$

with eigenvalues

$$\{0, 2\Delta_{\Omega}^{f}, 2\Delta_{\Omega}^{f} - \sqrt{g_{f}^{2} + g_{e}^{2}}, 2\Delta_{\Omega}^{f} + \sqrt{g_{f}^{2} + g_{e}^{2}}\}$$
 (c.17)

leading to resonance at $\Delta_{\Omega}^f=0,\pm\sqrt{g_f^2+g_e^2}/2$. Choosing $\Delta_{\Omega}^f=+\sqrt{g_f^2+g_e^2}/2$, the Hamiltonian in the basis of dressed states is

$$\tilde{H} = \begin{pmatrix} 0 & -\Omega_f g_e / \sqrt{g_f^2 + g_e^2} & \Omega_f g_f / \sqrt{2g_f^2 + 2g_e^2} & \Omega_f g_f / \sqrt{2g_f^2 + 2g_e^2} \\ -\Omega_f g_e / \sqrt{g_e^2 + g_f^2} & \sqrt{g_f^2 + g_e^2} & 0 & 0 \\ \Omega_f g_f / \sqrt{2g_f^2 + 2g_e^2} & 0 & 0 & 0 \\ \Omega_f g_f / \sqrt{2g_f^2 + 2g_e^2} & 0 & 0 & 2\sqrt{g_f^2 + g_e^2} \end{pmatrix}, \text{ (c.18)}$$

from which we see Rabi cycling between $|0\rangle$ and the third basis state in Eq. c.16, with Rabi frequency $\Omega_f g_f / \sqrt{2g_f^2 + 2g_e^2}$. Excitation from $|11\rangle$ follows the same dynamics, though with g_f and g_e swapped. The results in this section are actually a generalized version of the results in the previous two sections: setting $g_e \to g_f$ reproduces the results from Section c.2, and setting $g_e \to 0$ reproduces those from Section c.1.

d

The ultimate results of the classical gates using dissipation through engineered resonance in Section 6.3.2 are the measured truth tables shown in Figure 6.13. These truth tables are in reasonable agreement with the desired output (87% and 81% population fidelity for the OR and NOR gates, respectively). Section 6.3.3 discusses the main sources of error that lead to imperfections in the gate's process fidelity. This Appendix provides a quantitative overview of simulations that isolate the individual sources of error. A qualitative description of these errors is provided in Section 6.3.3.

Simulations are run with Python's Qutip package [250]. The simulation tracks the quantum dynamics of two four-level ions, with one shared mode of motion. The four levels of each ion are denoted $|0\rangle$, $|1\rangle$, $|e\rangle$, following the description of states in Section 6.3.1, schematically shown in Figure 6.8. The default simulation settings are given in Table d.1

Table d.1: Default parameters of the engineered resonance gate simulations

Parameter	Value
Pump coupling (g_f)	800 kHz
Probe coupling (Ω_f)	1.15 kHz
Dissipation rate (Γ_f)	$9\mathrm{ms}^{-1}$
Transfer coupling (Ω_e)	10 kHz
Initial mode occupation (\bar{n}_{init})	0 phonons
Heating rate $(\dot{\bar{n}})$	0 phonons per second
Sideband frequency (Ω_{SB})	550 kHz
Engineered state transfer duration, OR	909 μs
Engineered state transfer duration, NOR	642 μs
Dissipation duration	1 ms

The settings in Table d.1 are idealized, in the sense that the state population after the gate matches 100% with the desired output for each initial state.

We run the simulation for each of the four initial states ($|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$), for both gate types (OR and NOR). The simulation is repeated for isolated sources of error. A description of these errors, and their implementation in the simulation, is given below. The values used in the error simulations reflect those that we have experimentally determined.

- Ideal The default simulation settings are used, which produces an optimal gate.
- *Heating rate* The motional mode heating rate is set to $\dot{n} = 106$ phonons per second. The initial phonon number is kept at zero.
- *Initial phonon number* The phonon number at the start of the simulation is set to 0.14. The heating rate remains zero.
- Sideband cooling strength The ideal simulation has a high enough sideband cooling strength Γ_f such that any non-zero mode occupation after the state transfer step is fully depleted during the 1 ms dissipation step. In this simulation, we set $\Gamma_f/(2\pi)=4.5\,\mathrm{kHz}$.
- Finite pump coupling The pump coupling in the ideal simulation is set to be 100 times higher than the value used in our experiment. This ensures that the condition $g_f \gg \Omega_f$ is met, which minimizes off-resonant excitation of the undesired engineered resonance transfer processes. In the simulation, we set $g_f/(2\pi)=8\,\mathrm{kHz}$.

- Off-resonant carrier excitation The pump beam with coupling strength 8 kHz is detuned by 550 kHz from the carrier transition, which makes off-resonant excitation of the carrier nonnegligible. This is included in the simulation by including a detuned coupling of the $|1\rangle \leftrightarrow |f\rangle$ transition, with coupling strength $g_f/\eta=142$ kHz (with the Lamb-Dicke parameter $\eta=0.056$). Since the simulation of off-resonant carrier coupling require a realistic value of g_f , the errors include those of the 'finite pump coupling' described in the previous bullet point. Since this off-resonant coupling generates a stark shift on the $|0\rangle$ and $|1\rangle$ states, the detunings of the pump and probe beam have been appropriately compensated.
- *Total* All the above errors are included.

Figures d.1 through d.4 give an overview of the simulation results under various sources of error for the OR and NOR gates. The populations of the states $|0\rangle$, $|1\rangle$, and $|f\rangle$ are plotted for each ion individually, as a function of time during the engineered resonance state transfer pulse, and the dissipation pulse. The initial and final transfer pulses to and from $|e\rangle$ are not shown. Additionally, the population of the motional ground state, $P_{n=0}$ is shown. The final simulated input-to-output truth-table is shown on the right.

Figure d.5 summarizes the contribution of various sources of error, displayed as the absolute difference between the desired output population and the simulated output.




Figure d.1: Simulated engineered resonance OR gates with various sources of error




Figure d.2: Simulated engineered resonance OR gates with various sources of error (cont.)

NOR




Figure d.3: Simulated engineered resonance NOR gates with various sources of error

Time (µs)

Input

NOR (cont'd)




Figure d.4: Simulated engineered resonance NOR gates with various sources of error (cont.)




Figure d.5: Overview of engineered resonance gate errors

e

NOTES

- 1. Page 2 I'm not a big fan of the term QCCD to describe a segmented surface trap quantum computer. The acronym feels to me like a misnomer introduced as a buzz-word. None of the definitions of CCD that a quick Google search provide mimic the functionality proposed by the QCCD. Wikipedia (although no sources are cited) has the most generic definition: a set of capacitors that are able to transfer electric change between neighboring electrodes. Most other definitions place CCDs in the context of cameras, as being an array of light-sensitive metal oxide semiconductor (MOS) electrodes. I guess QCCD just refers to the fact that the trap is chip-like with many separated voltage carriers. It is, nevertheless, worth introducing the term QCCD, since that has become the default nomenclature.
- 2. Page 3 Customarily, it is asked to paint a literary "red line" in the introduction of a thesis that form a logical connection between works done during the PhD. I would forego that formality and just present each project as self-standing and in it's own way some contribution to the field of quantum computation, without creating a forced narrative to link it all. Although, it was suggested to me that this thesis would be more enjoyable if chapters were bridged with plot twists and cliffhangers.
- 3. Page 6 In a lot of literature on ion trapping, the equations for the Mathieu a and q parameters are written in terms of an electrode voltage and the distance between the trap center and the electrode, instead of field curvatures α_k and β_k , as I have done. Because electrodes are not perfectly hyperbolic and do not extend off to infinity, representing the a and q parameters in terms of electrode distance and voltage is an approximation. Often an additional geometric factor is introduced to correct for this discrepancy. For 3D blade traps, this factor is approximately 1. In segmented traps, or worse, segmented *planar* traps, this metric of the a and a parameter is less meaningful, since the geometric factor can be quite large and and potentials are generated by a large combination of electrode voltages.
- 4. Page 6 Most documents introduce ion trapping by taking the step from the Mathieu differential equation the stable solution equation of motion. In my opinion, this is basically answering the question "Why does an RF harmonic potential generate stable confinement?" with "Because the math says so." which I find pretty unsatisfying. Intuitively, it's not straightforward to reason why an electric field that is on average zero should produce a restoring force. The hand-wavy way that I look at it: Consider that the force from the RF field that drives the ion is both position and time dependent. The further away an ion is from the trap center, the higher the amplitude of the force. The crucial insight then comes by considering the force on an ion in a single period of the RF drive. Here, within our usual stability region ($q_k < 0.9$), the amplitude of an ion's motion is small compared to it's displacement from the minimum of the harmonic potential. During one oscillation, an ion moves a bit further from the center, and then a bit closer, exactly out of phase with the force being applied to it. The force of the RF field is slightly higher when the ion is further away in this cycle. The restoring force is therefore slightly higher than the repulsive force. The overall restoring force that makes up an ion's secular motion comes from this difference.
- 5. Page 6 The analytic expression for ω_k is given by

$$\omega_k = D_k \Omega_{\mathrm{RF}}$$
 (e.1)

where D_k can be found recursively using the continued fraction

$$D_k^2 = a_k + \frac{q_k^2}{(D_k + 2)^2 - a_k - \frac{q_k^2}{(D_k + 4)^4 - a_k - \dots}} + \frac{q_k^2}{(D_k - 2)^2 - a_k - \frac{q_k^2}{(D_k - 4)^4 - a_k - \dots}}.$$
 (e.2)

However, having the combination of both recursion and infinity division makes it not so straightforward to solve. I used the ansatz of $D_k^{(0)} = (1/2)\sqrt{a_k + q^2/2}$, and worked out the recursive function up until the 10th power, to solve for $D_k^{(1)}$. This then is used to solve for $D_k^{(2)}$, and so on. However, this method does not work for all values of a_k and q_k . As it turns out, parameter regions where this method does not result in complex numbers, and converges to a finite value, correspond exactly with the Mathieu stability diagram.

