Cryogenic, near-field quantum logic chips with passive field nulling on $^{43}Ca^{+}$

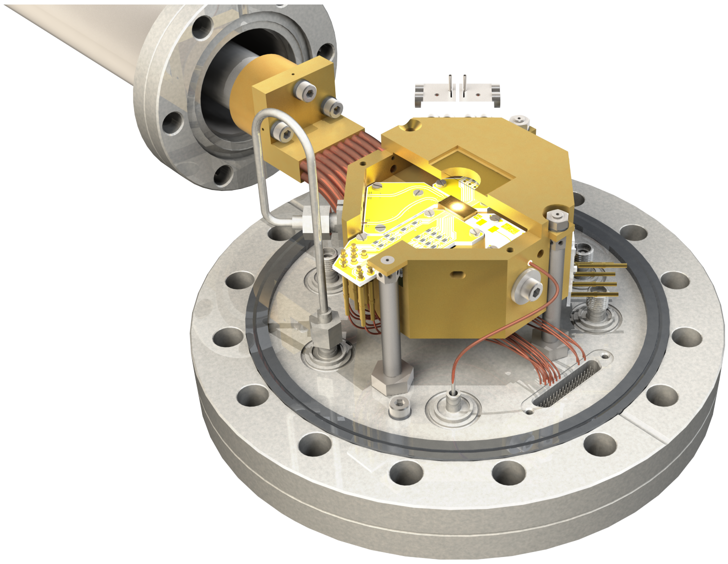

Jochen Wolf

Department of Physics, University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Atomic and Laser Physics

St Hugh's College

Trinity Term, 2019

Abstract

Cryogenic, near-field quantum logic chips with passive field nulling on $^{43}\mathrm{Ca}^{+}$
Jochen Wolf
A thesis submitted for the degree of Doctor of Philosophy
Trinity Term 2019
St Hugh's College, Oxford

This work describes the design and implementation of a new apparatus to achieve quantum information processing in trapped ions using microwave methods. The apparatus involves many design improvements over a previous quantum processing experiment which achieved the highest single-qubit gate fidelity, longest coherence time, and highest microwave based two-qubit fidelity [24][p. 38, 51][26]. The goal of this new experiment is a two-qubit gate speed and a fidelity improvement of a factor of 10, making microwave based gates much more feasible as quantum processors. To this end, a novel $4\mathrm{S}_{1/2}|3,+1\rangle \rightarrow 4\mathrm{S}_{1/2}|4,+1\rangle$ clock qubit within $^{43}\mathrm{Ca}^{+}$ is chosen, and the ion height is nearly halved. An anticipated heating rate increase is counteracted by the use of cryogenics. A novel ion chip design is implemented via wafer-scale fabrication and subsequently attached using a novel eutectic bonding technique. The implementation required many design and fabrication problems to be solved; these are described. The vacuum system reaches $< 10^{-11}$ mbar even at room temperature, lowering experimental difficulty and allowing for performance comparisons at a wide temperature range. A new experiment control system, ARTIQ, and corresponding Sinara hardware is used for control. Experimental results achieved so far include a study of contributing factors to the ion loading rate, and a comparison between the designed and measured microwave fields. All results achieved to date are compatible with the main speed and fidelity goal, which should be achieved in the near future.

Acknowledgements

I acknowledge funding by EPSRC, and more specifically the NQIT project, and the US Army. I would like to thank my supervisor, Professor Andrew Steane, for his guidance through each step of the process.

Additional supervision by Professor David Lucas and Dr Thomas Harty were greatly appreciated, as was that of Dr Joseph Goodwin in Dr Thomas Harty's absence. My research partners, Clemens Löschnauer, Marius Weber, and Ryan Hanley, were instrumental in getting the experiment to produce its results, for which I am grateful. I would also like to thank Dr James Tarlton for his insights regarding the initial design of the trap. I would also like to thank him and Keshav Thirumalai for both many fruitful discussions, and help in staying sane. Additional thanks go to Hannes Wolf, Markus Blöchl, and Alexander Hanel for proofreading. I would like to thank all members of the research group's cake club for their wonderful cakes. Also thanks to my family for their support. Lastly, I would like to express my sincere and very special thanks to Anna Moosbauer for providing absolutely invaluable feedback and support.

Contents

Nomenclature

amu	atomic mass unit, $\approx 1.66 \cdot 10^{-27}$ kg
AOM	Acousto-optic modulator
ARTIQ	Advanced Real-Time Infrastructure for Quantum physics
BSP	British standard plumbing
CF	ConFlat, a type of vacuum flange
COM	Spectroscopic centre-of-mass; the average of the energy levels weighted by their transition strength
DC	Direct Current
DDS	Direct digital synthesizer, a device for creating RF frequencies
DI	Deionised water, i.e. highly pure
EEM	Eurocard Extension Module, a PCB form factor
ENIG	Electroless nickel immersion gold, a PCB technique for adding a layer of Ni, and then Au on top of the conductor
EOM	Electro-optic modulator
FOFI	A hyperfine qubit with a first-order magnetic field insensitive transition
FPGA	Field-programmable gate array
HFSS	High Frequency Structure Simulator, a CAD software from ANSYS Inc used to simulate the ion trap
HPLC	high performance liquid chromatography, a technique requiring clean solvents
IPA	A type of solvent (2-Propanol)
Kapton	A UHV compatible plastic (Polyimide)
Macor	Trademark for a UHV compatible machineable glass-ceramic
micro-D	microminiature D, a type of electrical connector

NA	Numerical aperture
NIDC	the American National Isotope Development Center
NMP	A type of solvent (1-Methyl-2-pyrrolidone)
NQIT	National Quantum Information Technologies, a major funding body for this thesis
OFHC	Oxygen Free High thermal Conductivity, a pure grade of Copper with high RRR
OSA	Optical spectrum analyser
PCB	Printed Circuit Board
PDH	Pound–Drever–Hall
PEEK	A type of UHV compatible plastic (Polyether ether ketone)
PID	Proportional–integral–derivative, a control loop feedback system
PMT	Photomultiplier tube, a light detector
PP	A type of plastic (polypropylene)
PTFE	A type of UHV compatible plastic (Polytetrafluoroethylene, aka. Teflon)
QCCD	quantum charge-coupled device
RF	Radio frequency
RGA	A type of mass spectrometer used for detecting the composition of contaminants (Residual Gas Analyser)
RID	Run IDentification, used to identify ARTIQ experiments
RRR	Residual Resistance Ratio, resistance at 300 K divided by that at 0 K, measure of sample purity
SMA	SubMiniature version A, a type of coaxial RF connector
SNR	Signal to noise ratio
SPAM	State preparation and measurement
SS	Stainless steel
SS 304	A non-magnetic steel alloy, $\approx$18% Cr, $\approx$ 8% Ni
SS 316	A non-magnetic steel alloy, $\approx$17% Cr, $\approx$12% Ni, $\approx$ 2% Mo
subs.	subsection

TTL	Transistor-Transistor Logic
UHV	Ultra high vacuum, $1 \cdot 10^{-9}$ to $1 \cdot 10^{-12}$ mbar
USB	Universal Serial Bus
VNA	Vector Network Analyser

Chapter 1

Introduction

Moore's law stated that the number of transistors in a viable computer doubles every two years. Although this has held for about four decades, such exponential growth cannot go on indefinitely; in fact Gordon Moore himself is predicting an end to his law by 2025 [14]. One approach to this problem could be to change the computing basis from classical bits, that can either be 1 or 0, to their quantum analogue, qubits. These have a state space with a tensor product structure, which can be used to accelerate certain computing tasks by potentially many orders of magnitude. Some of the most prominent ones are Grover's search algorithm, useful for faster searches of a virtual database, and Shor's algorithm, which can be used to factorize integers, and thus can break certain types of encryption.

Quantum computing research has been going on for about four decades [19]. Significant progress has been made, and quantum supremacy could be demonstrated within the next few years; although a large-scale, universal, fault-tolerant quantum computer still requires significant advances with a much longer time frame [p. 177f, 28]. For constructing a quantum computer, "five plus two" requirements need to be met, known as the DiVincenzo criteria [18]. So far, the most used technological implementations on the path towards simultaneously achieving all of these criteria are trapped ions, and superconducting circuits [p. 2, 28]. While both of these systems have their merits, this work focuses on trapped ions.

The DiVincenzo criteria, and the trapped ion path towards their implementation, are as follows:

1. *A scalable physical system with well characterized qubits.* Using trapped ions, scalability can be achieved by using ion trapping chips, which can be scaled up using a combination of placing many ions onto a single chip, the quantum charge-coupled device (QCCD) approach, and interconnecting chips, the modular approach. This work's ion trap chip is introduced in sect. 3.2. The qubits are energy levels within specific isotopes, as such they are fixed and well known, see sect. 2.5. While there are significant engineering challenges to scaling up, there are fundamentally solvable.

2. *The ability to initialize the state of the qubits to a simple fiducial state.* For ions, an unknown state can be transferred to a known one via optical pumping, see sect. 2.7, effectively initializing the state.

3. *Long relevant decoherence times.* An isolated ion in a vacuum is well decoupled from its environment; the longest decoherence time of any type of qubit has been achieved with such a system, see sect. 2.3.

4. *A universal set of quantum gates.* The schemes for the microwave based gates are introduced in sect. 2.8 and 2.9; they do form a universal set.

5. *A qubit-specific measurement capability.* An ion can be made to fluoresce depending on the state of its qubit, see sect. 2.7. Combined with criterion #2, this is known as state preparation and measurement (SPAM), which can be achieved to high fidelities.

While the above criteria are sufficient for localised computing, requirements for quantum communication should also be taken into account:

6. *The ability to interconvert stationary and flying qubits.* Typical ions used for quantum computing have transitions in the optical domain; these can then be used to create photonic qubits out of stationary ones.

7. *The ability to faithfully transmit flying qubits between specified locations.* To achieve this, photons created by ions are coupled into fibres, and then transmitted.

For criterion #4, while some implementations of both single- and two-qubit gates are beyond the error correction threshold, only single-qubit gates are *well* beyond it. For two-qubit gates it would be useful to reduce the error rates by an additional order of magnitude, while also increasing its speed by an order of magnitude; this is the goal of this experiment, see sect. 2.3.

To focus on improving on criterion #4, and to reduce the technical difficulty of the experiment, criterion #1 is only partially met. While the ion trapping chip technology is fundamentally scalable, the bulk optics used here are not. This is an active area of research, and various ways are being investigated for achieving this scalability, for example using waveguides and gratings to distribute laser light [p. 30-33, 9].

Furthermore, no attempt is made to meet criteria #6 and #7 in this experiment. Again, this is an active area of research, and it seems possible to create such quantum communications while still meeting the other criteria, as demonstrated by a recent experiment remotely entangling two ion trap chip system via a fibre optic link [54].

All in all, this work describes the next iteration of the ion traps designed to implement part of a quantum computer. Improvements over the last iteration, [51], include cryogenics, which is used to reduce the heating rate of the ions, a major contributor to two-qubit gate errors. Furthermore, a novel ion trapping chip design is used, with the capability of shuttling ions, splitting ion chains, and passive partial microwave field nulling, which should reduce the experimental complexity, and increase the overall scalability.

Chapter 2

Experimental methodology

2.1 Qubit Criteria

The physical implementation of a ion trap quantum information processor starts with the choice of which element – and which isotope of that element – to trap. This choice is an important one that has far-reaching implications for apparatus design, and ultimate performance as a qubit. As such, this choice is made before any of the subsequent design aspects can be determined.

For trapped ions, there are generally four types choices of qubits [p. 1, 9]. As an example, they are introduced on Ca^+, whose energy level diagram is shown in fig. 2.2 & 2.3:

Optical The two qubit states are separated by an optical transition with a frequency of order $100\ldots1000\,\mathrm{THz}$, created by a pair of states in the groundstate and a metastable level, with lifetimes of order $\sim 1\,\mathrm{s}$. In Ca^+, this is typically a state in $4S_{1/2}$ and one in $4D_{5/2}$.

Fine Fine structure frequency splittings of order $1\ldots10\,\mathrm{THz}$ are typical for this qubit. In Ca^+, this would be a state in $4D_{3/2}$ and one in $4D_{5/2}$.

Hyperfine Both of the qubit states fall within the hyperfine structure of a particular fine structure level. This implies an isotope with a non-zero nuclear spin $I \neq 0$. The typical frequency splitting is on the order of $1 \ldots 10\,\mathrm{GHz}$. In $^{43}\mathrm{Ca}^+$, this is typically a pair of states in $4\mathrm{S}_{1/2}$, with one having a total angular momentum F of $F = 3$ and one with $F = 4$.

Zeeman The qubit is split by a frequency determined by an applied static magnetic field, typically of order $1 \ldots 100\,\mathrm{MHz}$. In Ca^+, this is a pair of states in $4\mathrm{S}_{1/2}$, with the same total angular momentum F[1].

As can be seen, the used frequency ranges span many orders of magnitude. Notably, the choice of a type of qubit using a metastable state places a fundamental limit on the achievable coherence time. Choosing two states within a ground state hyperfine manifold eliminates the energy relaxation time T_1, which is a major advantage over many other implementations of quantum computing. Before discussing their relative strengths and weaknesses, it is necessary to look at the ways in which operations on the qubit can be performed:

Optical The transition is addressed directly with a single laser; this is used exclusively with the optical qubit type. Due to the nature of their metastable states, this is an electric-dipole forbidden transition. To address this electric-quadrupole transition[2], a laser with a linewidth of $\sim 1\,\mathrm{Hz}$ is needed.