- 6. Page 6 I describe the pseudopotential as being an effective field that would explain the the derived secular motional frequency ω_k . In most literature, this is done the other way around: The pseudopotential is derived by integrating average forces due to the oscillating field. The secular frequency is then calculated as consequence of the pseudopotential.
- 7. Page 14 The Bloch sphere is almost universally depicted with $|0\rangle$ on the north-pole, and $|1\rangle$ on the south. Accordingly, $|0\rangle$ often is taken to represent the higher-energy excited state. Although the labeling convention is arbitrary, this choice always feels awkward to me. Zero is, after all, less than one. In this thesis, $|0\rangle$ is always taken to be the ground state.
- 8. Page 40 My first weeks in the group were hectic: to assure renewal of funding for our project, it was required to demonstrate that the experimental setup was at a stage that ions could be trapped. The review meeting was at the end of May, and I arrived at the start of May to an experiment where already for several months attempts were made to trap ions. My first weeks were spent working on helping Matthias and Rafael swap out traps in the experiment. Towards the end of the month, literally a few hours before the review meeting during which the project's progress was discussed, we were able to send pictures of trapped ions.
- 9. Page 40 Fun anecdote: There were a few issues with the version of the Berkeley trap that was shipped to us, which required us to re-coat the silica in-house from scratch. This led to a pretty embarrassing first impression story: During one of our group's retreats, I fully let loose in pointing out in what state the trap was sent to us. It was already mentioned to me that the project would be getting a new supervisor, his term starting after the group retreat. What was not mentioned to me, is that Philipp

Schindler had spent the past year and half in Berkeley, nor was it mentioned that he was sitting in the audience during the group retreat. So after about the third slide of me complaining about the trap, Philipp spoke up with a hint of outrage. My casual reply was "You're not from Berkeley, are you?" We got awkwardly formally introduced directly after the talk.

10. Page 43 If the contents of this thesis were distributed according to real-world time spent on each topic, the topic of acquiring liquid Helium would take up two or three chapters. As it turns out, Helium is already scarce and expensive enough in gaseous form, let alone in liquid form. Simply ordering liquid Helium to be delivered to the University on a regular basis was not an option.

Initially, the plan was to order a transport van on a bi-weekly or monthly basis, load it with empty dewars and a bundle of recovered gaseous helium, and drive it to Garching, Germany. There, at the Walther-Meißner-Institut, a two hour drive away, gaseous Helium is deposited, liquid Helium is filled, and one-to-two hours away we'd be on our way for another two-hour drive back to Innsbruck. Not to mention the time-budget involved in wheeling over the gas-bundles to and from an industrial forklift to get it loaded and unloaded from the van. All-in-all, a full day's work to get a full dewar of Helium, including the van pick-up and drop-off on the opposite side of Innsbruck. Not the best investment of a Ph.D. student's time.

Fortunately, our group not being the only one requiring liquid Helium, our institute decided to purchase a second-hand liquefier, so that our recovered gaseous Helium could be re-liquified on campus. What a time saver!

Unfortunately, it wasn't. The liquefaction plant still took an *immense* amount of time for regular operation, and that's not including the amount of time spent on repairs, servicing, and refilling with newly ordered gaseous helium. The most intense repair occurred in early 2021, when the liquefier's expansion machine malfunctioned. New strict Covid border restrictions prevented the company that supplied the liquefier to drive over from Germany for repairs. Instead, we got a Word document that detailed how to take apart the liquefier ourselves to extract the expansion machine, which we loaded into a van and drove over to Germany. When I arrived at the company on a Friday at 16:00, there was no one there. So I rang the neighbors doorbell, a farmer by the looks of it, and dropped off the expansion machine in his barn. All ended well with those repairs, but it took several months, obscene helium losses, and an immense bill before we were back in full operation.

- 11. Page 49 The toroidal wire coil has served us reliably for multiple traps. For a brief time we also tried another interesting option: a yttrium barium copper oxide (YBCO), high-temperature superconducting, spiral coil. It was a neat idea, since at cryogenic temperatures the coil's resistance would be zilch. However, the thing turned out to be a nightmare. As thermal stress easily destroys the high-temperature superconductor (HTS) properties, we could not solder wire connectors onto the YBCO coil with regular tin-based flux, but used indium instead. Indium melts at much lower temperatures than tin, but there's a reason its not used regularly for soldering. It does not 'latch' to metal properly, making it almost useless to solder with. In the end we managed to make a connection by sticking the connecting wire into a excessively large blob of indium, but destroyed two YBCO coils in the process. Also, having indium as a solder joint meant that we were restricted in to what temperature we could bake our setup after closing. Another problem was that we couldn't test the resonance properties or the matching network outside of vacuum at room temperature. Dipping it in a liquid nitrogen bath gets us some information, but were unable to extrapolate to 30 K operation, like we did with the other coils. So we had to ballpark the capacitances of the matching network. While we weren't too far off in the end, the advantage of having a potentially high-Q resonator was subdued by having imperfect matching. In the end, the YBCO resonator simply gave up one day and stopped resonating, and was promptly replaced with a wire coil.
- 12. Page 54 We have an in-vacuum lens with a ridiculously high NA, custom milled slits, placed on a >15 kEuro translation stage stack, shielded from the trap by a custom gold-wire mesh. What do we use it for? Check to see if our beams are aligned properly. Yes, this endeavor was a bit of an overkill. On the other hand, we did use the lens to show that state detection can be done from ions trapped in a cryogenic surface trap within 50 µs. This is, at the very least, a solid proof-of-principle demonstration that such an optical setup is compatible with surface traps in cryogenic environments.
- 13. Page 54 Word count in this sub-section for the word 'control,' or some variation thereof: 21
- 14. Page 56 One critical difference between the Berkeley trap and the Golden Gate trap is the way they were/are mounted onto the trap holder. The Golden Gate trap is wire-bonded to a chip carrier, which is connected to the trap mount through a ceramic pin grid array (PGA) interposer, which electrically connects to the trap PCB; a simple design carried over from the previously installed Sandia trap. The Berkeley trap, however, was directly wirebonded onto the trap PCB, with the trap itself sitting on a trapezoidal pedestal. It looks rather cool, like the trap is sitting on an elevated throne. The thing we didn't take into account with this design is the travel range of the wire-bonding machine, which should be able to bridge the height gap between the PCB and the top of the pedestal. It didn't. We therefore resorted to mounting the PCB onto translation stage that would allow the assembly to move up and down to help bridge this gap. So for each bond wire, the micrometer positioner needed to be turned to move the stage up by about a centimeter, and then back down for the next wire. With each pad double-bonded, we needed to twist the stage more than 100 times up and down. And with 'we,' I mean Kirill Lakhmanskiy, who gets full credit and acknowledgment for helping us not only with bonding the trap, but also to clean and coat it.
- 15. Page 60 After gold coating, two or three electrodes were shorted with their neighbors. This is apparently not entirely uncommon, as the deposition flux is not 100% collimated. Alternatively, some grains or chunks between electrodes may bridge them electrically. Either way, such shorts are not necessarily mean a lost cause for the trap, as there are subtle and not so subtle ways to remove them. If there are visible (by microscope) grains near the surface, they can still be removed mechanically with ultra-thin plastic filament, and a steady hand. If whatever causes the short is deeper down in the trench, this could be incinerated by passing a current through it. Two fine probes touch the two neighboring electrodes, connected by a current limited voltage supply. In the best case scenario, the cause of the short has the relatively highest resistance in this circuit, heats up, and disintegrates. This is accompanied by a very satisfyingly small *crack*, and a drop in current to zero. In the worst case scenario, and certainly not a scenario that hasn't happened in the past, a mini-explosion causes the electrode coating to completely flake off.
- 16. Page 69 A seemingly more descriptive way to put it is non-periodic Coulomb forces. However, this phrasing would suggest that the Coulomb interaction is periodic when ions are crystallized. Counter-intuitively, this is not the case: the Coulomb potential is a function of three periodic signals, but sums and products of of periodic signals are not periodic, unless the individual periods are rationally related. So a safer way to describe the distinction between the Coulomb potential in an ion crystal versus an ion cloud is with regular and irregular. Another way to characterize the difference is to use the term chaotic for an ions motion in an ion cloud, whereas motion in a crystal is predictable.

- 17. Page 70 I've been warned to steer clear of using the term chaotic, because apparently this term is reserved for mathematicians who know what they're talking about. However, I've seen the term pop up in enough ion cloud literature [138, 150] and I've ran numerical tests for chaos on the simulation data [251], so I feel the term's usage is justified. I take it to mean that if I run the same simulation twice with marginally different starting values, the motion of ions between the two simulations are quickly uncorrelated. Comically enough, the article on 'Chaos' on MathWorld [252], spends most of its introduction on describing how ill-defined chaos is in mathematics, and concludes with "So a simple, if slightly imprecise, way of describing chaos is 'chaotic systems are distinguished by sensitive dependence on initial conditions and by having evolution through phase space that appears to be quite random.' " That fits the bill in our simulations, so I'm going to stick with that terminology.
- 18. Page 71 Other literature points [151] to another possible discrete transition between crystal and cloud. This transition can occur with a multi-ion crystal that expands in a direction with an RF potential (i.e., not a linear chain along the RF-free axial direction). The Mathieu stability conditions for the trap control parameters (the *q* and *a* parameters) for such a crystal are not the same as for a single trapped ion. If one were to change the trap control parameters (for example, by increasing the trap drive power), the crystal is brought outside of the stability region, and it 'shatters.' Individual ions have more relaxed Mathieu stability conditions, and thus remain trapped, albeit in a melted state. This kind of melting event happens under very specific conditions, and are not applicable to the trap parameters described in this section.
- 19. Page 72 Generally speaking, RF *heating* is an energy exchange between secular and RF energy. It can thus, in theory, just as well lead to cooling. However, the average trend is for secular energy to increase. So where, then, is the break in symmetry that leads to heating, on average? An intuitive answer is that energy is an absolute property that can't drop below zero; randomly adding or removing energy, starting from a low energy with the boundary condition that energy can't be negative, will thus on average lead to an increase in energy, i.e. heating. The same argument can be made for surface noise-induced heating. This argument requires a bit more nuance, since ions require a non-zero threshold energy in order for RF heating to kick in, but in essence remains similar: RF heating shuts off below a threshold, so on average the process that adds energy is favored.

Unlike surface-noise induced heating, which reaches a thermal equilibrium with the electric field "bath," I don't suppose that RF heating leads to a thermal equilibrium: If we consider the "bath" to be the energy available in E_{RF} and note that the energy of this bath scales roughly with the square of the displacement of the ion, we notice that while the ion gains energy, the bath gets hotter. There is therefore no theoretical upper limit to the heating process of RF heating, except, of course, the trap depth.

- 20. Page 80 There is actually nothing surprising about the fact that collisions seem to occur only near the trap center: the two ions' center-of-mass motion follows a trajectory that is identical to that of a single trapped ion. This trajectory is not affected by the chaotic behavior of the ions, nor by RF heating. The amplitudes of this trajectory thus remain constant indefinitely. These amplitudes are set by the background-particle collision. A typical collision (usually with a hydrogen molecule) induces oscillation amplitudes of no more than a few micrometers. When ions subsequently collide with each other, naturally they are both close to their center-of-mass, which is restricted to being close to the center of the trap, even as the individual ion amplitudes grow into tens of micrometers. It's rather surprising to see that in a simulation of an ion cloud, properties like ion position, separation, and energy all behave so chaotically, while there's an underlying property, center-of-mass motion, that remains very stable. In hindsight the stability of this property is an indication that the numerical simulation is working properly.
- 21. Page 81 A majority of collisions lead to a in my words near-negligible energy changes. I also claim that collisions that have a peak Coulomb energy below a threshold of 0.5 meV result in a negligible change in energy. Furthermore, though not plotted, the majority of collisions occur with peak Coulomb energies below this threshold. One might wonder that since an overwhelming number of collisions occur below this threshold, that the sheer number of 'negligible' energy changes might still result in an overall energy change, and should thus be taken into account. I wanted to avoid polluting this section with too many statistical checks, but you can rest assured that this question has been checked and debunked: the overall change in energy comes predominantly from collisions above the peak Coulomb collision threshold.
- 22. Page 86 It's a seemingly simple problem: there are three pulse waves, each with a known frequency. What then, is the (average) frequency of the product of those pulse waves? Coming up with a formal answer has been less trivial than expected. Firstly, it isn't obvious that the product of multiple periodic functions isn't necessarily periodic. The product will only be periodic if the ratio of periods of the separate functions is rational (or in other words, if the periods share a common multiple). The periods of secular motion of ions can not be assumed to have this relation. Despite this non-periodicity, the combined function still exhibits regularity, so one would expect a simple derivation to find it's average pulse rate.

In the end, neither extensive Googling, asking of colleagues, nor putting pencil to paper resulted in a mathematically sound derivation. Appendix a shows how the solution is logically deduced using a geometric argument, of which I have mixed feelings: on the one hand, it works, and that should be all that matters. On the other hand, it would have been more satisfying to back up with a proper mathematical derivation.

—Last-minute update— My older brother, while reading through my thesis, came up with a much more intuitive derivation of the mean collision rate than the rather convoluted geometrical proof of Appendix a (I mean really... three time dimensions??). Instead of considering a collision to be the single case where $B_xB_yB_z=1$, one could analyze three independent scenarios that constitute a collision: 1) B_x changes from 0 to 1 under the condition that $B_yB_z=1$, 2) B_y changes while $B_xB_z=1$, and 3) B_z changes while $B_xB_y=1$. Since these scenarios are independent, the total average collision rate $\bar{f}_{\rm coll}$ is the sum of the mean rates of each, which I'll refer to as \bar{f}_x , \bar{f}_y , and \bar{f}_z . Let's look at \bar{f}_x : A change of B_x from 0 to 1 occurs at a frequency of $1/T_x$. The probability that both B_y and B_z are 1 during a randomly selected B_x change is given by the product of the individual probabilities, $(\Delta t_y/T_y) \cdot (\Delta t_z/T_z)$. The rate \bar{f}_x is therefore $\frac{1}{T_x} \left(\frac{\Delta t_z \Delta t_z}{T_y T_z} \right)$. It's then straightforward to see that $\bar{f}_x + \bar{f}_y + \bar{f}_z$ is the same as the collision rate given in Eq. 4.50. Note, as was the case in Appendix a, this derivation still requires that T_x , T_y , and T_z are irrationally related.

23. Page 86 Simplified ion dynamics simulations are initialized with ions being given random amplitudes. The way we've chosen these amplitudes, uniformly at random, may not be physically accurate. After a few simulated collisions, these values get scrambled enough that this initial distribution is irrelevant. While it wouldn't have affected the outcome of the simulation if all the amplitudes $a_{i,k}$ were set to be equal, or if all $(1/2)m\omega_k^2 a_{i,k}^2$ were set to be equal, or if only one of the two ions was given an initial motional energy, it seemed fairer to give them some initial distribution.

- 24. Page 99 As nice as it would be to show the scaling of $D \propto \omega_r^a$, with a = 2.5 as shown in Table 4.1, with three data points and their errors, it is rather pointless. If one stubbornly attempts to fit these data points, however, the exponent a = 2.5 does lie within the error range of that fit result.
- 25. Page 101 The analogy that I like to draw is the RF heating models that I've presented are like understanding climate instead of weather. The models are not useful for predicting rapid energy changes that ions will undergo during RF heating, but more for estimating an average trend of energy change. In the same way, certain meteorological models are not intended to predict if it's going to be warm next week, but rather what global warming will do to the climate over the next decades, which incidentally is also a study of heating rates of a chaotic system.
- 26. Page 103 Collision times are likely dominated by the motional periods of the lighter mass. If we assume the limiting case that heavier mass is *much* heavier than the lighter one, the lighter ion is oscillating much faster than the heavier one. In this case, when calculating collision times, we could model the heavier one as effectively standing still, and ignore whatever frequency it is oscillating at. This is similar to setting the motional amplitudes $a_{i,k}$, with i the index of the heavier ion, to zero. Collision times then follow directly from 4.50, and is only dependent on the motional frequencies of the lighter ion. This is speculative, and remains to be confirmed by simulation.
- 27. Page 104 A common misconception about ion traps is that at the center of the trap, the RF-null, the electric field potential is zero and constant, as it is often drawn to be the case in the oscillating saddle potential picture. In fact, the potential here, due to an applied RF voltage, is roughly half of the potential at the RF blades, since this potential component is zero at the remaining blades. This means that the potential at the RF-null is oscillating at half the amplitude of the applied voltage. From an ions point of view, of course, this is irrelevant since it only experiences a force due to a potential gradient, $F = -q\nabla V$.
- 28. Page 115 The operation is generally referred to as crystal separation, or ion chain splitting, but when speaking colloquially in the lab I often simplify it to "ion splitting." That terminology can be humorously misunderstood to refer to a fission process, which if I'm not mistaken would turn ⁴⁰Ca⁺into Argon. It would not surprise me, however, if that term accidentally found its way into this thesis in a few places.
- 29. Page 125 The axes are labeled potential curvature because this is what the second-order $Y_{l,n}$ represent. This can admittedly be misleading, since in reality the potential curvature of each $Y_{2,n}$ term depends on along which axis you describe it (for example, the $Y_{2,0}$ term has half the field curvature in the two radial directions as it does in the axial direction). The axis describes the value $v_{2,n}$, being the contribution of each potential in the linear combination $V_{DC} = \sum_n v_{2,n} Y_{2,n}$. Should you wish to see what $Y_{2,0}$ actually means in terms of terms of curvature in the axial direction ψ_z , we can use

$$V_{DC} = \nu_{2,0} Y_{2,0}$$

= $\nu_{2,0} (2z^2 + f(x,y))$
= $\frac{1}{2} \psi_z z^2 + \tilde{f}(x,y)$,

from which we see that the axial curvature is $\psi_z = 4\nu_{2,0}$.

30. Page 131 Machine-learning is a good candidate for the rotation parameter optimization, but I do not claim it to be necessarily the best. Truth be told, I chose it partly out of an interest to see if neural networks could be applied in our experiment. There is somewhat of a hype surrounding neural networks, and their usefulness and applicability has certainly been demonstrated. It's a powerful tool since it requires very little knowledge from the user's side, both from the point of view of understanding what the network does and how it is optimized, and from the point of view of understanding the model that the network is trying to describe. The latter is incidentally one of the critiques that some have of neural networks: a well-trained network can be used effectively as a predictor of some output given some input, but it remains a black-box in the sense that it reveals very little about the model it's emulating. For example, a well-trained image-detection neural network might be able to tell a self-driving car that a person is crossing the street in front of it, but it's a scary thought that the people who designed it have no idea how it works. Likewise, with the neural network optimization for ion crystal rotations, I learn very little about why parameters require adjustment from the initially simulated values, though that may be less critical than in pedestrian-detection.

In the end, while perhaps not the most efficient method, the neural network-based optimization managed to improve a procedure that resulted in tens of induced phonons to less than one, which is something I couldn't manage manually. The method requires very little mathematical knowledge or intervention.

I also attempted to use this same machine-learning approach to optimize MS gates, to contrast it with the Bayesian optimization method used in [111]. The Bayesian method is a polar opposite of the neural network approach, since it demands a high-level of prior knowledge of the model describing the system. Also here, the neural network approach resulted in a high-fidelity gate as an output. Though optimization took notably more time than in [111], it can be argued that the method is more robust to model errors, and more aptly finds an optimum in the presence of noise

- 31. Page 131 There's probably a million good YouTube videos that describe how a neural network works, but a delightful one that can be enjoyed by physicists, and their grandmothers alike, is by Grant Sanderson from 3Blue1Brown [253]
- 32. Page 136 Is it fair to assume that the phonon state post-rotation is thermally distributed? We are, after all, not applying electric field noise, but a sequence of well-controlled fields. The measured excitation profiles match reasonably with the model that assumes a thermal distribution, even when considering the profiles of |11\), |01 + 10\), and |00\ separately. While that in itself is not a proof that it's a thermal distribution, it at least distills confidence that the calculated phonon number is a reasonable estimate. Other works [124, 178] go the extra mile to include coherent phonon displacement in the heating model.
- 33. Page 138 We took a data set that shows the Ramsey contrast for various initial ion separations, by changing the initial axial confinement. A rotation sequence then consists of first bringing ions to the confinement at which rotations were optimized, applying the rotation, and bringing the ions back to their initial separation. Sandwiched between two Ramsey pulses, one can see that the cyclical drop and recovery of Ramsey contrast, as $\phi_{pos}^{(1)}$ and $\phi_{pos}^{(2)}$ move in and out of phase with respect to each other.
- 34. Page 139 Why use only 100 rotations to show success rate of Ca Sr rotations? That is a classic case of saying on the day itself "I'll measure that properly later." By the time it was 'later' the setup was disassembled and traps were exchanged. Remeasuring

just Ca - Sr success rate would not fit in the narrative of the rest of the chapter since trap electrodes and filters are completely different.