Raman The transition is addressed indirectly, via a third energy level, with the help of a pair of lasers coupling to the electric-dipoles. Their frequency difference δ is equal to the frequency splitting of the qubit. Additionally, they are off-resonant by a fairly large frequency Δ from the third energy level; in Ca^+, this is typically the $4\mathrm{P}_{1/2}$ level. This method has been used on all of the above qubit types. The two lasers are typically derived from the same laser source as their phase and amplitude relationship needs to be well controlled.

[1] In $^{40}\mathrm{Ca}^+$, this is simply both states of the $4\mathrm{S}_{1/2}$ manifold.

[2] Electric-octupole qubit transitions exist as well, but typically do not offer higher performance than their electric-quadrupole counterparts due to technical limitations on the quality of available lasers.

Microwave The transition is addressed directly with a microwave or RF drive. This method is applicable to the hyperfine and Zeeman qubit types. As their states are in the same fine structure manifold, these are magnetic-dipole transitions, see sect. 2.8.1. As the transition wavelengths are large compared with the size of an ion, the coupling of the drive to the ion's motion is very low in free space. To achieve higher coupling, the ions can be brought close to the source of the drive field; this vastly increases the coupling strength, see sect. 2.9.1.

Table 2.1 presents the ions most commonly used in quantum information[3]. The considerations for choosing the ions are:

Element	m amu	I	$\lambda_{1/2}, \lambda_{3/2}, \lambda_{5/2}$ nm	τ s	B_0 mT	Qubits	Gates
Be^+	9	3/2	313, 313, N.A.	N.A.	11.94	H	R, M
Mg^+	25	5/2	280, 280, N.A.	N.A.	10.9	H	R
Ca^+	40	0	397, 393, 729	1.17		Z, F, O	R, M
	43	7/2			14.61	H, O	O, R
Sr^+	87	9/2	422, 408, 674	0.34			
	88	0				O	O, R
Cd^+	111	1/2	226, 214, N.A.	N.A.	0.	H	R
Ba^+	133	1/2	493, 455, 1762	35.0	0.	H, O	
	137	3/2			0.	H, O	
	138	0				O	R
Yb^+	171	1/2	369, 329, 411	0.007	0.	H	R,M
	172, 174	0				O	

Table 2.1: Properties of the ions most commonly used in ion trapping for quantum information. The listed properties are: their mass m, nuclear spin I, mean lifetime τ, and lowest magnetic field B_0 for which a first-order magnetic field insensitive exists that has been used in an experiment. The transition wavelengths λ from the ground fine structure states to the $P_{1/2}$, $P_{3/2}$, and $D_{5/2}$ levels are denoted as $\lambda_{1/2}$, $\lambda_{3/2}$, and $\lambda_{5/2}$, respectively. Elements that do not have a $D_{5/2}$ level are marked by a "N.A.". Lastly, typical qubits used are denoted as Z (Zeeman), H (hyperfine), F (fine), and O (optical) and typical gates used are R (Raman), O (optical), M (magnetic [AC or static gradient]). Reproduced from [p. 34, 9].

1. In general, the energy of the states of all types of qubits are dependent on

[3]Note that all elements (that have a stable isotope) belonging to group II of the periodic table are present; it is useful to have a single valence electron, and singly ionising group II elements results in just that.

the static magnetic field. Any changes in the magnetic therefore induce a change in the transition energy and thus a change in frequency with respect to the drive field. This then leads to both increased gate errors and decreased coherence times. However, for hyperfine qubits, there exist pairs of states that are to first-order magnetic field insensitive (FOFI). In practise, this increases the coherence times by some orders of magnitude, e.g. from $\sim 0.1\,\mathrm{s}$ to $50\,(10)$ s.

2. The higher the mass of the ion, the slower the achievable gate speeds, with other parameters being equal, so a low-mass ion is preferable.

3. For the prospect of a large-scale quantum processor, a second set of ions of a different element – but of a similar mass – is very useful for sympathetic cooling, error-correction, and remote-entanglement [p. 28, 9].

4. Generally, the smaller the wavelengths of the lasers used, the larger the technical difficulty of the experiments. In particular, lasers in the UV range become problematic due to their tendency to charge the substrates, and generally have worse performance compared to their IR range counterparts.

5. The existence of metastable states in an ion is useful for increasing readout fidelities. The downside is a more complicated level structure which opens the laser cooling cycle to additional levels.

All of these points will have to be considered for the choosing a suitable element, and isotope, for quantum information processing experiments.

2.2 Qubit Choice

For this experiment, Ca^+ is the element of choice. The energy level structure is shown in figs. 2.2 & 2.3. The reasons are outlined as follows:

1. It is an ion with the lowest mass of the ions commonly used in ion trapping that still has metastable excited states, see tab. 2.1. It also permits a second set of ions of a different element, but similar mass to be used for ancillary purposes such as sympathetic cooling. Two choices have already been demonstrated, $^{9}Be^{+}$ and $^{88}Sr^{+}$ [41, 31].

2. While its wavelengths of 393 nm and 397 nm are technically in the $10 \dots 400$ nm range of UV light, they do not typically cause substrate charging, and diode lasers and optical components of sufficient performance are available.

3. It has a set of metastable excited states, $3D_{3/2}$, and $3D_{5/2}$. The scheme to mitigate the impact of the more complicated level structure on the readout and laser cooling are introduced in sects. 2.6 & 2.7.

4. Ca^{+} is very competitive in published fidelities and gate speeds [p. 17, 9].

All four types of qubits and all three types of operations can be performed on Ca^{+}.

Considering types of qubits, the limited lifetime of the optical and fine structure qubits severely limits their use as quantum memory. Zeeman qubits place high requirements on the stability magnetic field, which may require superconducting magnets running in persistent mode to achieve; this is beyond the scope of this experiment. Hyperfine qubits are attractive, particularly due to the existence of FOFI states.

Looking at operations, optical types need lasers that are technically demanding, represent the state of the art, and typically suffer from reliability issues. In the case of Raman stimulation, the need of the two lasers to be interferometrically stable introduces a significant source of errors. Both optical and Raman stimulation also suffer from errors due to off-resonant excitation, which places a limit of $\sim 99.99\%$ on the fidelity of the operations. As will be apparent in the next section, this fidelity limit is problematic. Microwave stimulation offers comparatively much cheaper, more reliable, and higher quality[4] drive field sources than their laser counterparts, and do not have the off-resonant excitation and interferometric-stability errors associated with Raman stimulations.

Therefore, an attractive choice is the use of a hyperfine FOFI qubit directly stimulated via microwaves. This implies the use of $^{43}Ca^{+}$, with its complicated level structure of 144 states involved in its cooling and readout cycles. Previous work has used the ground state hyperfine FOFI qubit with the lowest magnetic field, the $4S_{1/2}|F = 4, m_F = 0\rangle \rightarrow |3, +1\rangle$ qubit at a magnetic field of 14.609 42 mT [51]. For reasons explained in the next section and in sect. 2.9.4, this work uses the $4S_{1/2}$ hyperfine FOFI qubit with the second lowest magnetic field, $|4, +1\rangle \rightarrow |3, +1\rangle$, field insensitive at 28.787 27 mT.

2.3 Comparison with previous experiment

This experiment is intended to be an improvement over a previous microwave-driven experiment on a magnetic field insensitive qubit in $^{43}Ca^{+}$ [51]. Its performance characteristics were a state preparation and measurement fidelity of 99.93%, a memory coherence time of $T_2^* = 50\,\mathrm{s}$, a single-qubit gate fidelity of 99.9999% at a gate duration of $t_{g1} = 1 \ldots 2\,\mu\mathrm{s}$, and a two-qubit gate fidelity of $99.7\,(1)\,\%$ at a gate duration of $t_{g2} = 3.25\,\mathrm{ms}$ [24][p. 38, 51][26].

The main goal of this experiment is to improve the two-qubit gate results in both speed and fidelity by an order of magnitude[5]. To achieve this, the ion-to-surface distance d and the temperature of the trap T are reduced, a different qubit transition is chosen, the polarisation of the qubit transition is now π instead of σ, and the magnetic field B is increased:

d The ion's height above the surface is reduced from $75\,\mu\mathrm{m}$ to $40\,\mu\mathrm{m}$.

[4] As measured by the phase, amplitude, and frequency control.

[5] As such it would be problematic to use Raman stimulation as the basis for two-qubit gates as the fidelity limit from off-resonant excitation nearly coincides with the targeted fidelity, and any other sources of errors would then likely make the target unattainable.

All else being equal, this should increase the strength of the microwave field by a factor of ~ 2, at the cost of an expected increase in the heating rate by a factor of $\left(\frac{75\,\mu\mathrm{m}}{40\,\mu\mathrm{m}}\right)^4 = 12$, which is dealt with in by lowering the trap temperature T.

T The trap, together with its mounting structure, could be cooled to $\approx 4\,\mathrm{K}$ using liquid helium, as opposed to the previous experiment limited to room temperature. As roughly 1/3 of the previous experiment's error budget was due to the heating rate, the above factor of 12 increase needs to be counteracted, which is done by cooling the experiment to $\approx 16\,\mathrm{K}$[6]. This should decrease the heating rate by a factor of 100.

π While both experiments use the $4S_{1/2}$ hyperfine manifold, the previous experiment used $|4,0\rangle$ to $|3,+1\rangle$ as a qubit, a $\sigma_\pm$-polarised transition. However, due to other design constraints explained in sect. 3.2, microwaves created in the near-field regime are π-polarised, resulting in a $\sqrt{2}$-decrease in coupling efficiency to that qubit. In this experiment, a different qubit is used, $|4,+1\rangle$ to $|3,+1\rangle$, which is π-polarised.

B In order to use their respective qubit's first order magnetic field insensitive points, the previous experiment used 14.6 mT, while this experiment uses 28.8 mT. This increases the energy splitting between $|\mathrm{F}, \mathrm{m_F}\rangle$ to $|\mathrm{F}, \mathrm{m_F} \pm 1\rangle$ from ≈ 50 MHz to ≈ 100 MHz, decreasing errors due to off-resonant excitation by a factor of 4, potentially lowering two-qubit gate error rates. However, typical natural line widths of the laser based transitions are $\approx 20\,\mathrm{MHz}$, and thus optical state-preparation, see sect. 2.7, is significantly slower.

More details on the speed increase are derived in sect. 2.9.4.

[6]Which is a more economical choice over 4 K, see sect. 3.6, but still provides cryogenic pumping of most gasses.

2.4 Photoionisation

In this work, neutral calcium atoms are ionised using an existing two step resonant photoionisation scheme [37]. The first step is achieved with a 423 nm laser, exciting the $4s^2\,{}^1S_0 \rightarrow 4s4p\,{}^1P_1$ transition in neutral Ca, at a nominal frequency of $f_{^{40}\mathrm{Ca}} = 709.078\,24$ THz [34]. From this level, the ionisation wavelength required is 389.81 nm [34], achieved via a 378 nm laser, exciting the electron out into the continuum. The isotope shift for the 423 nm transition is ${}_{\Delta}f_{iso} = f_{^{43}\mathrm{Ca}^+} - f_{^{40}\mathrm{Ca}^+} = +612\,\mathrm{MHz}$ [37]; this shift is irrelevant for the second step photoionisation.

When ionising ^{43}Ca in a magnetic field, the Zeeman effect in the ground state $4s^2\,{}^1S_0$ is negligible, due to the lack of electronic angular momentum and the weakness of the nuclear magnetic moment [p. 23f, 23]. The splitting of the excited state $4s4p\,{}^1P_1$ can be seen in fig. 2.1. For a magnetic field of $B = 28.8\,\mathrm{mT}$, the states are in the high field regime; they form three separate manifolds with $m_J = \{-1, 0, 1\}$, separated by $\approx 400\,\mathrm{MHz}$; the $m_J = \pm 1$ manifolds are distributed over $\approx 100\,\mathrm{MHz}$. Therefore, any π-polarised 423 nm light should be unaffected by the magnetic field, while $\sigma_\pm$-polarised should be detuned by $\pm 400\,\mathrm{MHz}$.

The laser setup involved is an existing group shared resource, a Toptica DL100 extended cavity diode laser, and a free-running second laser. These lasers are distributed among multiple fibres, one of which goes to the experiment featured in this work. Notably, both the 423 nm and the 378 nm lasers are coupled into the same fibre to facilitate alignment of the photoionisation lasers. They are switched on and off via a mechanical shutter as opposed to e.g. an AOM, as the switching speeds required are very slow, <1 s.