- 35. Page 146 Whether or not one should omit the motional mode in the state notation has been a recurring discussion. I repeatedly left them out where I felt the motional mode wasn't important in the context of what is being explained, and could instead be distracting or confusing. However, notational correctness outweighs notational convenience, so I reluctantly reintroduced it in many places.
- 36. Page 155 Heating rates in our trap are notably reduced when adding external filters to the DC connections of the trap electrodes. Having low heating rates is quite crucial in the engineered resonance dissipation experiments. The fact that we require a low axial motional frequency does not simplify this. It is therefore important for the functionality of the sequence that the external filters are used. However, with a pretty aggressive frequency cutoff of about 12 Hz, it is not possible to apply ion reshuffling sequences on every cycle of an experiment, since they'd have to occur in a matter of at most tens of milliseconds.

Instead, we have a reordering sequence that lasts about half of a second. We could, in principle, run this sequence between each subsequent set of cycles. However, we have found that since the reorder sequences is not extremely gentle with ions (it involves pushing the Sr ion about 10 micrometers from the trap center), the ion lifetimes are noticeably reduced by performing this reorder so frequently.

We therefore elected to only run the sequence when necessary. In practice, this means monitoring fluorescence, excitation, and/or ion position (when using the CCD camera) to detect changes in crystal orientation and quickly clicking the "reorder" button whenever that happens. There'll be a few faulty data points, but can be removed/ignored. For the data presented in this section, if a reorder occurred during a scan, the whole scan was discarded and retaken.

- [1] H. Kragh. *Quantum Generations: A History of Physics in the Twentieth Century*. Princeton University Press, 2002.
- [2] E. Schrödinger. "Are There Quantum Jumps?" In: *British Journal for the Philosophy of Science* 3.11 (1952), pp. 233–242.
- [3] F. Scheck. "Applications of Quantum Mechanics." In: *Quantum Physics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 255–308.
- [4] J. I. Cirac. "Where we stand and what awaits us." In: Nanophotonics 10.1 (2021), pp. 453–456.
- [5] A. Sandberg. "The physics of information processing superobjects: daily life among the jupiter brains." In: *Journal of Evolution and Technology* 5.1 (1999).
- [6] R. P. Feynman. "Simulating physics with computers." In: *International Journal of Theoretical Physics* 21.6 (1982), pp. 467–488.
- [7] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer. "Elucidating reaction mechanisms on quantum computers." In: *Proceedings of the National Academy of Sciences* 114.29 (2017), pp. 7555–7560.
- [8] B. X. Wang, M. J. Tao, Q. Ai, T. Xin, N. Lambert, D. Ruan, Y. C. Cheng, F. Nori, F. G. Deng, and G. L. Long. "Efficient quantum simulation of photosynthetic light harvesting." In: *npj Quantum Information* 4.1 (2018), p. 52.
- [9] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien. "Quantum computers." In: *Nature* 464.7285 (2010), pp. 45–53.
- [10] P. W. Shor. "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer." In: SIAM Journal on Computing 26.5 (1997), pp. 1484–1509.
- [11] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2011.
- [12] T. Monz, D. Nigg, E. A. Martinez, M. F. Brandl, P. Schindler, R. Rines, S. X. Wang, I. L. Chuang, and R. Blatt. "Realization of a scalable Shor algorithm." In: *Science* 351.6277 (2016), pp. 1068–1070.
- [13] E. Martín-López, 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." In: *Nature Photonics* 6.11 (2012), pp. 773–776.
- [14] U. Skosana and M. Tame. "Demonstration of Shor's factoring algorithm for N=21 on IBM quantum processors." In: *Scientific Reports* 11.1 (2021), p. 16599.
- [15] M. H. Devoret and R. J. Schoelkopf. "Superconducting circuits for quantum information: An outlook." In: *Science* 339.6124 (2013), pp. 1169–1174.
- [16] J. W. Silverstone, D. Bonneau, J. L. O'Brien, and M. G. Thompson. "Silicon quantum photonics." In: *IEEE Journal of Selected Topics in Quantum Electronics* 22.6 (2016), pp. 390–402.
- [17] D. S. Weiss and M. Saffman. "Quantum computing with neutral atoms." In: *Physics Today* 70.7 (2017), pp. 44–50.
- [18] P. C. Maurer, G Kucsko, C Latta, L Jiang, N. Y. Yao, S. D. Bennett, F Pastawski, D Hunger, N Chisholm, M Markham, D. J. Twitchen, J. I. Cirac, and M. D. Lukin. "Room-temperature quantum bit memory exceeding one second." In: *Science* 336.6086 (2012), pp. 1283–1286.
- [19] P. L. Knight, E. A. Hinds, M. B. Plenio, D. J. Wineland, M Barrett, J Britton, J Chiaverini, B DeMarco, W. M. Itano, B Jelenković, C Langer, D Leibfried, V Meyer, T Rosenband, and T Schätz. "Quantum information processing with trapped ions." In: *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 361.1808 (2003), pp. 1349–1361.

- [20] J. Weiner and P. T. Ho. Light-Matter Interaction, Fundamentals and Applications. Wiley, 2008.
- [21] G Werth, V. N. Gheorghe, and F. G. Major. *Charged Particle Traps II: Applications*. Springer Berlin, Heidelberg, 2009.
- [22] W. Paul and H. Steinwedel. "Ein neues Massenspektrometer ohne Magnetfeld." In: Zeitschrift für Naturforschung A 8.7 (1953), pp. 448–450.
- [23] J. D. Siverns and Q. Quraishi. "Ion trap architectures and new directions." In: *Quantum Information Processing* 16.12 (2017), p. 314.
- [24] D. Kielpinski, C. Monroe, and D. J. Wineland. "Architecture for a large-scale ion-trap quantum computer." In: *Nature* 417.6890 (2002), pp. 709–711.
- [25] K. R. Brown, J. Kim, and C. Monroe. "Co-designing a scalable quantum computer with trapped atomic ions." In: *npj Quantum Information* 2.1 (2016), p. 16034.
- [26] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. "High-fidelity universal gate set for ⁹Be⁺ ion qubits." In: *Physical Review Letters* 117.6 (2016), p. 060505.
- [27] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. "Demonstration of the trapped-ion quantum CCD computer architecture." In: *Nature* 592.7853 (2021), pp. 209–213.
- [28] J. Preskill. "Quantum computing in the NISQ era and beyond." In: Quantum 2 (2018), p. 79.
- [29] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. "Ion-trap measurements of electric-field noise near surfaces." In: *Reviews of Modern Physics* 87.4 (2015), p. 1419.
- [30] P. C. Holz, K. Lakhmanskiy, D. Rathje, P. Schindler, Y. Colombe, and R. Blatt. "Electric-field noise in a high-temperature superconducting surface ion trap." In: *Phys. Rev. B* 104.6 (2021), p. 064513.
- [31] M. Teller, D. A. Fioretto, P. C. Holz, P. Schindler, V. Messerer, K. Schüppert, Y. Zou, R. Blatt, J. Chiaverini, J. Sage, and T. E. Northup. "Heating of a trapped ion induced by dielectric materials." In: *Phys. Rev. Lett.* 126.23 (2021), p. 230505.
- [32] 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." In: *New Journal of Physics* 13.12 (2011), p. 123023.
- [33] D. A. Hite, Y. Colombe, A. C. Wilson, D. T. C. Allcock, D. Leibfried, D. J. Wineland, and D. P. Pappas. "Surface science for improved ion traps." In: *MRS Bulletin* 38.10 (2013), pp. 826–833.
- [34] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. "Scaling and suppression of anomalous heating in ion traps." In: *Phys. Rev. Lett.* 97.10 (2006), p. 103007.
- [35] D. A. Hite, K. S. McKay, and D. P. Pappas. "Surface science motivated by heating of trapped ions from the quantum ground state." In: *New Journal of Physics* 23.10 (2021), p. 103028.
- [36] R. McConnell, C. Bruzewicz, J. Chiaverini, and J. Sage. "Reduction of trapped-ion anomalous heating by in situ surface plasma cleaning." In: *Phys. Rev. A* 92.2 (2015), p. 020302.
- [37] E. Kim, A. Safavi-Naini, D. A. Hite, K. S. McKay, D. P. Pappas, P. F. Weck, and H. R. Sadeghpour. "Electric-field noise from carbon-adatom diffusion on a Au(110) surface: First-principles calculations and experiments." In: *Phys. Rev. A* 95.3 (2017), p. 033407.
- [38] S. X. Wang, Y. Ge, J. Labaziewicz, E. Dauler, K. Berggren, and I. L. Chuang. "Superconducting microfabricated ion traps." In: *Applied Physics Letters* 97.24 (2010), p. 244102.
- [39] Z. C. Mao, Y. Z. Xu, Q. X. Mei, W. D. Zhao, Y. Jiang, Y. Wang, X. Y. Chang, L. He, L. Yao, Z. C. Zhou, Y. K. Wu, and L. M. Duan. "Experimental realization of multi-ion sympathetic cooling on a trapped ion crystal." In: *Phys. Rev. Lett.* 127.14 (2021), p. 143201.
- [40] H. Rohde, S. T. Gulde, C. F. Roos, P. A. Barton, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt. "Sympathetic ground-state cooling and coherent manipulation with two-ion crystals." In: *Journal of Optics B: Quantum and Semiclassical Optics* 3.1 (2001), S34.

- [41] 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 ⁹Be⁺ and ²⁴Mg⁺ for quantum logic." In: *Phys. Rev. A* 68.4 (2003), p. 042302.
- [42] J. P. Home, M. J. McDonnell, D. J. Szwer, B. C. Keitch, D. M. Lucas, D. N. Stacey, and A. M. Steane. "Memory coherence of a sympathetically cooled trapped-ion qubit." In: *Phys. Rev. A* 79.5 (2009), p. 50305.
- [43] J. Labaziewicz, Y. Ge, D. R. Leibrandt, S. X. Wang, R. Shewmon, and I. L. Chuang. "Temperature dependence of electric field noise above gold surfaces." In: *Physical Review Letters* 101.18 (2008), p. 180602.
- [44] J Ekin. Experimental Techniques for Low-Temperature Measurements: Cryostat Design, Material Properties and Superconductor Critical-Current Testing. Oxford University Press, 2006.
- [45] A. Steane. "Multiple-particle interference and quantum error correction." In: *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 452.1954 (1996), pp. 2551–2577.
- [46] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill. "Topological quantum memory." In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4452–4505.
- [47] H. Bombin and M. A. Martin-Delgado. "Topological quantum distillation." In: *Phys. Rev. Lett.* 97.18 (2006), p. 180501.
- [48] H. Bombin. "Single-shot fault-tolerant quantum error correction." In: *Phys. Rev. X* 5.3 (2015), p. 031043.
- [49] D. G. Cory, M. D. Price, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, and S. S. Somaroo. "Experimental quantum error correction." In: *Phys. Rev. Lett.* 81.10 (1998), pp. 2152–2155.
- [50] B. M. Terhal. "Quantum error correction for quantum memories." In: *Rev. Mod. Phys.* 87.2 (2015), pp. 307–346.
- [51] E. T. Campbell, B. M. Terhal, and C. Vuillot. "Roads towards fault-tolerant universal quantum computation." In: *Nature* 549.7671 (2017), pp. 172–179.
- [52] D. Jordan and P. Smith. *Nonlinear Ordinary Differential Equations: An Introduction for Scientists and Engineers*. Oxford University Press, 2007.
- [53] N. W. McLachlan. Theory and Application of Mathieu Functions. Clarendon Press, 1947.
- [54] P. K. Ghosh. *Ion Traps*. Clarendon Press, 1995.
- [55] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. "Quantum dynamics of single trapped ions." In: *Reviews of Modern Physics* 75.1 (2003), pp. 281–324.
- [56] H. G. Dehmelt. "Radiofrequency Spectroscopy of Stored Ions I: Storage." In: vol. 3. Advances in Atomic and Molecular Physics. Academic Press, 1968, pp. 53–72.
- [57] C Monroe and J Kim. "Scaling the ion trap quantum processor." In: *Science* 339.6124 (2013), pp. 1164–1169.
- [58] S. A. Schulz, U. Poschinger, F. Ziesel, and F. Schmidt-Kaler. "Sideband cooling and coherent dynamics in a microchip multi-segmented ion trap." In: *New Journal of Physics* 10.4 (2008), p. 045007.
- [59] S. Ragg, C. Decaroli, T. Lutz, and J. P. Home. "Segmented ion-trap fabrication using high precision stacked wafers." In: *Review of Scientific Instruments* 90.10 (2019), p. 103203.
- [60] C. Decaroli, R. Matt, R. Oswald, M. Ernzer, J. Flannery, S. Ragg, and J. Home. *Design, fabrication and characterisation of a micro-fabricated double-junction segmented ion trap*.
- [61] Y. Wan, R. Jördens, S. D. Erickson, J. J. Wu, R. Bowler, T. R. Tan, P. Y. Hou, D. J. Wineland, A. C. Wilson, and D. Leibfried. "Ion transport and reordering in a 2D trap array." In: *Advanced Quantum Technologies* 3.11 (2020), p. 2000028.

- [62] R. B. Blakestad, C. Ospelkaus, A. P. Vandevender, J. H. Wesenberg, M. J. Biercuk, D. Leibfried, and D. J. Wineland. "Near-ground-state transport of trapped-ion qubits through a multidimensional array." In: *Physical Review A Atomic, Molecular, and Optical Physics* 84.3 (2011), p. 032314.
- [63] J. P. Home. "Quantum science and metrology with mixed-species ion chains." In: *Advances in Atomic, Molecular and Optical Physics* 62 (2013), pp. 231–277.
- [64] D. F. James. "Quantum dynamics of cold trapped ions with application to quantum computation." In: *Applied Physics B: Lasers and Optics* 66.2 (1998), pp. 181–190.
- [65] J. B. Wübbena, S. Amairi, O. Mandel, and P. O. Schmidt. "Sympathetic cooling of mixed-species two-ion crystals for precision spectroscopy." In: *Physical Review A Atomic, Molecular, and Optical Physics* 85.4 (2012), p. 043412.
- [66] J. P. Home, D. Hanneke, J. D. Jost, D. Leibfried, and D. J. Wineland. "Normal modes of trapped ions in the presence of anharmonic trap potentials." In: *New Journal of Physics* 13 (2011), p. 073026.
- [67] J. I. Cirac and P Zoller. "Quantum computations with cold trapped ions." In: *Phys. Rev. Lett.* 74.20 (1995), pp. 4091–4094.
- [68] D. P. DiVincenzo. "The physical implementation of quantum computation." In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783.
- [69] C. Roos. "Thesis: Controlling the quantum state of trapped ions." PhD thesis. University of Innsbruck, 2000, p. 145.
- [70] N. B. Delone and V. P. Krainov. "AC Stark shift of atomic energy levels." In: *Physics-Uspekhi* 42.7 (1999), p. 669.
- [71] C. Hempel. "Digital quantum simulation, Schrödinger cat state spectroscopy and setting up a linear ion trap." PhD thesis. University of Innsbruck, 2014.
- [72] D. Manzano. "A short introduction to the Lindblad master equation." In: *AIP Advances* 10.2 (2020), p. 025106.
- [73] E. T. Jaynes and F. W. Cummings. "Comparison of quantum and semiclassical radiation theories with application to the beam maser." In: *Proceedings of the IEEE* 51.1 (1963), pp. 89–109.
- [74] G. Tommaseo, T. Pfeil, G. Revalde, G. Werth, P. Indelicato, and J. P. Desclaux. "The g_J -factor in the ground state of Ca⁺." In: *The European Physical Journal D Atomic, Molecular, Optical and Plasma Physics* 25.2 (2003), pp. 113–121.
- [75] M. Chwalla, J. Benhelm, K. Kim, G. Kirchmair, T. Monz, M. Riebe, P. Schindler, A. S. Villar, W. Hänsel, C. F. Roos, R. Blatt, M. Abgrall, G. Santarelli, G. D. Rovera, and P. Laurent. "Absolute frequency measurement of the 40 Ca $^+$ $4s^2S_{1/2} 3d^2D_{5/2}$ clock transition." In: *Phys. Rev. Lett.* 102.2 (2009), p. 023002.
- [76] J. Bergquist, J. Bollinger, W. Itano, and D. Wineland. *Trapped Ions and Laser Cooling, VI.* URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=105136.
- [77] C. J. Foot. Atomic Physics. Oxford University Press, 2005.
- [78] K. Kim, C. F. Roos, L. Aolita, H. Häffner, V. Nebendahl, and R. Blatt. "Geometric phase gate on an optical transition for ion trap quantum computation." In: *Phys. Rev. A* 77.5 (2008), p. 050303.
- [79] K. Mølmer and A. Sørensen. "Multiparticle entanglement of hot trapped ions." In: *Phys. Rev. Lett.* 82.9 (1999), pp. 1835–1838.
- [80] V. Nebendahl, H. Haeffner, and C. Roos. "Optimal control of entangling operations for trapped ion quantum computing." In: *Physical Review A* 79 (2008), p. 012312.
- [81] H. J. Metcalf and P. van der Straten. "Laser Cooling and Trapping." In: *The Optics Encyclopedia*. American Cancer Society, 2007.
- [82] S. Stenholm. "The semiclassical theory of laser cooling." In: *Rev. Mod. Phys.* 58.3 (1986), pp. 699–739.

- [83] E. W. Weber. "Optical pumping of ions." In: Physics Reports 32.3 (1977), pp. 123–167.
- [84] G. Morigi, J. Eschner, and C. H. Keitel. "Ground state laser cooling using electromagnetically induced transparency." In: *Phys. Rev. Lett.* 85.21 (2000), pp. 4458–4461.
- [85] M. K. Joshi, A. Fabre, C. Maier, T. Brydges, D. Kiesenhofer, H. Hainzer, R. Blatt, and C. F. Roos. "Polarization-gradient cooling of 1D and 2D ion Coulomb crystals." In: *New Journal of Physics* 22.10 (2020), p. 103013.
- [86] D. C. Marinescu and G. M. Marinescu. "Chapter 1 Preliminaries." In: *Classical and Quantum Information*. Boston: Academic Press, 2012, pp. 1–131.
- [87] G. B. Arfken, H. J. Weber, and F. E. Harris. "Chapter 12 Further Topics in Analysis." In: *Mathematical Methods for Physicists (Seventh Edition)*. Seventh Ed. Boston: Academic Press, 2013, pp. 551–598.
- [88] M. W. van Mourik, E. A. Martinez, L. Gerster, P. Hrmo, T. Monz, P. Schindler, and R. Blatt. "Coherent rotations of qubits within a surface ion-trap quantum computer." In: *Physical Review A* 102.2 (2020), p. 022611.
- [89] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Society for Industrial and Applied Mathematics, 1995.
- [90] J. Keller, H. L. Partner, T. Burgermeister, and T. E. Mehlstäubler. "Precise determination of micromotion for trapped-ion optical clocks." In: *Journal of Applied Physics* 118.10 (2015), p. 104501.
- [91] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. "Minimization of ion micromotion in a Paul trap." In: *Journal of Applied Physics* 83.10 (1998), pp. 5025–5033.
- [92] R. Dubessy, T. Coudreau, and L. Guidoni. "Electric field noise above surfaces: A model for heating-rate scaling law in ion traps." In: *Phys. Rev. A* 80.3 (2009), p. 031402.
- [93] P. Dutta and P. M. Horn. "Low-frequency fluctuations in solids: 1/*f* noise." In: *Rev. Mod. Phys.* 53.3 (1981), pp. 497–516.
- [94] 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." In: *Physical Review A Atomic, Molecular, and Optical Physics* 61.6 (2000), p. 063418.
- [95] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, S. Schulz, F. Schmidt-Kaler, and H. Häffner. "Fabrication and heating rate study of microscopic surface electrode ion traps." In: *New Journal of Physics* 13.1 (2011), p. 013032.
- [96] D. Lidar, I. Chuang, and K. Whaley. "Decoherence free subspaces for quantum computation." In: *Physical Review Letters* 81 (1998), p. 2594.
- [97] H. Y. Carr and E. M. Purcell. "Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments." In: *Phys. Rev.* 94.3 (1954), pp. 630–638.
- [98] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger. "Experimental Uhrig dynamical decoupling using trapped ions." In: *Phys. Rev. A* 79.6 (2009), p. 062324.
- [99] G. J. Milburn, S. Schneider, and D. F. V. James. "Ion trap quantum computing with warm ions." In: *Fortschritte der Physik* 48.9-11 (2000), pp. 801–810.
- [100] P. J. Lee, K. A. Brickman, L. Deslauriers, P. C. Haljan, L. M. Duan, and C. Monroe. "Phase control of trapped ion quantum gates." In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (2005), S371–S383.
- [101] 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." In: *Physical review letters* 91.15 (2003), p. 157901.