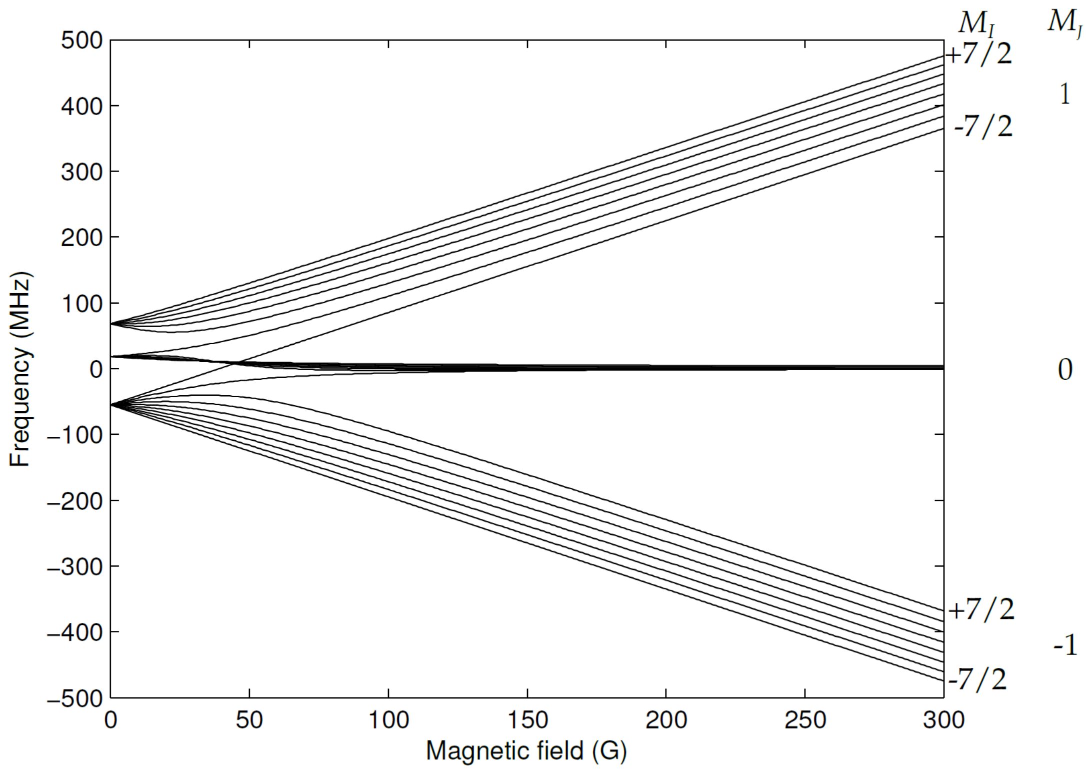

Figure 2.1: Energy splitting of $4s4p\,^1P_1$ in ^{43}Ca as a function of magnetic field [p. 24, 23].

2.5 Energy Levels

As this work features both $^{43}Ca^+$ and $^{40}Ca^+$, it is important to understand their energy levels. Both have the manifolds $4S_{1/2}$, $4P_{1/2}$, $4P_{3/2}$, $3D_{3/2}$, and $3D_{5/2}$. Due to the lack of nuclear spin I, the level structure of $^{40}Ca^+$ has only 18 states involved in its readout and cooling cycles, see fig. 2.2. This is in sharp contrast to the level structure of $^{43}Ca^+$, which has a nuclear spin of $I = 7/2$, resulting in 144 states involved in its readout and cooling cycles, see fig. 2.3. The energy levels are calculated using the following Hamiltonian H_{tot}:

$$H_{tot} = H_{GS} + H_{FS} + H_{HFS} + H_Z \tag{2.1}$$

where H_{GS} is the gross interaction, H_{FS} is the fine interaction, H_{HFS} is the hyperfine interaction, and H_Z is the Zeeman interaction. H_{HFS} can be expressed as

$$H_{HFS} = A\boldsymbol{I} \cdot \boldsymbol{J} \tag{2.2}$$

with an experimentally-determined and manifold-specific hyperfine structure constant A, and the nuclear and electronic angular momentum operators $\boldsymbol{I}$ and $\boldsymbol{J}$. H_Z is given by

$$H_Z = -\boldsymbol{\mu_B} \cdot \boldsymbol{B} - \boldsymbol{\mu_I} \cdot \boldsymbol{B} \tag{2.3}$$

with the Bohr and nuclear magneton operators $\boldsymbol{\mu_B}$ and $\boldsymbol{\mu_N}$, and constants μ_B and μ_N, the Landé g-factor g_J, and the nuclear g-factor g_I. Without loss of generality, the magnetic field operator $\boldsymbol{B}$ is aligned with the z-axis

$$H_Z = -g_J \mu_B B_z J_z - g_I \mu_N B_z I_z \tag{2.4}$$

The subsequent calculations are performed using a Python toolkit [25]. The combined contribution of H_{GS} and H_{FS} are added as experimentally determined constants, the sources for which – as well as those of the above constants – are quoted in the toolkit.

At the adopted field of $\sim 28.8\,\mathrm{mT}$, the manifolds $4\mathrm{S}_{1/2}$ and $4\mathrm{P}_{1/2}$ are still in the low-field regime, while $4\mathrm{P}_{3/2}$ and $3\mathrm{D}_{5/2}$ are in the high-field regime. $3\mathrm{D}_{3/2}$ is in the intermediate-field regime. For transitions between manifolds of different regimes, the electric dipole selection rules reduce to only forbidding transitions that do not follow $\Delta m_F = 0, \pm 1$, $\Delta F = 0, \pm 1$ (additionally, $F = 0 \rightarrow 0$ is forbidden), as well as between states of the same parity.

To avoid having to switch basis between the fine structure manifolds, the $|\mathrm{F}, \mathrm{m_F}\rangle$ state basis is used throughout this thesis. The Zeeman splitting of the $4\mathrm{S}_{1/2}$ manifold of $^{43}\mathrm{Ca}^+$ is shown in fig. 2.4.

The isotope shift of the transitions between manifolds of $^{40}\mathrm{Ca}^+$ compared to $^{43}\mathrm{Ca}^+$ is $_\Delta f{=}f_{^{43}\mathrm{Ca}^+} - f_{^{40}\mathrm{Ca}^+}$, and can be found in tab. 2.2.

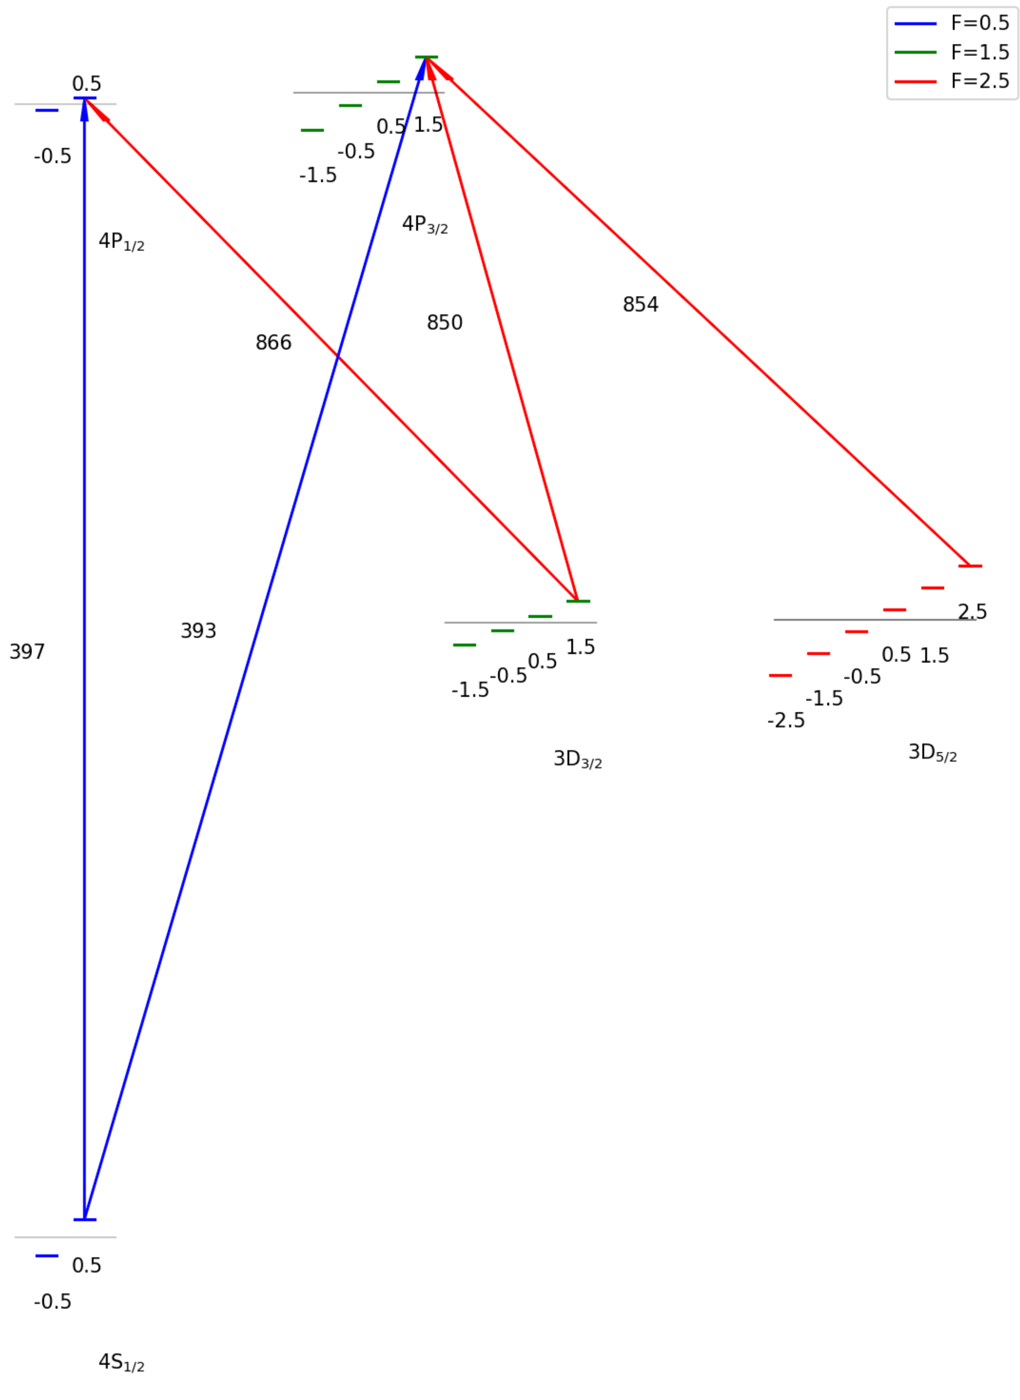

Figure 2.2: Energy levels of $^{40}Ca^+$ at a magnetic field of 28.787 27 mT, more details in fig. 2.3. Due to the lack of nuclear spin I, the total electronic angular momentum J and the total angular momentum F are the same in $^{40}Ca^+$, $F = J$. For consistency between $^{40}Ca^+$ and $^{43}Ca^+$, F is used to describe states in this thesis.