- [102] C. H. Baldwin, B. J. Bjork, M. Foss-Feig, J. P. Gaebler, D. Hayes, M. G. Kokish, C. Langer, J. A. Sedlacek, D. Stack, and G. Vittorini. "High-fidelity light-shift gate for clock-state qubits." In: *Phys. Rev. A* 103.1 (2021), p. 012603.
- [103] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L. M. Duan, and C. Monroe. "Entanglement of single-atom quantum bits at a distance." In: *Nature* 449.7158 (2007), pp. 68–71.
- [104] C. Cabrillo, J. Cirac, P. Garcia-Fernandez, and P. Zoller. "Creation of entangled states of distant atoms by interference." In: *Physical Review A* 59 (1998), pp. 1025–1033.
- [105] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. "Deterministic entanglement of ions in thermal states of motion." In: *New Journal of Physics* 11.2 (2009), p. 023002.
- [106] M. F. Brandl et al. "Cryogenic setup for trapped ion quantum computing." In: *Review of Scientific Instruments* 87.11 (2016), p. 113103.
- [107] M. F. Brandl. "Towards Cryogenic Scalable Quantum Computing with Trapped Ions." PhD thesis. University of Innsbruck, 2016.
- [108] M. Niedermayr. "Cryogenic surface ion traps." PhD thesis. University of Innsbruck, 2015.
- [109] P. Maunz. "High Optical Access Trap 2.0." In: *Technical Report SAND2016-0796R* (2016), p. 618951.
- [110] Quantenbit AG Schmidt-Kaler. URL: https://www.quantenbit.physik.uni-mainz.de/.
- [111] L. Gerster, F. Martínez-García, P. Hrmo, M. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz. "Experimental Bayesian calibration of trapped ion entangling operations." In: (2021), p. 020350.
- [112] P. Hrmo, B. Wilhelm, L. Gerster, M. W. van Mourik, M. Huber, R. Blatt, P. Schindler, T. Monz, and M. Ringbauer. "Native qudit entanglement in a trapped ion quantum processor." In: *Nature Communications* 14.1 (2023), p. 2242.
- [113] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. "Measurement of ion motional heating rates over a range of trap frequencies and temperatures." In: *Physical Review A Atomic, Molecular, and Optical Physics* 91.4 (2015), 041402(R).
- [114] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. A. Trainor, and W. Kells. "Thousandfold improvement in the measured antiproton mass." In: *Phys. Rev. Lett.* 65.11 (1990), pp. 1317–1320.
- [115] F. Jessen et al. "Trapping of ultracold atoms in a 3He/4He dilution refrigerator." In: *Applied Physics B* 116.3 (2014), pp. 665–671.
- [116] D. De Motte, A. R. Grounds, M. Rehák, A. Rodriguez Blanco, B. Lekitsch, G. S. Giri, P. Neilinger, G. Oelsner, E. Il'ichev, M. Grajcar, and W. K. Hensinger. "Experimental system design for the integration of trapped-ion and superconducting qubit systems." In: *Quantum Information Processing* 15.12 (2016), pp. 5385–5414.
- [117] R. G. Scurlock. History and Origins of Cryogenics. Clarendon Press, 1992.
- [118] P. Micke, J. Stark, S. A. King, T. Leopold, T. Pfeifer, L. Schmöger, M. Schwarz, L. J. Spieß, P. O. Schmidt, and J. R. Crespo López-Urrutia. "Closed-cycle, low-vibration 4 K cryostat for ion traps and other applications." In: *Review of Scientific Instruments* 90.6 (2019), p. 065104.
- [119] R. Haefer. Kryo-Vakuumtechnik Grundlagen und Anwendungen. 1st ed. Springer Berlin Heidelberg, 1981.
- [120] J. G. Hust. "Thermal anchoring of wires in cryogenic apparatus." In: Review of Scientific Instruments 41.5 (1970), pp. 622–624.
- [121] L. Gerster. "Towards scalability of quantum processors." PhD thesis. University of Innsbruck, 2023.

- [122] Z. Chen, W. Li, R. Li, Y. Zhang, G. Xu, and H. Cheng. "Fabrication of highly transparent and conductive indium–tin oxide thin films with a high figure of merit via solution processing." In: *Langmuir* 29.45 (2013), pp. 13836–13842.
- [123] F. R. Ong, K. Schüppert, P. Jobez, M. Teller, B. Ames, D. A. Fioretto, K. Friebe, M. Lee, Y. Colombe, R. Blatt, and T. E. Northup. "Probing surface charge densities on optical fibers with a trapped ion." In: *New Journal of Physics* 22.6 (2020), p. 063018.
- [124] V. Kaushal, B. Lekitsch, A. Stahl, J. Hilder, D. Pijn, C. Schmiegelow, A. Bermudez, M. Müller, F. Schmidt-Kaler, and U. Poschinger. "Shuttling-based trapped-ion quantum information processing." In: *AVS Quantum Science* 2.1 (2020), p. 014101.
- [125] M. F. Brandl, P. Schindler, T. Monz, and R. Blatt. "Cryogenic resonator design for trapped ion experiments in Paul traps." In: *Applied Physics B: Lasers and Optics* 122.6 (2016), p. 157.
- [126] R. E. Honig. *Vapor pressure data for the more common elements*. David Sarnoff Research Center, 1957.
- [127] K. Vant, J. Chiaverini, W. Lybarger, and D. J. Berkeland. "Photoionization of strontium for trapped-ion quantum information processing." In: *arXiv preprint quant-ph/0607055* (2006).
- [128] L. Postler. "Ein schmalbandiges Halbleiter-Lasersystem für Quanteninformationsverarbeitung und Spektroskopie." Master's thesis. University of Innsbruck, 2015.
- [129] R. Stricker. "Gatteroperationen höher güte in einem optischen quantenbit." Master's thesis. University of Innsbruck, 2017.
- [130] M. Guggemos, D. Heinrich, O. A. Herrera-Sancho, R. Blatt, and C. F. Roos. "Sympathetic cooling and detection of a hot trapped ion by a cold one." In: *New Journal of Physics* 17.10 (2015), p. 103001.
- [131] L. S. Ma, P. Jungner, J. Ye, and J. Hall. "Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path." In: *Optics Letters* 19 (1994), p. 1777.
- [132] D. Heinrich. "Ultrafast coherent excitatation of a 40 Ca $^+$ ion." PhD Thesis. University of Innbsruck, 2020.
- [133] P. T. T. Pham. "A general-purpose pulse sequencer for quantum computing." Master's thesis. Massachusetts Institute of Technology, 2005.
- [134] P. Schindler. "Frequency synthesis and pulse shaping for quantum information processing with trapped ions." Master's Thesis. 2008.
- [135] J. Chiaverini and J. M. Sage. "Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range." In: *Phys. Rev. A* 89.1 (2014), p. 012318.
- [136] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt. "Characterizing large-scale quantum computers via cycle benchmarking." In: *Nature Communications* 10.1 (2019), p. 5347.
- [137] Ph.D. theses Quantum Optics and Spectroscopy. URL: https://quantumoptics.at/en/publications/ph-d-theses.html.
- [138] R. Blümel, C. Kappler, W. Quint, and H. Walther. "Chaos and order of laser-cooled ions in a Paul trap." In: *Physical Review A* 40.2 (1989), pp. 808–823.
- [139] M. W. Van Mourik, P. Hrmo, L. Gerster, B. Wilhelm, R. Blatt, P. Schindler, and T. Monz. "Rf-induced heating dynamics of noncrystallized trapped ions." In: *Physical Review A* 105.3 (2022), p. 033101.
- [140] A. M. Hankin, E. R. Clements, Y. Huang, S. M. Brewer, J. S. Chen, C. W. Chou, D. B. Hume, and D. R. Leibrandt. "Systematic uncertainty due to background-gas collisions in trapped-ion optical clocks." In: *Physical Review A* 100.3 (2019), p. 033419.
- [141] R. Blümel, J. M. Chen, E. Peik, W. Quint, W. Schleich, Y. R. Shen, and H. Walther. "Phase transitions of stored laser-cooled ions." In: *Nature* 334.6180 (1988), pp. 309–313.

- [142] M. Cetina, A. T. Grier, and V. Vuletić. "Micromotion-induced limit to atom-ion sympathetic cooling in Paul traps." In: *Physical Review Letters* 109.25 (2012), p. 253201.
- [143] L. H. Nguyên, A. Kalev, M. D. Barrett, and B. G. Englert. "Micromotion in trapped atom-ion systems." In: *Physical Review A Atomic, Molecular, and Optical Physics* 85.5 (2012), p. 052718.
- [144] I. Siemers, R. Blatt, T. Sauter, and W. Neuhauser. "Dynamics of ion clouds in Paul traps." In: *Phys. Rev. A* 38.10 (1988), pp. 5121–5128.
- [145] C. Champenois. "About the dynamics and thermodynamics of trapped ions." In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 42.15 (2009), p. 154002.
- [146] Y. S. Nam, D. K. Weiss, and R. Blümel. "Explicit, analytical radio-frequency heating formulas for spherically symmetric nonneutral plasmas in a Paul trap." In: *Physics Letters, Section A: General, Atomic and Solid State Physics* 381.40 (2017), pp. 3477–3481.
- [147] V. L. Ryjkov, X. Zhao, and H. A. Schuessler. "Simulations of the rf heating rates in a linear quadrupole ion trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 71.3 (2005), p. 033414.
- [148] J. D. Tarnas, Y. S. Nam, and R. Blümel. "Universal heating curve of damped Coulomb plasmas in a Paul trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 88.4 (2013), 041401(R).
- [149] C. B. Zhang, D. Offenberg, B. Roth, M. A. Wilson, and S. Schiller. "Molecular-dynamics simulations of cold single-species and multispecies ion ensembles in a linear Paul trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 76.1 (2007), p. 012719.
- [150] R. G. Brewer, J. Hoffnagle, R. G. DeVoe, L. Reyna, and W. Henshaw. "Collision-induced two-ion chaos." In: *Nature* 344.6264 (1990), pp. 305–309.
- [151] J. W. Emmert, M. Moore, and R. Blümel. "Prediction of a deterministic melting transition of two-ion crystals in a Paul trap." In: *Physical Review A* 48.3 (1993), pp. 1757–1760.
- [152] F. Diedrich, E. Peik, J. M. Chen, W. Quint, and H. Walther. "Observation of a Phase Transition of Stored Laser-Cooled Ions." In: *Physical Review Letters* 59.26 (1987), pp. 2931–2934.
- [153] D. H. Dubin and T. M. O'Neil. "Trapped nonneutral plasmas, liquids, and crystals (the thermal equilibrium states)." In: *Reviews of Modern Physics* 71.1 (1999), pp. 87–172.
- [154] J. P. Schiffer, M. Drewsen, J. S. Hangst, and L. Hornekær. "Temperature, ordering, and equilibrium with time-dependent confining forces." In: *Proceedings of the National Academy of Sciences of the United States of America* 97.20 (2000), pp. 10697–10700.
- [155] Ion trajectory simulation source code. URL: https://github.com/MartinvMourik/TrappedIonDynamics/tree/v1.0.
- [156] M. G. House. "Analytic model for electrostatic fields in surface-electrode ion traps." In: *Physical Review A Atomic, Molecular, and Optical Physics* 78.3 (2008), p. 033402.
- [157] J. D. Prestage, A. Williams, L. Maleki, M. J. Djomehri, and E. Harabetian. "Dynamics of charged particles in a Paul radio-frequency quadrupole trap." In: *Physical Review Letters* 66.23 (1991), pp. 2964–2967.
- [158] V. Roberdel, D. Leibfried, D. Ullmo, and H. Landa. "Phase-space study of surface-electrode Paul traps: Integrable, chaotic, and mixed motions." In: *Physical Review A* 97.5 (2018), p. 053419.
- [159] D. Freedman, R. Pisani, and R. Purves. Statistics. WW Norton & Company New York, 2007.
- [160] R. Lyons. Sum of Two Sinusoids. URL: https://dspguru.com/files/Sum_of_Two_Sinusoids.pdf.
- [161] L. Devroye. Non-Uniform Random Variate Generation. Springer-Verlag, 1986.
- [162] R. Blatt, P. Zoller, G. Holzmüller, and I. Siemers. "Brownian motion of a parametric oscillator: A model for ion confinement in radio frequency traps." In: *Zeitschrift für Physik D Atoms, Molecules and Clusters* 4.2 (1986), pp. 121–126.
- [163] D. J. Wineland and W. M. Itano. "Laser cooling of atoms." In: *Physical Review A* 20.4 (1979), pp. 1521–1540.

- [164] L. L. Yan, S. L. Su, and M. Feng. "Analytical investigation of one-dimensional Doppler cooling of trapped ions with Λ -type configuration." In: *Physical Review A* 100.3 (2019), p. 033418.
- [165] M. Hettrich, T. Ruster, H. Kaufmann, C. F. Roos, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. "Measurement of dipole matrix elements with a single trapped ion." In: *Physical Review Letters* 115.14 (2015), p. 143003.
- [166] F. Vedel, M. Vedel, and R. Evans March. "New Schemes for resonant ejection in r.f. quadrupolar ion traps." In: *International Journal of Mass Spectrometry and Ion Processes* 99.1-2 (1990), pp. 125–138.
- [167] D. T. Snyder, W. P. Peng, and R. G. Cooks. "Resonance methods in quadrupole ion traps." In: *Chemical Physics Letters* 668 (2017), pp. 69–89.
- [168] B. A. Collings, W. R. Stott, and F. A. Londry. "Resonant excitation in a low-pressure linear ion trap." In: *Journal of the American Society for Mass Spectrometry* 14.6 (2003), pp. 622–634.
- [169] Y. Ibaraki, U. Tanaka, and S. Urabe. "Detection of parametric resonance of trapped ions for micromotion compensation." In: *Applied Physics B: Lasers and Optics* 105.2 (2011), pp. 219–223.
- [170] X. Zhao, V. L. Ryjkov, and H. A. Schuessler. "Parametric excitations of trapped ions in a linear rf ion trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 66.6 (2002), p. 063414.
- [171] M. Sudakov, N. Konenkov, D. J. Douglas, and T. Glebova. "Excitation frequencies of ions confined in a quadrupole field with quadrupole excitation." In: *Journal of the American Society for Mass Spectrometry* 11.1 (2000), pp. 10–18.
- [172] X. Zhou, C. Xiong, S. Zhang, N. Zhang, and Z. Nie. "Study of nonlinear resonance effect in paul trap." In: *Journal of the American Society for Mass Spectrometry* 24.5 (2013), pp. 794–800.
- [173] J. Franzen. "The non-linear ion trap. Part 5. Nature of non-linear resonances and resonant ion ejection." In: *International Journal of Mass Spectrometry and Ion Processes* 130.1-2 (1994), pp. 15–40.
- [174] R. Alheit, S. Kleineidam, F. Vedel, M. Vedel, and G. Werth. "Higher order non-linear resonances in a Paul trap." In: *International Journal of Mass Spectrometry and Ion Processes* 154.3 (1996), pp. 155–169.
- [175] C. D. Bruzewicz, R. McConnell, J. Stuart, J. M. Sage, and J. Chiaverini. "Dual-species, multiqubit logic primitives for Ca+/Sr+ trapped-ion crystals." In: *npj Quantum Information* 5.1 (2019), p. 102.
- [176] C. C. Bradley, J. G. Story, J. J. Tollett, J. Chen, N. W. M. Ritchie, and R. G. Hulet. "Laser cooling of lithium using relay chirp cooling." In: *Optics Letters* 17.5 (1992), p. 349.
- [177] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. "Coherent diabatic ion transport and separation in a multizone trap array." In: *Physical Review Letters* 109.8 (2012), p. 080502.
- [178] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. "Controlling fast transport of cold trapped ions." In: *Physical Review Letters* 109.8 (2012), p. 080501.
- [179] M. Palmero, R. Bowler, J. P. Gaebler, D. Leibfried, and J. G. Muga. "Fast transport of mixed-species ion chains within a Paul trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 90.5 (2014), p. 053408.
- [180] M. A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. "Transport of quantum states and separation of ions in a dual RF ion trap." In: *Quantum Information and Computation* 2.4 (2002), pp. 257–271.
- [181] X. J. Lu, J. G. Muga, X. Chen, U. G. Poschinger, F. Schmidt-Kaler, and A. Ruschhaupt. "Fast shuttling of a trapped ion in the presence of noise." In: *Physical Review A Atomic, Molecular, and Optical Physics* 89.6 (2014), p. 063414.
- [182] A. Pauli. "Classical control of an ion in a surface trap." Master's thesis. University of Innsbruck, 2011.

- [183] H. Kaufmann, T. Ruster, C. Schmiegelow, M. Luda, K. Vidyut, J. Schulz, D. Lindenfels, F. Schmidt-Kaler, and U. Poschinger. "Scalable creation of long-lived multipartite entanglement." In: *Physical Review Letters* 119 (2017), p. 150503.
- [184] S. D. Fallek, C. D. Herold, B. J. McMahon, K. M. Maller, K. R. Brown, and J. M. Amini. "Transport implementation of the Bernstein–Vazirani algorithm with ion qubits." In: *New Journal of Physics* 18.8 (2016), p. 083030.
- [185] D. Hanneke, J. P. Home, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. "Realization of a programmable two-qubit quantum processor." In: *Nature Physics* 6.1 (2010), pp. 13–16.
- [186] R. Reichle, D. Leibfried, R. B. Blakestad, J. Britton, J. D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. "Transport dynamics of single ions in segmented microstructured Paul trap arrays." In: *Fortschritte der Physik* 54.8-10 (2006), pp. 666–685.
- [187] H. A. Fürst, M. H. Goerz, U. G. Poschinger, M. Murphy, S. Montangero, T. Calarco, F. Schmidt-Kaler, K. Singer, and C. P. Koch. "Controlling the transport of an ion: classical and quantum mechanical solutions." In: *New Journal of Physics* 16.7 (2014), p. 075007.
- [188] J. Alonso, F. M. Leupold, B. C. Keitch, and J. P. Home. "Quantum control of the motional states of trapped ions through fast switching of trapping potentials." In: *New Journal of Physics* 15 (2013), p. 023001.
- [189] Y. Ding, T. Y. Huang, K. Paul, M. Hao, X. Chen, and X. Chen. "Smooth bang-bang shortcuts to adiabaticity for atomic transport in a moving harmonic trap." In: *Physical Review A* 101.6 (2020), p. 063410.
- [190] W. C. Burton, B. Estey, I. M. Hoffman, A. R. Perry, C. Volin, and G. Price. "Transport of multispecies ion crystals through a junction in an RF Paul trap." In: *ArXiv* (2022).
- [191] P. C. Holz. "Towards ion-lattice quantum processors with surface trap arrays." PhD thesis. University of Innsbruck, 2019.
- [192] A. C. Wilson, Y. Colombe, K. R. Brown, E. Knill, D. Leibfried, and D. J. Wineland. "Tunable spin-spin interactions and entanglement of ions in separate potential wells." In: *Nature* 512.1 (2014), pp. 57–60.
- [193] F. Hakelberg, P. Kiefer, M. Wittemer, U. Warring, and T. Schaetz. "Interference in a prototype of a two-dimensional ion trap array quantum simulator." In: *Phys. Rev. Lett.* 123.10 (2019), p. 100504.
- [194] J. P. Home and A. M. Steane. "Electrode configurations for fast separation of trapped ions." In: *Quantum Information and Computation* 6.4-5 (2006), pp. 289–325.
- [195] M. Palmero, S. Martínez-Garaot, U. G. Poschinger, A. Ruschhaupt, and J. G. Muga. "Fast separation of two trapped ions." In: *New Journal of Physics* 17.9 (2015), p. 093031.
- [196] H. Kaufmann, T. Ruster, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. "Dynamics and control of fast ion crystal splitting in segmented Paul traps." In: *New Journal of Physics* 16 (2014), p. 073012.
- [197] T. Ruster, C. Warschburger, H. Kaufmann, C. T. Schmiegelow, A. Walther, M. Hettrich, A. Pfister, V. Kaushal, F. Schmidt-Kaler, and U. G. Poschinger. "Experimental realization of fast ion separation in segmented Paul traps." In: *Physical Review A Atomic, Molecular, and Optical Physics* 90.3 (2014), p. 033410.
- [198] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland. "Multi-element logic gates for trapped-ion qubits." In: *Nature* 528.7582 (2015), pp. 380–383.
- [199] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk. "T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation." In: *Applied Physics Letters* 88.3 (2006), p. 034101.
- [200] R. B. Blakestad, C. Ospelkaus, A. P. Vandevender, J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. "High-fidelity transport of trapped-ion qubits through an X-junction trap array." In: *Physical Review Letters* 102.15 (2009), p. 153002.