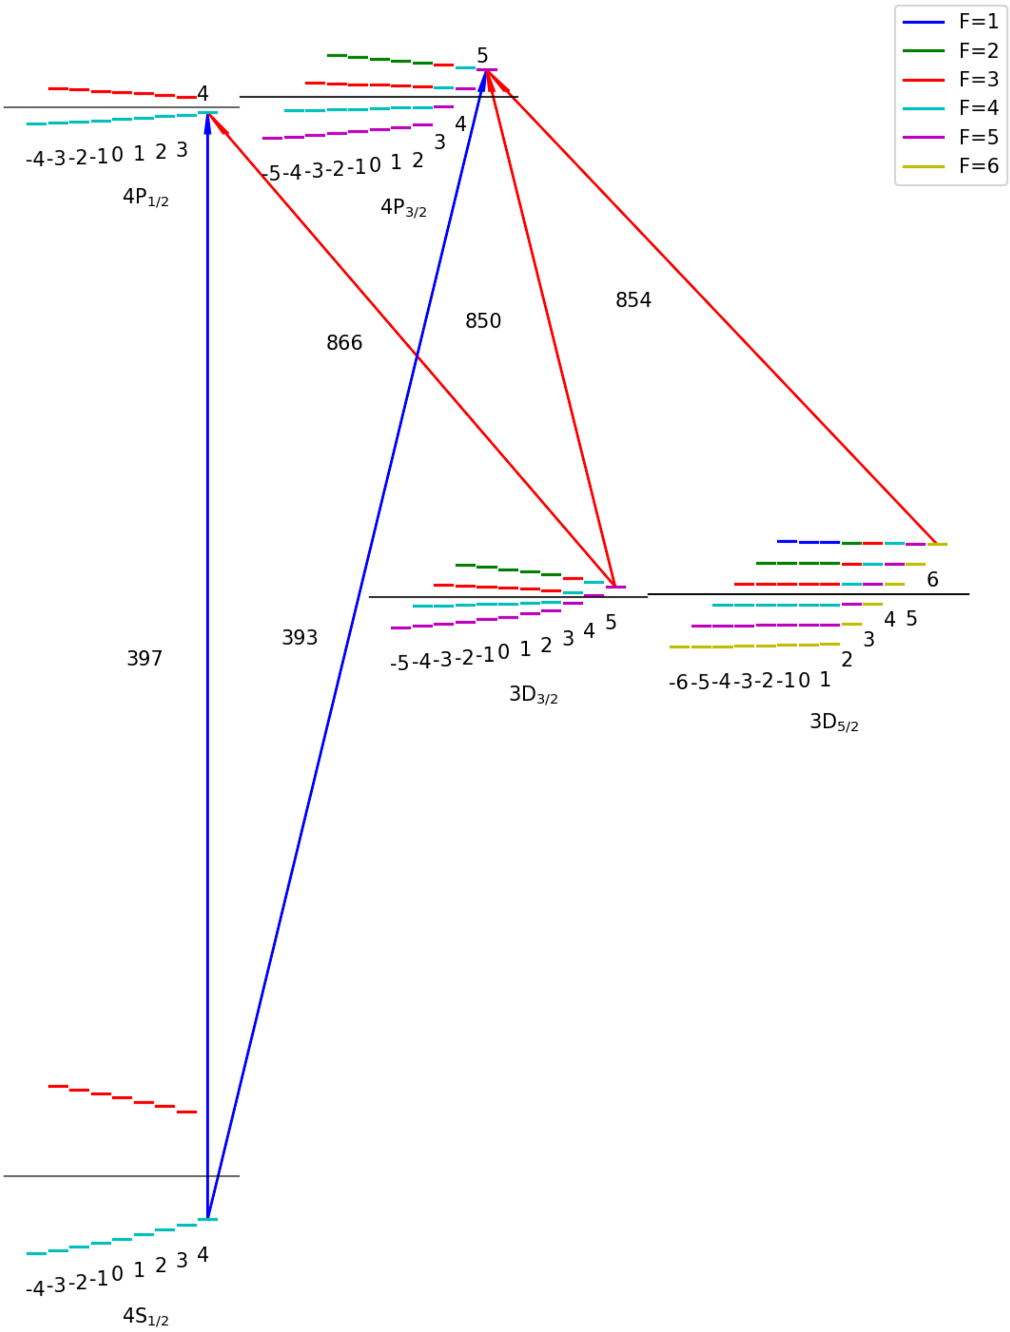

Figure 2.3: Energy levels of $^{43}Ca^{+}$ at a magnetic field of 28.787 27 mT. The horizontal black levels indicate the spectroscopic centre-of-mass (COM) of each of the fine structure manifolds. The colours of the individual energy levels indicate their total angular momentum F as indicated by the legend. The colour of the transitions indicate whether the lasers fall into the “blue” range, [378 nm . . . 423 nm] or “red” range [850 nm . . . 866 nm]; the former includes the photoionisation lasers from sect. 2.4. The Zeeman effect and Hyperfine splitting are exaggerated by a factor of 30 000; otherwise they would be indistinguishable from the COM line on this scale. The fine structure splitting is drawn on the same scale as the gross structure; note the slightly lower energy of $3D_{3/2}$ compared to $3D_{5/2}$.

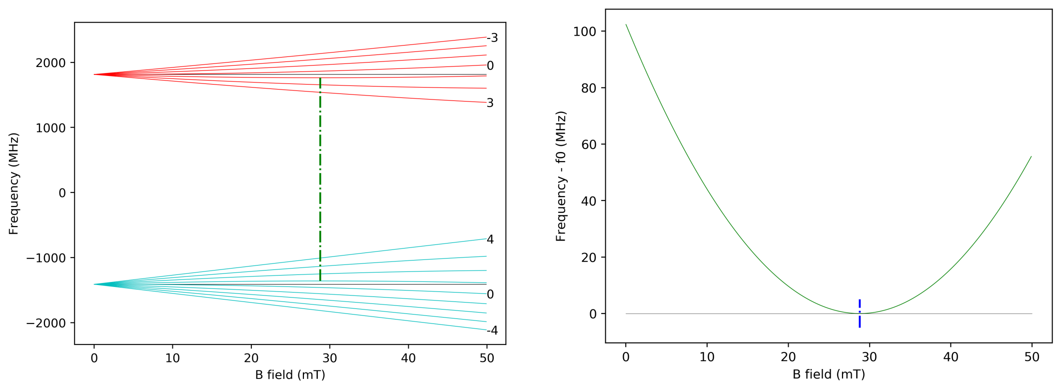

(a) (b)

Figure 2.4: Zeeman splitting of the $4S_{1/2}$ manifold of $^{43}Ca^+$. (a) The colour coding of the manifolds is the same as in fig. 2.3, i.e. F=3 is shown in red, and F=4 in cyan. The qubit transition $|3,+1\rangle \rightarrow |4,+1\rangle$ is shown in green at a magnetic field of 28.787 27 mT. The Zeeman effect is included to all orders, but in the range of magnetic fields shown it remains small compared to the hyperfine structure so the manifold is in the low-field regime. The black lines indicate the energy level of the hyperfine manifold in the absence of a magnetic field. (b) The qubit transition frequency (green) is shown as a function of magnetic field; it is subtracted from its frequency at the clock qubit magnetic field of 28.787 27 mT. It has a turning point, which makes this qubit insensitive to the magnetic field to first order; the second order sensitivity is $0.025\,14\,{}^{\mathrm{kHz}}/_{\mathrm{mT}^2}$. The frequencies of the other relevant transitions are shown in fig. 2.7.

Transition	$\lambda\left({}^{40}\text{Ca}^{+}-\text{COM}\right)$	$f\left({}^{40}\text{Ca}^{+}-\text{COM}\right)$		Isotope shift $_\Delta f_{\text{iso}}$		Hyperfine	$_\Delta f\,(B=0)$	Zeeman	$_\Delta f\,(B=B_0)$
	nm	THz	Ref	MHz	Ref	MHz	MHz	MHz	MHz
$4\text{S}_{1/2}\lvert 3,+3\rangle \rightarrow 4\text{P}_{1/2}\lvert 4,+4\rangle$	396.958 978 8(9)	755.222 766 2(17)	[58]	+688.(17)	[37]	$-(+1814.4)+(-254.4)$	-1 380.8	$-(+275.3)+(+134.4)$	-1 522.7
$4\text{S}_{1/2}\lvert 4,+4\rangle \rightarrow 4\text{P}_{3/2}\lvert 5,+5\rangle$	393.477 471 6(3)	761.905 012 7(5)	[59]	+692.(19)	[37]	$-(-1411.2)+(-164.9)$	+1 938.3	$-(+403.5)+(+806.3)$	+2 341.1
$3\text{D}_{3/2}\lvert 4,+4\rangle \rightarrow 4\text{P}_{1/2}\lvert 4,+4\rangle$	866.452	346.000	[34]	-3464.3(3.0)	[42]	$-(-11.8)+(-254.4)$	-3 706.9	$-(+361.1)+(+134.4)$	-3 933.6
$3\text{D}_{3/2}\lvert 4,+4\rangle \rightarrow 4\text{P}_{3/2}\lvert 5,+5\rangle$	850.036	352.682	[34]	-3462.4(2.6)	[42]	$-(-11.8)+(-164.9)$	-3 615.5	$-(+361.1)+(+806.3)$	-3 170.3
$3\text{D}_{5/2}\lvert 5,+5\rangle \rightarrow 4\text{P}_{3/2}\lvert 5,+5\rangle$	854.444	350.863	[34]	-3465.4(3.7)	[42]	$-(-10.7)+(-164.9)$	-3 619.6	$-(+1195.7)+(+806.3)$	3 230.2

Table 2.2: Transition properties of ${}^{40}\text{Ca}^{+}$, their isotope shifts to ${}^{43}\text{Ca}^{+}$, combined with the corresponding Hyperfine & Zeeman shifts at a magnetic field of $B_0 = 28.778\,27$ mT. The isotope shifts are $_\Delta f_{\text{iso}} = f_{{}^{43}\text{Ca}^{+}} - f_{{}^{40}\text{Ca}^{+}}$, referenced to their respective centre-of-masses and without Zeeman splitting. For the hyperfine and Zeeman splittings, first (second) numbers correspond to the shift of the lower (upper) state, while the sign in front of the brackets shows the effect on the laser's frequency. Note that in the table, the overall energy shift $_\Delta f\,(B=B_0)$ is taken with reference to ${}^{+}\text{Ca}^{40}$ at $B=0$. Calculated using [25], unless otherwise specified.

2.6 Cooling Scheme

Ions held in a trap exhibit a non-zero motional heating rate. Left unchecked, this heating gives the ion enough energy to escape the trap. The scheme used for cooling ions is explained in this section. As shown in fig. 2.3, there is a large number of states in $^{43}Ca^{+}$ that the ion could be in. Choosing a cycling transition in $^{43}Ca^{+}$, while minimising the number of lasers involved is therefore important. To this end, an existing scheme is used, [p. 24-34, 23], explained as follows, and shown in fig. 2.5.

The scheme involves a 397 nm laser, with $2/3 - \sigma_{+}$-polarisation, and $1/3 - \pi-$polarisation[7]. Without loss of generality, let us assume that the ion starts out in $4S_{1/2}|4,+4\rangle$. The 397 nm laser is tuned slightly[8] red of resonance to $4P_{1/2}|4,+4\rangle$, which has a 94% branching ratio back to $4S_{1/2}|4,3\ldots4\rangle$, resulting in Doppler cooling. To also cool the $F=3$ manifold, sidebands[9] are put onto the 397 nm laser using an electro-optic modulator (EOM). $4P_{1/2}$ can also decay into $3D_{3/2}$, with a 6% branching ratio; to recover this population, an 866 nm laser is used. To make this laser efficient at repumping all states within $3D_{3/2}$ back to $4P_{1/2}$, a comparatively high intensity is used to power broaden across all states in $3D_{3/2}$, indicated by the multitude of drawn transitions in fig. 2.5. For the 397 nm laser, the intensity is a trade-off between power broadening over the transitions and increasing the background during readout[10]. The power broadening of the 397 nm and 866 nm lasers is used to address the states other than the previously assumed $4S_{1/2}|4,+4\rangle$. Due to the predominantly σ_{+}-polarisation of the 397 nm laser, the ion tends to increase in the magnetic quantum number m_F, in a way that will be explained as part of the state-preparation in sect. 2.7.

[7] A higher ratio of σ_{+} to π-polarisations would be preferable, but this ratio is limited by the laser having a 45° angle to the magnetic field B.

[8] Typically half a natural line-width.

[9] The frequency splitting of these sidebands is experimentally optimised to be 2653 ± 50 MHz see sect. 5.2.3.

[10] To partially counteract this, specialised sets of parameters - laser intensity and detuning - are used for cooling, state-prep, and readout

While this scheme does not need the 854 nm laser during cooling operations, it is also shown in fig. 2.5 as the scheme is also involved in state-preparation, see sect. 2.7.

Note that the transitions drawn in fig. 2.5 are just the ones that are typically in use during steady-state[11].

[11] During optical pumping, other states are involved as well, but only transiently.

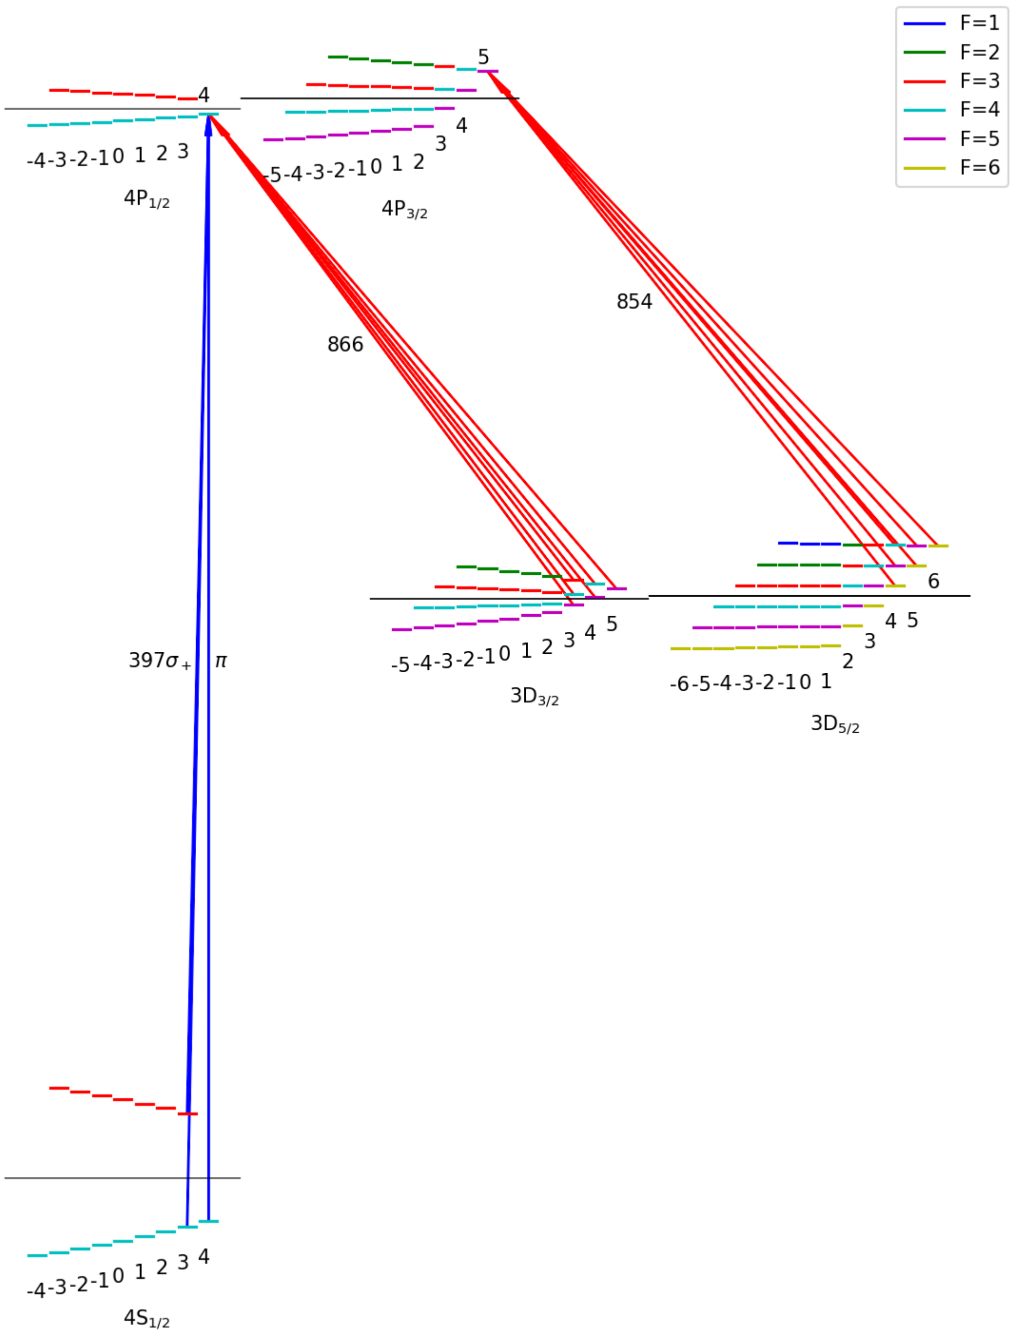

Figure 2.5: Cooling scheme of $^{43}Ca^+$, explained in sect. 2.6. Two other schemes, sect. 2.7, both have just one conceptual difference to this cooling scheme: Firstly, the state-preparation scheme uses a purely σ_+-polarised 397 nm laser, preparing the ion into $4S_{1/2}|4,+4\rangle$. Secondly, the readout scheme omits the 854 nm laser, i.e. only gives fluorescence if the ion was not shelved into $3D_{5/2}$. Note that the red laser transitions are shown to start at $3D_{3/2}|3\ldots5,3\ldots5\rangle$ and $3D_{5/2}|4\ldots6,4\ldots6\rangle$, while the blue laser starts at $4S_{1/2}|3\ldots4,3\ldots4\rangle$, limited by selection rules. Also see fig. 2.3.

2.7 SPAM Scheme

While the chosen qubit can be manipulated using microwaves, lasers are still necessary for other operations, namely state preparation and measurement (SPAM). The SPAM scheme is explained in this section. The individual steps - state preparation (state-prep), shelving, and readout - are explained in this section.

State-prep Starting with an unknown state, the ion could be in any of the 96 states within $4S_{1/2}$, $3D_{3/2}$, and $3D_{5/2}$[12]. To prepare the ion in a known state, a method very similar to the cooling scheme of sect. 2.6 is used, except that the 397 nm laser has pure σ_+-polarisation, see fig. 2.5. This empties the $3D_{5/2}$ states, while involving the $4S_{1/2}$, $4P_{1/2}$, and $3D_{3/2}$ states in a cycling transition. This means that within the $4S_{1/2}$ states, starting with an unknown m_F, the σ_+-polarisation excites this to $m_F + 1$ within $4P_{1/2}$, which in turn usually decays (with a 94% branching ratio) back to $4S_{1/2}$ with $m'_F = (m_F, m_F + 1, m_F + 2)$, with no preference between these three states. Therefore on average, this increases m_F by one per cycle[13]. Once a cycle ends up in $4S_{1/2}|4, +4\rangle$, it remains there because there is no $m_F = +5$ state within $4P_{1/2}$. This state-preparation pulse is applied for long enough to reach a target probability of successful preparation.

Shelving Any algorithm performed is assumed to leave the ion's state in either $4S_{1/2}|4, +4\rangle$, $4S_{1/2}|3, n\rangle$, or a combination thereof. A σ_+-polarised 393 nm laser is then applied, ideally together with a σ_+-polarised 850 nm laser. As opposed to the 397 nm laser, this 393 nm does not go through an EOM, and is therefore tuned to be only resonant to the

[12] While it could technically also initially start out in $4P_{1/2}$ and $4P_{3/2}$, it would decay into the above mentioned states within a few microseconds. While in principle the $3D_{3/2}$ and $3D_{5/2}$ states would also decay into the $4S_{1/2}$ states, this process takes ~ 1 s, which is too long.

[13] If the $4P_{1/2}$ state instead decays to $3D_{3/2}$, it is re-pumped back into $4P_{1/2}$, until it decays to the $4S_{1/2}$ states. While this process could drastically change m_F, the branching ratio of 6% is small enough that this effect can be neglected after cycling many times.

$4S_{1/2}|4,+4\rangle$ to $4P_{3/2}|5,+5\rangle$ transition. Therefore, only the $4S_{1/2}|4,+4\rangle$ population is transferred to $4P_{3/2}|5,+5\rangle$, while $4S_{1/2}|3,n\rangle$ is unaffected, see fig. 2.6. After cycling on the $4S_{1/2}|4,+4\rangle$ to $4P_{3/2}|5,+5\rangle$ transition, the ion eventually decays from $4P_{3/2}|5,+5\rangle$, with a 90% branching ratio, to $3D_{5/2}$[14], which are the states that are unaffected by the following readout. In this case the ion is said to be put "on a shelf" or "shelved". To recover the 10% population that decays from $4P_{3/2}|5,+5\rangle$ to $3D_{3/2}|5,+5\,\text{or}+4\rangle$, a σ_+-polarised 850 nm laser can be used to repump this population back to $4P_{3/2}|5,+5\rangle$.[15] Depending on the intensity of the applied 393 nm laser and the magnitude of the Zeeman splitting, the other states $4S_{1/2}|4,m_F \neq +4\rangle$ could partially be transferred to the $3D_{5/2}$ shelf, too, but it is more reliable to transfer them first to $4S_{1/2}|4,+4\rangle$ instead.

Readout After shelving, the ion's state is either within the $3D_{5/2}$ shelf, or $4S_{1/2}|3,n\rangle$. To determine which one, the cooling scheme of sect. 2.6 is used, but without the 854 nm laser, see fig. 2.5. For the $4S_{1/2}|3,n\rangle$ states, this creates a cycling transition involving the $4S_{1/2}$, $4P_{1/2}$, and $3D_{3/2}$ states, scattering the applied 397 nm laser light. This is collected by an imaging system detailed in sect. 3.13. For the $3D_{5/2}$-shelf, there is no scattering. To distinguish the two states, a threshold is used after collecting photons for a given time. As $3D_{5/2}$ decays within $\sim 1\,\text{s}$, there is a fundamental limit to the readout fidelity that can be achieved with this scheme, but this limit is high.

[14]Specifically, this could be $3D_{5/2}|4,+4\rangle$, $3D_{5/2}|5,+4\ldots+5\rangle$, and $3D_{5/2}|6,+4\ldots+6\rangle$

[15]Without this 850 nm laser the readout contrast is limited to 90%, which is still a reliable readout, although less useful for some purposes. Due to time limitations, this laser's beam path is not included in this work, although it will be added to the experiment in the near future.

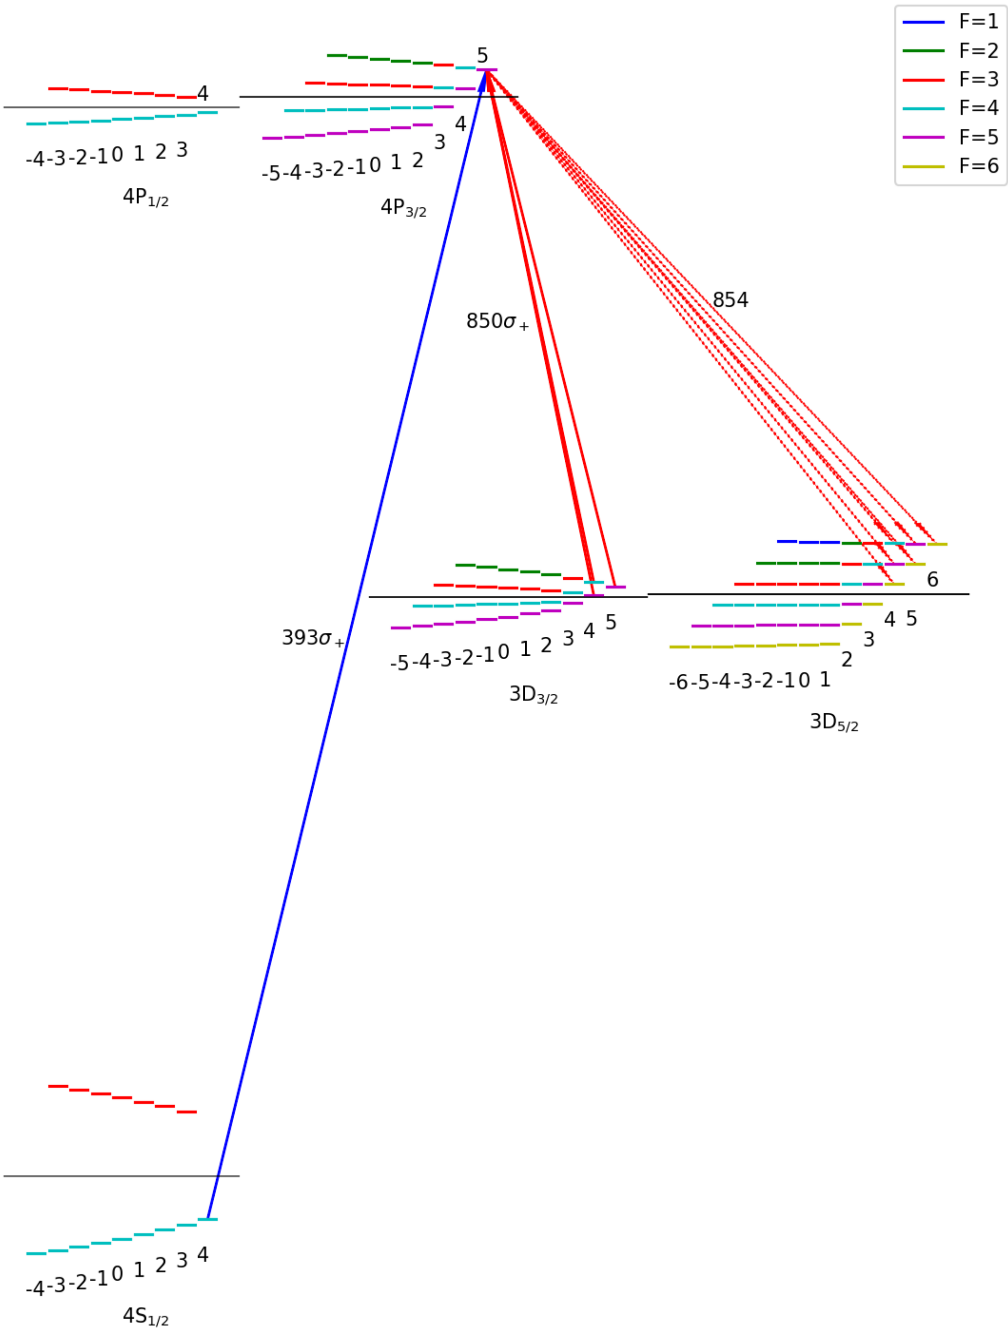

Figure 2.6: Shelving scheme of $^{43}Ca^{+}$, explained in sect. 2.7. The dotted 854 nm lines indicate the decay paths that are not re-pumped, i.e. the $3D_{5/2}$ shelf. Also see fig. 2.3.

2.8 Single-qubit gates

2.8.1 Magnetic dipole transitions

The parity of all states within the $4S_{1/2}$ manifold is $(-1)^l = +1$, i.e. even. As such, the selection rules forbid electric dipole transitions between them, see sect. 2.5. Additionally, all higher orders (quadrupole, octupole, etc.) of electrical interactions are also forbidden, as $L = 0$. Instead, *magnetic* dipole interaction is used, with the magnetic moment of the atom $\boldsymbol{\mu}$[16] and an applied radiation field $\boldsymbol{B}(t)$; the interaction is described by $-\boldsymbol{\mu} \cdot \boldsymbol{B}(t)$. The following derivation follows [p. 17ff, 51]. The Rabi frequency Ω with which two states $|F_1, m_{F,1}\rangle$ and $|F_2, m_{F,2}\rangle$ interact is

$$\Omega = -\frac{1}{\hbar}\langle F_1, m_{F,1}|\boldsymbol{\mu} \cdot \boldsymbol{B}(t)|F_2, m_{F,2}\rangle \tag{2.5}$$

The applied radiation field $\boldsymbol{B}(t)$ can be chosen to be aligned to the axis e_z. Introducing a spherical basis vector set

$$\begin{aligned} e_+ &= e_1 = +\tfrac{1}{\sqrt{2}}\left(e_x - ie_y\right) \\ e_- &= e_{-1} = -\tfrac{1}{\sqrt{2}}\left(e_x - ie_y\right) \\ e_0 &= e_z \end{aligned} \tag{2.6}$$

brings eq. 2.5 then into the form

$$\Omega = -\frac{1}{\hbar}\sum_{Q=\{\pm1,0\}}(-1)^{Q+1}\langle F_1, m_1|\mu_Q|F_2, m_2\rangle B_{-Q}(t) \tag{2.7}$$

This basis change coincides with the notion of polarisations; $\sigma_\pm$-polarised interactions are given by the corresponding basis vectors $e_\pm$, i.e. $Q = \pm1$ and are the case if $m_1 - m_2 = \pm1$, whereas π-polarised interactions have $Q = 0$ and $m_1 - m_2 = 0$.

---

[16] Quantities printed in bold are operators.

$$\langle 3, m_1|\mu_\pm|4, m_2\rangle = \pm\frac{1}{\sqrt{2}}\left\{\mu_B g_J \alpha_{m_1}\beta_{m_2} \mp \mu_N g_I \left(\beta_{m_1}\beta_{m_2}\sqrt{16-m_1^2} + \alpha_{m_1}\alpha_{m_2}\sqrt{15-2m_1-m_1^2}\right)\right\} \tag{2.8}$$

$$\langle 3, m_1|\mu_0|4, m_2\rangle = \frac{1}{2}\mu_B g_J \left(\alpha_{m_1}\alpha_{m_2} - \beta_{m_1}\beta_{m_2}\right) + \mu_N g_I \left\{\alpha_{m_1}\alpha_{m_2}\left(m_1+\frac{1}{2}\right) - \beta_{m_1}\beta_{m_2}\left(m_1-\frac{1}{2}\right)\right\} \tag{2.9}$$

Used here are Clebsch-Gordan expansion coefficients α and β, the Bohr magneton μ_B, the nuclear magneton μ_N, the Landé g-factor $g_J = 2.002\,256\,64(9)$ [57], and the nuclear g-factor $g_I = -1.315\,348(9)$ [p. 158, 23]. The Clebsch-Gordan coefficients are defined as a combination of high field states

$$|F, m_F\rangle = \alpha|I, J, m_F + 1/2, -1/2\rangle + \beta|I, J, m_F - 1/2, 1/2\rangle \tag{2.10}$$

The matrix elements $\langle 3, m_1|\mu_Q|4, m_2\rangle$ resulting from eqs. 2.8 and 2.9 are tabulated in fig. 2.7. Two-qubit gate speeds are linearly proportional to the corresponding matrix element.

State	α	β	State	α	β
			$\|4,-4\rangle$	1.	0.
$\|3,-3\rangle$	0.2880	-0.9576	$\|4,-3\rangle$	0.9576	0.2880
$\|3,-2\rangle$	0.4155	-0.9096	$\|4,-2\rangle$	0.9096	0.4155
$\|3,-1\rangle$	0.5202	-0.8541	$\|4,-1\rangle$	0.8541	0.5202
$\|3,0\rangle$	0.6154	-0.7882	$\|4,0\rangle$	0.7882	0.6154
$\|3,+1\rangle$	0.7071	-0.7071	$\|4,+1\rangle$	0.7071	0.7071
$\|3,+2\rangle$	0.7992	-0.6011	$\|4,+2\rangle$	0.6011	0.7992
$\|3,+3\rangle$	0.8953	-0.4455	$\|4,+3\rangle$	0.4455	0.8953
			$\|4,+4\rangle$	0.	1.

Table 2.3: Clebsch-Gordan expansion coefficients α and β for the $4S_{1/2}$ hyperfine manifold of $^{43}Ca^+$ at a magnetic field of 28.7827 mT [p. 170, 51].

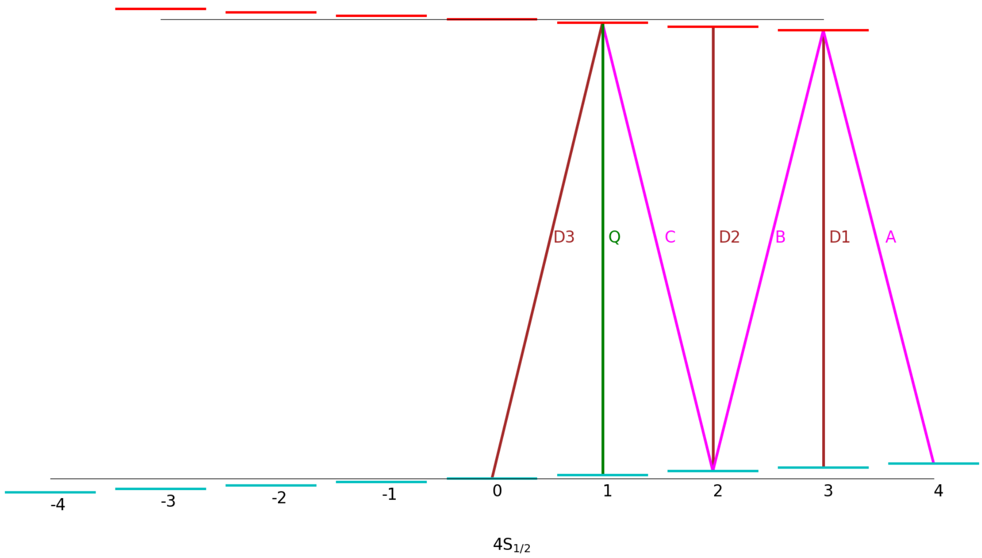

	Transition	Frequency (GHz)	$^{df}/_{dB}$ (MHz/mT)	R (μ_B)
A	$\|4,+4\rangle \leftrightarrow \|3,+3\rangle$	2.546 787 218	-0.2246 218	-1.267 363
D1	$\|4,+3\rangle \leftrightarrow \|3,+3\rangle$	2.674 533 102	-0.1689 743	-0.798 542
B	$\|4,+2\rangle \leftrightarrow \|3,+3\rangle$	2.791 111 000	-0.1233 169	-0.379 132
B'	$\|4,+3\rangle \leftrightarrow \|3,+2\rangle$	2.790 946 119	-0.1233 742	-1.012 843
D2	$\|4,+2\rangle \leftrightarrow \|3,+2\rangle$	2.907 524 018	-0.0777 169	-0.961 755
C	$\|4,+2\rangle \leftrightarrow \|3,+1\rangle$	3.015 270 465	-0.0388 871	-0.799 970
C'	$\|4,+1\rangle \leftrightarrow \|3,+2\rangle$	3.015 435 346	-0.0388 298	+0.601 736
Q	$\|4,+1\rangle \leftrightarrow \|3,+1\rangle$	3.123 181 794	0.	-1.001 026
D3	$\|4,+0\rangle \leftrightarrow \|3,+1\rangle$	3.224 113 796	+0.0340 094	+0.789 034
D3'	$\|4,+1\rangle \leftrightarrow \|3,+0\rangle$	3.223 948 915	+0.0339 521	+0.616 028

Figure 2.7: Energy levels of the $4S_{1/2}$ manifold of $^{43}Ca^+$ at the clock qubit ($|3,+1\rangle \rightarrow |4,+1\rangle$) magnetic field of 28.787 27 mT. The hyperfine splitting is shown on the same scale as the Zeeman splitting. Laser based state-prep leaves the population in $|4,+4\rangle$. To transfer the population to the qubit states, preparation pulses are shown magenta, and labelled A-C indicating the order in which they are applied when transferring the population. The qubit Q itself is shown in green. The other highlighted transitions, shown in brown, are used for diagnostics. The black lines indicated the energy level of the hyperfine manifold in the absence of a magnetic field. The transition frequencies and their linear magnetic field sensitivities are shown in the table, as well as their matrix elements [p. 283, 23]. The qubit has a quadratic magnetic field sensitivity of $^{d^2f}/_{dB^2}$ of $0.025\,14\,^{kHz}/_{\mathrm{mT}^2}$. The transitions marked with an apostrophe are the near-degenerate transition pairs, with C' being of particular importance as explained in the text.

2.8.2 Single-Qubit Gates

After an ion's state is prepared using the SPAM scheme, sect. 2.7, it is in $4\mathrm{S}_{1/2}|4,+4\rangle$. Hereafter, in this section, all states are within $4\mathrm{S}_{1/2}$, so all states are described by $|\mathrm{F},\mathrm{m_F}\rangle$ alone. Transferring this population to one of the qubit states ($|3,+1\rangle, |4,+1\rangle$) is done using three microwave pulses, see fig. 2.7. They transfer the population from $|4,+4\rangle$ to $|3,+3\rangle$, then to $|4,+2\rangle$, and finally to $|3,+1\rangle$. Either of the two-qubit's states can be defined as $|0\rangle$, and $|1\rangle$, although typically the state with higher energy is called $|1\rangle$, i.e. $|1\rangle = |3,+1\rangle$; so to prepare into $|0\rangle = |4,1\rangle$, another microwave pulse is needed. All these pulse times are calibrated to transfer the entire population, i.e. so-called "π-pulses" or pulses of duration t_π[17].

Notably, $\sigma_\pm$-polarisation is needed to transfer population into – and out of – the qubit's states, but within the qubit π-polarisation is used.

For readout, one of the qubit states needs to be mapped back to $|4,+4\rangle$. However, as $|3,+1\rangle \to |4,+2\rangle$ is nearly degenerate to $|4,+1\rangle \to |3,+2\rangle$, simply applying a π-pulse would lead to significant off-resonant excitation errors.

To avoid this, a composite pulse sequence of two $\frac{\pi}{2}$-pulses can be used, with a delay of $\tau \approx {}^1\!/_{2\delta}$ between them, with δ being the frequency splitting between the two transitions [p. 38f, 23]. This composite pulse is then followed by the other two pulses, $|4,+2\rangle \to |3,+3\rangle$ and $|3,+3\rangle \to |4,+4\rangle$. Therefore, the overall mapping is $|3,+1\rangle \to |4,+4\rangle$, while $|4,+1\rangle$ is nominally unaffected.

2.9 Two-qubit gates

Unless otherwise specified, the axes used in this thesis are:

x The trap's axis

y The in-plane axis perpendicular to x

[17] This "π" refers to the duration of the pulse and is not to be confused by the polarisation of the pulses.

z The out-of-plane axis

2.9.1 Lamb Dicke Regime

While this thesis only features work on a single ion, the plan for achieving two-qubit gates is to use the Molmer-Sorensen scheme [40], with some additional modifications [p. 44-57, 23]. These will be explained in the following sections.

The basic idea is to use the two ions' electrostatic repulsion as a bus to transport information between the ions, which are assumed to be in the same potential well. This type of gate assumes that the Lamb Dicke regime applies, where the coupling between the ion's internal and motional states is sufficiently small that the motional state quantum number n is only changed by one at most. This is the case if the interaction between ion and driving radiation gives a potential energy that is a linear function of position (and hence a force that is independent of position). In the case of driving by a standing wave standing wave, the Lamb Dicke parameter η_{sw} is defined as

$$\eta_{sw} = q_0 \frac{2\pi}{\lambda} \cos\theta \tag{2.11}$$

with the spread of the ion zero-point wave function $q_0 = \sqrt{\hbar/2m\omega_j} \sim 10 \ldots 30\,\mathrm{nm}$ (which uses the ion mass $m = 43\,\mathrm{amu}$ and the motional mode frequency $\omega_j \sim 1 \ldots 10\,\mathrm{MHz}$), the wavelength λ, and the angle θ between the laser and the chosen motional mode [107ff, 3].

The Lamb Dicke regime applies if

$$\eta\sqrt{n+1} \ll 1 \tag{2.12}$$

The magnitude of the Lamb Dicke parameter η also has implications on the speed of a two qubit gate. When using an unmodified Molmer-Sorensen scheme, the two-qubit Rabi frequency Ω_{MS} is [p. 2, 40]

$$\Omega_{MS} = -\frac{(\Omega\eta)^2}{2(\omega_j - \delta)} \tag{2.13}$$

using the single qubit Rabi frequency Ω, and the detuning of the field δ as defined in fig. 2.8. From the above equation, it is follows that it is desirable to maximise Lamb Dicke parameter η when trying to increase gate speeds – as long as the Lamb Dicke parameter η does not increase to beyond the Lamb Dicke regime.

At optical wavelengths standing wave Lamb Dicke parameter, can be as large as $\eta_{sw,opt} \sim 0.1$, which creates a reasonably fast gate. However, if one used microwave wavelengths, then the Lamb Dicke parameter decreases to an order of $\eta \sim 10^{-7}$, dropping the gate speed by *12* orders of magnitude. This makes microwave-based gates infeasible when using standing waves.

Fortunately, it is possible to circumvent this problem by placing ions into the vicinity of wires which carry an oscillating microwave current [43]. An effective Lamb Dicke η_{eff} parameter can be defined as the ratio of single qubit Rabi frequency Ω to the two-qubit Rabi frequncy Ω_{MS} [p. 55, 51]:

$$\eta_{eff} = \frac{\Omega_{MS}}{\Omega} \tag{2.14}$$

The length scale of the strength of the interaction is then given by the distance to the current carrying wires, and as a first approximation the following effective Lamb Dicke parameter can be found

$$\eta_{eff} \approx \frac{q_0}{d} = 0.3 \ldots 0.8 \cdot 10^{-3} \tag{2.15}$$

using the ion zero-point wave function q_0 as defined above, and the ion-to-electrode distance $d = 40\,\mu$m used in this experiment. While this Lamb Dicke parameter η_{eff} is still smaller than those that can be achieved using lasers, it is large enough to make microwave based gates feasible.

Using an oscillating magnetic field, the single qubit Rabi frequency is given by

$$\Omega = \frac{B}{2\hbar} R \tag{2.16}$$

while the two qubit Rabi frequency is

$$\Omega_{MS} = \frac{q_0 B'}{2\hbar} R \tag{2.17}$$

using the amplitude of the magnetic field B and its spacial derivate B', as well as the matrix element of the transition R which can be found in 2.7 for the transitions used in this thesis, as well as the reduced Planck constant $\hbar$ [p. 2, 43].

Using equation 2.14, and inserting the above Rabi frequencies, the effective Lamb Dicke parameter is then

$$\eta_{eff} = \frac{q_0 B'}{B} \tag{2.18}$$

Therefore, in order to target a faster gate, one goal of the ion trap design is to maximise the magnetic field gradient B', while minimising the magnetic field amplitude B.

2.9.2 The Two-Qubit Hamiltonian

For two ions in a trap, the overall system Hamiltonian is:

$$H(t) = H_0 + H_{motion} + H_I(t) \tag{2.19}$$

Here, H_0 is the internal energy of the ions, H_{motion} is the motion of the ions, and $H_I(t)$ is the interaction with an applied microwave field. The following derivation follows [p. 48-61, 51], where additional details can be found. The two ions are assumed to only have two internal states – the qubit states $4\mathrm{S}_{1/2}|3, +1\rangle =$

$|1\rangle$ and $4S_{1/2}|4,+1\rangle = |0\rangle$ apart from a shared motion. This means the internal energy H_0 is

$$H_0 = \frac{1}{2}\hbar\omega_0\sigma_{z,1} + \frac{1}{2}\hbar\omega_0\sigma_{z,2} \tag{2.20}$$

with the reduced Planck constant $\hbar$, the qubit transition frequency $\omega_0 = f_{1\to 1} = 3\,123.181\,794\,\mathrm{MHz}$, and the Pauli matrix $\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ acting on the the ion denoted in their index. The Hamiltonian describing the motion of the ions is

$$H_{motion} = \sum_{l=\{\tilde{x},\tilde{y},\tilde{z}\}} \sum_{j=1}^{N_l} \hbar\omega_{l,j} \left(\frac{1}{2} + a^{\dagger}_{l,j}a_{l,j}\right) \tag{2.21}$$

with the three axes l of the motional modes[18], particular motional mode j of the total number of motional modes N_l within that axis l, the frequency $\omega_{l,j} \sim 1\ldots 10\,\mathrm{MHz}$ of a particular motional mode, as well as the creation and annihilation operators $a^{\dagger}_{l,j}$ and $a_{l,j}$ acting on a particular motional mode. The position operator $\boldsymbol{r}_n$ of a particular ion $n = 1, 2$ is

$$\boldsymbol{r}_n = \boldsymbol{r}_{0,n} + \sum_{l=\{\tilde{x},\tilde{y},\tilde{z}\}} \sum_{j=1}^{N_l} q_l b_n^{l,j} q_0 \left(a_{l,j} + a^{\dagger}_{l,j}\right) \tag{2.22}$$

using the ion's equilibrium location $\boldsymbol{r}_{0,n}$, the ion zero-point wave function q_0 as defined in sect. 2.9.1, the unit vector q_l based on the tilted coordinate system, and a displacement $b_n^{l,j}$. This displacement can indicate the centre of mass modes, $b^{l,1} = \frac{1}{\sqrt{2}}(1,1)$, or the breathing mode along $\tilde{x}$ and rocking mode along $\tilde{y}$ and $\tilde{z}$, $b^{l,2} = \frac{1}{\sqrt{2}}(-1,1)$.

Finally, a drive field $\boldsymbol{B}(r,t)$ is applied, which couples to the ions' dipole moments $\boldsymbol{\mu}$:

[18]As described in 2.9.4, $\tilde{y}$ and $\tilde{z}$ are tilted by the motional mode angle θ with respect to the axes y, and z, while $\tilde{x} = x$

$$H_I(t) = -\boldsymbol{\mu} \cdot \boldsymbol{B}(r_1)\cos(\omega t - \phi) - \boldsymbol{\mu} \cdot \boldsymbol{B}(r_2)\cos(\omega t - \phi) \tag{2.23}$$

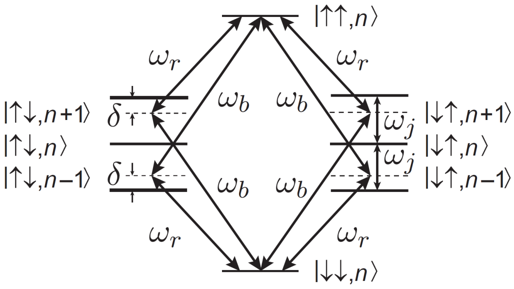

Figure 2.8: A diagram of a Molmer-Sorensen gate, showing the levels, microwave fields, and their detunings. In this case, the qubit states are $4S_{1/2}|4,+1\rangle = |\downarrow\rangle = |0\rangle$ and $4S_{1/2}|3,+1\rangle = |\uparrow\rangle = |1\rangle$, which are together with the motional quantum state $|n\rangle$ give a combined state description of e.g. $|00,n\rangle$. The motional quantum states are spaced apart by the motional frequency ω_j. The applied red and blue sidebands are ω_r and ω_b, as explained in the text. Taken from [fig. 6.13, 3][fig. 3, 49].

using the magnetic field frequency ω – which is tuned close to the qubit transition frequency $f_{1\to 1} = 3\,123.181\,794\,\mathrm{MHz}$ – and applied with a phase ϕ.

The magnetic field $\boldsymbol{B}(r,t)$ can be expressed as

$$\boldsymbol{B}(r,t) = \sum_{Q=\{\pm 1,0\}} (-1)^Q B_Q(r,t)\, e_{-Q} \tag{2.24}$$

using the spherical basis unit vectors e_Q as defined in sect. 2.8.1, the polarisation Q, and the magnetic field components $B_Q(r,t)$. Assuming we are in the Lamb Dicke regime, the following linear approximation of the magnetic field is valid

$$\boldsymbol{B}(r,t) = \sum_{Q=\{\pm 1,0\}} (-1)^Q \left\{ B_Q + \sum_{l=\{\tilde{y},\tilde{z}\}} \sum_{j=1}^{N_l} \frac{\partial B_Q}{\partial q_l} b_n^{l,j} q_l \left(a_{l,j} + a_{l,j}^\dagger \right) \right\} e_{-Q} \tag{2.25}$$

The interaction Hamiltonian for ion n can now be rewritten as

$$\begin{gathered} H_{I,n}(t) = (-1)^{Q+1} B_{-Q} \langle 1|\mu_Q|0\rangle \cos(\omega t - \phi)\, \sigma_{x,n} - \\ \tfrac{1}{2} B_0 \left(\langle 1|\mu_0|1\rangle - \langle 0|\mu_0|0\rangle\right) \cos(\omega t - \phi)\, \sigma_{z,n} \end{gathered} \tag{2.26}$$

It follows from equ. 2.9 and tab. 2.3, that $\langle 1|\mu_0|1\rangle = \langle 0|\mu_0|0\rangle = 0$, as is the case for all FOFI qubits. Therefore, the second term in $H_{I,n}(t)$ is zero, so

$$H_{I,n}(t) = (-1)^{Q+1} B_{-Q} \langle 1|\mu_Q|0\rangle \cos(\omega t - \phi)\, \sigma_{x,n} \tag{2.27}$$

Inserting the magnetic field $\boldsymbol{B}(r,t)$ into this results in

$$H_{I,n}(t) = -\left\{ B_{-Q} + \sum_{l=\{\tilde{y},\tilde{z}\}} \sum_{j=1}^{N_l} \frac{\partial B_{-Q}}{\partial q_l} b_n^{l,j} q_l \left(a_{l,j} + a_{l,j}^\dagger\right) \right\} \langle 1|\mu_Q|0\rangle \cos(\omega t - \phi)\, \sigma_{x,n} \tag{2.28}$$

To get to the effect of the this field on the motion and internal states, an interaction picture between $H_0 + H_{motion}$ is used and a rotating wave approximation is made, resulting in the interaction picture Hamiltonian $\check{H}_{I,n}(t)$

$$\begin{gathered} \check{H}_{I,n}(t) = \tfrac{1}{2}\hbar\Omega_Q e^{-i((\omega-\omega_0)t-\phi)} \sigma_{+,n} + \\ + \textstyle\sum_{l=\{\tilde{y},\tilde{z}\}} \sum_{j=1}^{N_l} \tfrac{1}{2}\hbar\Omega_{MS,Q,l,j,n,} e^{-i((\omega-\omega_0)t-\phi)} \sigma_{+,n} \left(a_{l,j} e^{-it\omega_{l,j}} + a_{l,j}^\dagger e^{it\omega_{l,j}}\right) + H.C. \end{gathered} \tag{2.29}$$

With the following Rabi frequencies

$$\Omega_Q = -\frac{1}{\hbar} (-1)^Q \langle 1|\mu_Q|0\rangle \tag{2.30}$$

$$\Omega_{Q,l,j,n} = -\frac{1}{\hbar}(-1)^Q \langle 1|\mu_Q|0\rangle \frac{\partial B_{-Q}}{\partial q_l} b_n^{l,j} q_{l,j} \tag{2.31}$$

which were already used at the end of the previous section. This interaction Hamiltonian $\check{H}_{I,n}(t)$ can be used to create the gate by applying using two frequency components, $\omega_r = \omega_0 - \omega_j - \delta$ and $\omega_b = \omega_0 + \omega_j + \delta$, see fig. 2.8. The

Molmer Sorensen Hamiltonian is then $\check{H}_{MS,n}(t) = \check{H}_{r,n}(t) + \check{H}_{b,n}(t)$, applying the red and blue sidebands, respectively.

Note that there is an intentional detuning δ from the nominal resonance $\omega = \omega_0 \pm \omega_j$. The gate is applied for a time $t = \frac{2\pi K}{\delta}$. The detuning is chosen such that $\delta = 2\Omega\sqrt{K}$, with the two-ion Rabi frequency Ω , and an integer $K \geq 1$, which determines the number of loops through phase space[19]. In the absence of errors, this gate transiently couples the motional state $|n\rangle$ with the ion's internal state, e.g. $|00\rangle$, but fully decouples the two at the end of the gate. The overall effect of the gate is

$$\begin{aligned} |00\rangle &\rightarrow \tfrac{1+i}{2}\left(e^{2i\phi_{avg}}|11\rangle - i|00\rangle\right) \\ |01\rangle &\rightarrow \tfrac{1+i}{2}\left(|10\rangle - i|01\rangle\right) \\ |10\rangle &\rightarrow \tfrac{1-i}{2}\left(|10\rangle + i|01\rangle\right) \\ |11\rangle &\rightarrow \tfrac{1-i}{2}\left(|00\rangle + e^{-2i\phi_{avg}}i|11\rangle\right) \end{aligned} \tag{2.32}$$

using the average phase of the red and blue sidebands[20] $\phi_{avg} = \frac{1}{2}(\phi_r + \phi_b)$. This two-qubit gate forms, together with the single qubit gates, a universal set of gates.

The highest fidelity achieved for microwave based two-qubit gates achieved a Rabi frequency of $\Omega \approx 2\pi \cdot 300\,\mathrm{Hz}$ [p. 4 , 26]; the goal of this experiment is to achieve a two-ion Rabi frequency $\sim$ 10 times as high, i.e. $\Omega \approx 2\pi \cdot 3\,\mathrm{kHz}$.

One feature of this scheme is that it does not require ground state cooling – a thermal state of motion is in principle sufficient – so this type of gate is more resistant to a non-zero motional heating rate $\dot{n}$ than other gates.

2.9.3 Two-qubit gate errors

So far, no sources of errors were considered affecting the fidelity of the gate. Previous work considered the following sources of errors the most relevant [p.

[19] All else being equal, higher number of loops lead to a slower gate, but one that is more resistant to certain types of errors.

[20] Note that despite the parameter $e^{\pm 2i\phi_{avg}}$ affecting only a single state ensemble $|11\rangle$, the overall state $\alpha|11\rangle + \beta|11\rangle$ is still normalised.

64-76, 51]:

- Rabi frequency Ω fluctuations, typically caused by changes in the sideband power. Previous work achieved a Rabi frequency stability "at the sub-Hertz level", leading to a negligible error of $1 \cdot 10^{-6}$.

- Timing miscalibrations such that $t \neq \frac{2\pi K}{\delta}$, causing some entanglement between the internal and motional states of the ions to remain. Previous work attributed an two-qubit gate error of 0.2 (1) % to a fractional offset of t from its nominal value by 0.021 (4).

- Differences in the Rabi frequencies between the two sidebands, Ω_r and Ω_b, caused by unbalanced microwave power. Previous work achieved a fractional difference of $3.4\,(6) \cdot 10^{-3}$, leading to a negligible error of $2.8\,(3) \cdot 10^{-7}$ after applying error suppression techniques.

- Off-resonant excitation of transitions outside those used for the qubit. These can be partially counteracted by pulse-shaping and spectral filtering of the microwave field. Additionally, the higher static magnetic field used in this work as compared to the previous work suppresses off-resonant effects.

- Off-resonant excitation leading to AC-Zeeman shifts. The static part of this effect can be calibrated for, but fluctuations in the AC-Zeeman shift need the application of error suppression techniques.

- Motional heating. This can be counteracted by using cryogenics.

- Motional decoherence, for example caused by fluctuations in the RF drive power.

Previous work used a modified Molmer Sorensen scheme with $K = 4$ loops through phase space [26], as a technique to suppress errors. The same scheme is planned to be used for this work, so the modifications are explained as follows:

The first modification of the scheme is a π-pulse applied on each ion halfway through the gate, i.e. after $K = 2$ loops. In an error-free Molmer Sorensen gate, this does not affect the states. However, small perturbations in the Hamiltonian from the ideal case lead to the loops not closing, i.e. the motion of the ions and their internal states are not fully disentangled. This leads to errors in the gate. The π-pulse suppresses the effect of typical perturbations that are varying slowly compared to the gate time. Note that this only works for an even number of loops, i.e. a minimum of 2 loops is needed.

The second modification of the scheme is related to imperfect nulling of the relevant magnetic field polarization. This σ-polarised field will off-resonantly drive transitions from the $4\mathrm{S}_{1/2}|4,+1\rangle \leftrightarrow 4\mathrm{S}_{1/2}|3,+1\rangle$ qubit to the neighbouring states, leading to population leakage away from the qubit, and therefore errors in the gates. Pulse shaping techniques were used to highly suppress these off-resonant Rabi-flops, but an a.c. Zeeman shift Δ is induced on the qubit transition due to this field. This frequency Δ is experimentally calibrated between gates. However, variations in Δ that are fast compared with the calibration cycles, Δ', remain, leading to errors. By applying an additional frequency component in between the red and blue sidebands, at the carrier frequency $\Omega_c = \omega_0 + \Delta$, single qubit rotations are continuously driven during the gate's operation. Without the first scheme modification, the gate time t_g would need to follow $\Omega_c t_g = n \cdot 2\pi$ to complete the single qubit rotations. However, the π-pulse in between the loops removes this constraint. If $\Omega_c \gg \Omega, \Delta'$, than this further suppresses errors; the carrier power is chosen such that this condition holds. Typically this power is then much lower than the sidebands, e.g. 2 W per sideband and 3 μW for the carrier [p. 4, 26].

After these modifications – and many other measures taken – the remaining errors were [p. 4, 26]:

$\leq 0.2\%$ due to the motional mode heating rate of $\dot{\bar{n}} = 5\,\mathrm{s}^{-1}$

$\sim 0.2\%$ due to fluctuations in the motional mode frequency of $\sim 30\,\mathrm{Hz\,rms}$

$\leq 0.1\%$ due to off-resonant excitation and a.c. Zeeman shifts

$\leq 0.1\%$ due to systematic miscalibration in the sideband Rabi frequencies or gate time

The overall two-qubit gate fidelity was measured as $99.7\,(1)\,\%$. An additional target for this experiment is to achieve a factor of ~ 10 lowering of the two-qubit gate error rate. To this end, several factors are used:

1. Compared to the above, the motional mode heating rate $\dot{\bar{n}}$ is expected to increase by a factor of $\left(\frac{75\,\mu\mathrm{m}}{40\,\mu\mathrm{m}}\right)^4 = 12$ due to the lower ion height, but decrease by a factor of ~ 100 due to the employed cryogenic cooling of the trap surface. However, as the previous experiment had atypically low heating rates, a return to the norm would suggest a roughly equal heating rate. Assuming this is the case, the speed improvement of a factor of ~ 10, discussed in sect. 2.9.4, should reduce the error associated with the heating rate by a factor of ~ 10.

2. The motional mode frequency fluctuations are measured in sect. 5.4, placing an upper bound of $2\pi \cdot 1\,\mathrm{kHz}$ on the fluctuations, limited by measurement noise. Assuming this upper bound, this would be a severe limitation on the two-qubit gate errors. An upgrade in the RF drive chain to increase the stability of the RF amplitude, and therefore of the secular frequencies, is already under way, see sect. 5.4.

3. The experimental control system, see sect. 3.11, offers more flexibility and ease-of-use than the previous system. Calibrations can be done more easily and more often. This should result in a lowering of systematic miscalibrations.

4. The FOFI qubit employed in this experiment necessitates a higher magnetic field than the previous work; this increases the Zeeman splitting between neighbouring states from $\approx 50\,\mathrm{MHz}$ to $\approx 100\,\mathrm{MHz}$, as shown in fig. 2.7. The transitions D3, D3', C, and C' – as defined in fig. 2.7 – are σ-polarised spectator transitions to the π-polarised qubit transition Q, leading to errors due to off-resonant excitation. The factor of 2 increase in the splitting between Q and the spectator transitions, increases the on-to-off resonant excitation ratio by a factor of 4. Additionally, the π-polarised microwave magnetic field now couples better to the π-polarised qubit, while the coupling worse to the σ-polarised off-resonant transitions, each increasing the on-to-off resonant excitation by a factor of 2. Finally, an analogous argument with the matrix elements yields another factor of ≈ 4. Overall, this should decrease off-resonant excitation by a factor of ≈ 64. However, the residual, non-nulled, microwave magnetic field component could be higher in this experiment, so that some of this reduction in off-resonant excitation could be counteracted by this increase. The net effect should still be an order of magnitude decrease in the off-resonant excitation.

2.9.4 Ion Trap Chip Design Considerations

The chosen qubit $4\mathrm{S}_{1/2}|4,+1\rangle \leftrightarrow 4\mathrm{S}_{1/2}|3,+1\rangle$ uses a π-polarised transition. To maximise the effective Lamb Dicke parameter η_{eff}, the amplitude of the microwave magnetic field in the π-polarisation, $B_{MW,\pi}$, needs to be minimised, while its gradient, $B'_{MW,\pi}$, needs to be maximised. This is the goal of the ion trap chip design, see sect. 3.2. Due to the high Zeeman splitting of $\sim 100\,\mathrm{MHz}$, the σ-polarised magnetic field, $B_{MW,\sigma}$, is far enough off-resonant that any residual errors can be made small by using the dynamical decoupling scheme introduced in sect. 2.9.3. Crucially, this means that the σ-polarised magnetic field , $B_{MW,\sigma}$, does *not* need to be minimised, facilitating the ion trap chip geometry.

To create the specified field, a two wire design – parallel to the x axis – is chosen, bent into a "U" shape, see fig. 3.4. Therefore, the same microwave current flows in both wires, but in opposite directions. This cancels out the $B_{MW,y}$ component in the symmetry (xz) plane between the wires.

Along with additional considerations relating to the beam paths, the static magnetic field $B \approx 28.8\,\mathrm{mT}$ is chosen to be parallel to the surface to allow for σ-polarised laser beams, which are used for SPAM, see sect. 2.7. These considerations then also result in $B_{MW,y} = B_{MW,\pi}$, i.e. the static magnetic field being parallel to the y-axis.

To achieve a higher coupling, the motional mode used for the gate also needs to be in the direction of the magnetic field gradient $B'_{MW,\pi}$. This is maximised for a motional mode angle of $\theta = 0°$. However, some non-zero motional mode angle θ is needed for a single in-plane laser beam path to be able to cool all three motional mode axes. The motional mode angle θ is therefore a trade-off between gate speed and cooling efficiency.

As already mentioned, the aim of this experiment is an order-of-magnitude improvement over the current highest fidelity microwave based two-qubit gate, in both gate speed and fidelity. To achieve this order-of-magnitude speed increase, the following sources of speed improvements are drawn from:

1. The previous experiment used the $|4,+0\rangle \leftrightarrow |3,+1\rangle$ transition as a qubit, with a matrix element of $R_{01}\,(14.6\,\mathrm{mT}) = -0.702\,127$ [p. 282, 23], while this experiment uses the $|4,+1\rangle \leftrightarrow |3,+1\rangle$ transition, with a matrix element $R_{11}\,(28.8\,\mathrm{mT}) = -1.001\,026$. The relative strengths gives a gate speed increase of a factor of $\frac{R_{11}(28.8\,\mathrm{mT})}{R_{01}(14.6\,\mathrm{mT})} = 1.425\,705$, i.e. an increase of 43%.

2. The geometry of this ion trap – and that of the previous experiment [p. 16, 51] – produces a linearly polarised magnetic field. In the previous experiment, the σ-polarised transition $4\mathrm{S}_{1/2}|4,0\rangle \leftrightarrow 4\mathrm{S}_{1/2}|3,+1\rangle$ was therefore slowed down by a factor of 2 in power and $\sqrt{2}$ in gate speed. With this

experiment's π-polarised qubit $4S_{1/2}|4,+1\rangle \leftrightarrow 4S_{1/2}|3,+1\rangle$, this is no longer the case, so a factor of $\sqrt{2}$ increase in gate speed can be anticipated.

3. The secular frequency used for two-qubit gate in the previous experiment was 3 MHz, while this experiment is planning to use ~ 6 MHz. This leads to a decrease of a factor of $\sqrt{2}$ in two-qubit gate speed. This is done to reduce the heating rate of the ions' motional mode $\dot{\bar{n}}$ as defined by $\dot{\bar{n}} = \frac{q^2}{4m\hbar\omega} S_E(\omega)$, using the ion's charge q and mass m, the trap frequency $f = {}^{\omega}/_{2\pi}$, and the electric-field noise spectral density $S_E(\omega)$ [p. 37f, 9]. It also increases the trap depth, which tends to decrease with the size of the trap. The scaling of $S_E(\omega)$ is typically $S_E(\omega) \propto \omega^{-1 \ldots -1.5}$, resulting in a two-qubit gate error $\propto \omega^{-2 \ldots 2.5}$, while ignoring effects on the gate speed. However, as the Lamb Dicke parameter η also has a dependence of $\eta \propto \omega^{-0.5}$, the two-qubit gate error scales roughly as $\propto \omega^{-1.5}$. Therefore, doubling the secular frequency should result in a factor of $2^{1.5} \approx 2.8$ reduction in the heating rate $\dot{\bar{n}}$. However, the trade-off between gate speed and gate fidelity when choosing the secular frequency will be revisited once the heating rate $\dot{\bar{n}}$ of the system is measured.

4. The previous experiment used a motional mode angle θ of $40°$ [p. 48, 51], while this experiment is planning to use $\sim 15°$. The effect on the two-qubit gate speed is a factor of $\frac{\cos(15°)}{\cos(40°)} = 1.26$.

The above factors increase the two-qubit gate speed by a well-defined amount, a factor of $1.43 \cdot \sqrt{2} \cdot \frac{1}{\sqrt{2}} \cdot 1.26 = 1.8$. Increasing the magnetic field gradient $B'_{MW,\pi}$ is planned to be used to further increase the speed:

1. By decreasing the distance between the ion and the the microwave electrodes increases the magnetic field gradient $B'_{MW,\pi}$. For an infinitely long, straight, current carrying wire, the produced magnetic field is $B = \frac{\mu_0 I}{2\pi d}$, using the current I, the vacuum permeability μ_0, and the ion to surface

distance d. This means that $B \propto d^{-1}$, i.e. by decreasing the ion height from 75 μm to 40 μm, there is an anticipated field strength increase of a factor of ~ 2. However, this is only a very rough first approximation, as the exact geometry of the wires and resonator on the chip modifies this value. A comparative design study between the different trap geometries and their respective magnetic field properties has not been undertaken.

2. Both the ion trap in the previous experiment, and in this one are Au plated. However, as this trap will be operated at cryogenic temperatures, the resistance of the Au drops by a factor of ~ 10 (depending on the purity). The exact factor of resistance drop is hard to predict however. The same microwave current will therefore dissipate less power, but the dissipation of microwaves in the ion trap system is not as straightforward, and would need to be simulated. While this has not been undertaken, the amount of reduction in dissipated power P is unknown, but estimated to give rise to a small reduction in dissipated power P. In the case of a fixed limit on the dissipated power P, this increases the limit on the microwave current I. This increases the magnetic field gradient $B'_{MW,\pi}$, leading to a faster gate.

3. Advances in heat sinking capabilities, see sect. 4.2.2, should increase the power dissipation capabilities of the ion trap chip over the previous experiment, allowing for a higher current I to be applied, and thus magnetic field gradient $B'_{MW,\pi}$ to be created.

4. Wafer scale production of ion trap chips, see sect. 4.1.2, mean that an ion trap damaged by overheating can be replaced more readily and reliably than with ion traps produced one at a time. This means that a higher current I can be applied, even at the risk of damaging the trap, as it is easier to replace; it is generally not known what the current limit I_{lim} of a particular trap is. One could even deliberately apply a slow current ramp, reaching a

high enough current to destroy the trap to establish the current limit I_{lim}; the replacement trap would then be used at a current suitably lower than I_{lim}.

Therefore, the lower bound of the overall speed increase is a factor of $1.43 \cdot \sqrt{2} \cdot \frac{1}{\sqrt{2}} \cdot 1.26 = 1.8$, from the first four sources, respectively. The use of cryogenic Au as a conductor increases the current for a given power, so this will give some contribution to the speed increase. Additionally, a future measurement of the motional heating rate $\dot{\bar{n}}$ will influence the final choice of the secular frequency f, potentially changing the speed further. On top of that, there is an unknown speed-up from the change in ion trap design. Also, a higher power could be allowed to dissipate.

It remains to be seen how much of an two-qubit gate speed improvement will be possible with this experiment. However, the results achieved to date suggest the target speed increase could be possible.

Chapter 3

Design

3.1 Overall Design

In this chapter, the parts of the experiment are introduced. Fig. 3.1 shows a schematic overview of the whole experiment; the individual parts will be explained in the following sections. Please note that this schematic only provides a snapshot, and parts of the experiment are changed and updated regularly.

A cut-out rendering of the fully assembled vacuum system placed on an optics table, model RS 4000 with S-2000A isolators both from Newport, can be seen in fig. 3.2. The cryostat and the liquid helium dewar are connected via a flexible transfer line. Due to the length of the cryostat, there is a large leverage arm that the end of the transfer line has onto both the vacuum system, and onto the braid. The concerns are that the former could spring a leak, and the latter could move the pillbox and ion in relation to the lasers and imaging system. Due to these concerns, the optics table is not floated; in contrast to laser based gates, on-chip microwave based ones are much less prone to vibration induced errors.