- [201] K. Wright, J. M. Amini, D. L. Faircloth, C. Volin, S. Charles Doret, H. Hayden, C. S. Pai, D. W. Landgren, D. Denison, T. Killian, R. E. Slusher, and A. W. Harter. "Reliable transport through a microfabricated X-junction surface-electrode ion trap." In: New Journal of Physics 15 (2013), p. 033004.
- [202] G. Shu, G. Vittorini, A. Buikema, C. S. Nichols, C. Volin, D. Stick, and K. R. Brown. "Heating rates and ion-motion control in a y-junction surface-electrode trap." In: *Physical Review A Atomic, Molecular, and Optical Physics* 89.6 (2014), p. 062308.
- [203] H. Kaufmann, T. Ruster, C. T. Schmiegelow, M. A. Luda, V. Kaushal, J. Schulz, D. Von Lindenfels, F. Schmidt-Kaler, and U. G. Poschinger. "Fast ion swapping for quantum-information processing." In: *Physical Review A* 95.5 (2017), p. 052319.
- [204] F. Splatt, M. Harlander, M. Brownnutt, F. Zähringer, R. Blatt, and W. Haänsel. "Deterministic reordering of ⁴⁰Ca⁺ ions in a linear segmented Paul trap." In: *New Journal of Physics* 11 (2009), p. 103008.
- [205] A. Bermudez, P. Schindler, T. Monz, R. Blatt, and M. Müller. "Micromotion-enabled improvement of quantum logic gates with trapped ions." In: *New Journal of Physics* 19.11 (2017), p. 113038.
- [206] C. Marquet, F. Schmidt-Kaler, and D. F. James. "Phonon-phonon interactions due to non-linear effects in a linear ion trap." In: *Applied Physics B: Lasers and Optics* 76.3 (2003), pp. 199–208.
- [207] C. W. Chou, D. B. Hume, J. C. Koelemeij, D. J. Wineland, and T. Rosenband. "Frequency comparison of two high-accuracy Al+ optical clocks." In: *Physical Review Letters* 104.7 (2010), p. 070802.
- [208] M. J. Kochenderfer and T. A. Wheeler. Algorithms for Optimization. MIT Press, 2019.
- [209] K. P. Murphy. Machine learning: a probabilistic perspective. MIT press, 2012.
- [210] D. Stathakis. "How many hidden layers and nodes?" In: *International Journal of Remote Sensing* 30.8 (2009), pp. 2133–2147.
- [211] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland. "Laser Cooling to the Zero-Point Energy of Motion." In: *Physical Review Letters* 62.4 (1989), pp. 403–406.
- [212] I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. MIT Press, 2016.
- [213] F. A. Potra and S. J. Wright. "Interior-point methods." In: *Journal of Computational and Applied Mathematics* 124.1 (2000), pp. 281–302.
- [214] A. D. Patterson, J. Rahamim, T. Tsunoda, P. A. Spring, S. Jebari, K. Ratter, M. Mergenthaler, G. Tancredi, B. Vlastakis, M. Esposito, and P. J. Leek. "Calibration of a cross-resonance two-qubit gate between directly coupled transmons." In: *Phys. Rev. Applied* 12.6 (2019), p. 064013.
- [215] J. Kelly et al. "Scalable in situ qubit calibration during repetitive error detection." In: *Phys. Rev. A* 94.3 (2016), p. 032321.
- [216] L. Gerster, F. Martinez-Garcia, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Muller, R. Blatt, P. Schindler, and T. Monz. "Experimental bayesian calibration of trapped-ion entangling operations." In: *PRX Quantum* 3.2 (2022), p. 020350.
- [217] D. Finkel and C. Kelley. "Convergence analysis of the DIRECT algorithm." In: *N. C. State Univ.* 14 (2004).
- [218] S. Ejtemaee and P. C. Haljan. "3D Sisyphus cooling of trapped ions." In: *Phys. Rev. Lett.* 119.4 (2017), p. 043001.
- [219] 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." In: *Physical Review Letters* 117.6 (2016), p. 060504.
- [220] M. Müller, K. Hammerer, Y. L. Zhou, C. F. Roos, and P. Zoller. "Simulating open quantum systems: from many-body interactions to stabilizer pumping." In: *New Journal of Physics* 13.8 (2011), p. 085007.
- [221] J. Foldager, A. Pesah, and L. K. Hansen. "Noise-assisted variational quantum thermalization." In: *Scientific Reports* 12.1 (2022), p. 3862.

- [222] D. Zhu, S. Johri, N. M. Linke, K. A. Landsman, C. H. Alderete, N. H. Nguyen, A. Y. Matsuura, T. H. Hsieh, and C. Monroe. "Generation of thermofield double states and critical ground states with a quantum computer." In: *Proceedings of the National Academy of Sciences* 117.41 (2020), pp. 25402–25406.
- [223] H. Wang, S Ashhab, and F. Nori. "Quantum algorithm for simulating the dynamics of an open quantum system." In: *Phys. Rev. A* 83.6 (2011), p. 062317.
- [224] L. Del Re, B. Rost, A. F. Kemper, and J. K. Freericks. "Driven-dissipative quantum mechanics on a lattice: Simulating a fermionic reservoir on a quantum computer." In: *Phys. Rev. B* 102.12 (2020), p. 125112.
- [225] A. Bermudez, X. Xu, R. Nigmatullin, J. O'Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. "Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation." In: *Physical Review X* 7.4 (2017), p. 041061.
- [226] P. W. Shor. "Scheme for reducing decoherence in quantum computer memory." In: *Phys. Rev. A* 52.4 (1995), R2493–R2496.
- [227] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne. "Benchmarking quantum computers: The five-qubit error correcting code." In: *Phys. Rev. Lett.* 86.25 (2001), pp. 5811–5814.
- [228] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland. "Realization of quantum error correction." In: *Nature* 432.7017 (2004), pp. 602–605.
- [229] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. "Realization of three-qubit quantum error correction with superconducting circuits." In: *Nature* 482.7385 (2012), pp. 382–385.
- [230] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt. "Experimental repetitive quantum error correction." In: *Science* 332 (2011), pp. 1059–1062.
- [231] N. M. Linke, M. Gutierrez, K. A. Landsman, C. Figgatt, S. Debnath, K. R. Brown, and C. Monroe. "Fault-tolerant quantum error detection." In: *Science Advances* 3.10 (2017), e1701074.
- [232] D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt. "Quantum computations on a topologically encoded qubit." In: *Science* 345.6194 (2014), pp. 302–305.
- [233] F. Reiter, A. S. Sørensen, P. Zoller, and C. A. Muschik. "Dissipative quantum error correction and application to quantum sensing with trapped ions." In: *Nature Communications* 8.1 (2017), p. 1822.
- [234] A. Peres. "Reversible logic and quantum computers." In: Phys. Rev. A 32.6 (1985), pp. 3266–3276.
- [235] J. Kelly et al. "State preservation by repetitive error detection in a superconducting quantum circuit." In: *Nature* 519.7541 (2015), pp. 66–69.
- [236] J. R. Wootton and D. Loss. "Repetition code of 15 qubits." In: Phys. Rev. A 97.5 (2018), p. 052313.
- [237] J. Herrmann et al. "Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases." In: *Nature Communications* 13.1 (2022), p. 4144.
- [238] I. Cong, S. Choi, and M. D. Lukin. "Quantum convolutional neural networks." In: *Nature Physics* 15.12 (2019), pp. 1273–1278.
- [239] G. Mazzola, P. J. Ollitrault, P. K. Barkoutsos, and I. Tavernelli. "Nonunitary operations for ground-state calculations in near-term quantum computers." In: *Phys. Rev. Lett.* 123.13 (2019), p. 130501.
- [240] Z. Hu, R. Xia, and S. Kais. "A quantum algorithm for evolving open quantum dynamics on quantum computing devices." In: *Scientific Reports* 10.1 (2020), p. 3301.
- [241] J. Wu and T. H. Hsieh. "Variational thermal quantum simulation via thermofield double states." In: *Phys. Rev. Lett.* 123.22 (2019), p. 220502.

- [242] H. B. Enderton. "Chapter o Useful Facts about Sets." In: *A Mathematical Introduction to Logic (Second Edition)*. Boston: Academic Press, 2001, pp. 1–10.
- [243] J. D. Jost, J. P. Home, J. M. Amini, D. Hanneke, R. Ozeri, C. Langer, J. J. Bollinger, D. Leibfried, and D. J. Wineland. "Entangled mechanical oscillators." In: *Nature* 459.7247 (2009), pp. 683–685.
- [244] M. DeCross, E. Chertkov, M. Kohagen, and M. Foss-Feig. *Qubit-reuse compilation with mid-circuit measurement and reset*. URL: https://arxiv.org/abs/2210.08039.
- [245] J. Yirka and Y. Subacsı. "Qubit-efficient entanglement spectroscopy using qubit resets." In: *Quantum* 5 (2021), p. 535.
- [246] E. Chertkov, J. Bohnet, D. Francois, J. Gaebler, D. Gresh, A. Hankin, K. Lee, D. Hayes, B. Neyenhuis, R. Stutz, A. C. Potter, and M. Foss-Feig. "Holographic dynamics simulations with a trapped-ion quantum computer." In: *Nature Physics* 18.9 (2022), pp. 1074–1079.
- [247] J. Rickert. "Simultaneous and individual ion addressing for quantum information processing." Master's thesis. University of Innsbruck, 2018.
- [248] I. Pogorelov et al. "Compact ion-trap quantum computing demonstrator." In: *PRX Quantum* 2.2 (2021), p. 020343.
- [249] A. Maitra, D. Leibfried, D. Ullmo, and H. Landa. "Far-from-equilibrium noise-heating and laser-cooling dynamics in radio-frequency Paul traps." In: *Physical Review A* 99.4 (2019), p. 043421.
- [250] J. R. Johansson, P. D. Nation, and F. Nori. "QuTiP: An open-source Python framework for the dynamics of open quantum systems." In: *Computer Physics Communications* 183.8 (2012), pp. 1760–1772.
- [251] G. A. Gottwald and I. Melbourne. "The o-1 Test for Chaos: A Review BT Chaos Detection and Predictability." In: ed. by C. H. Skokos, G. A. Gottwald, and J. Laskar. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016. Chap. Gottwald20, pp. 221–247.
- [252] E. W. Weisstein. Chaos. url: https://mathworld.wolfram.com/Chaos.html.
- [253] Neural networks. URL: https://youtube.com/playlist?list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi.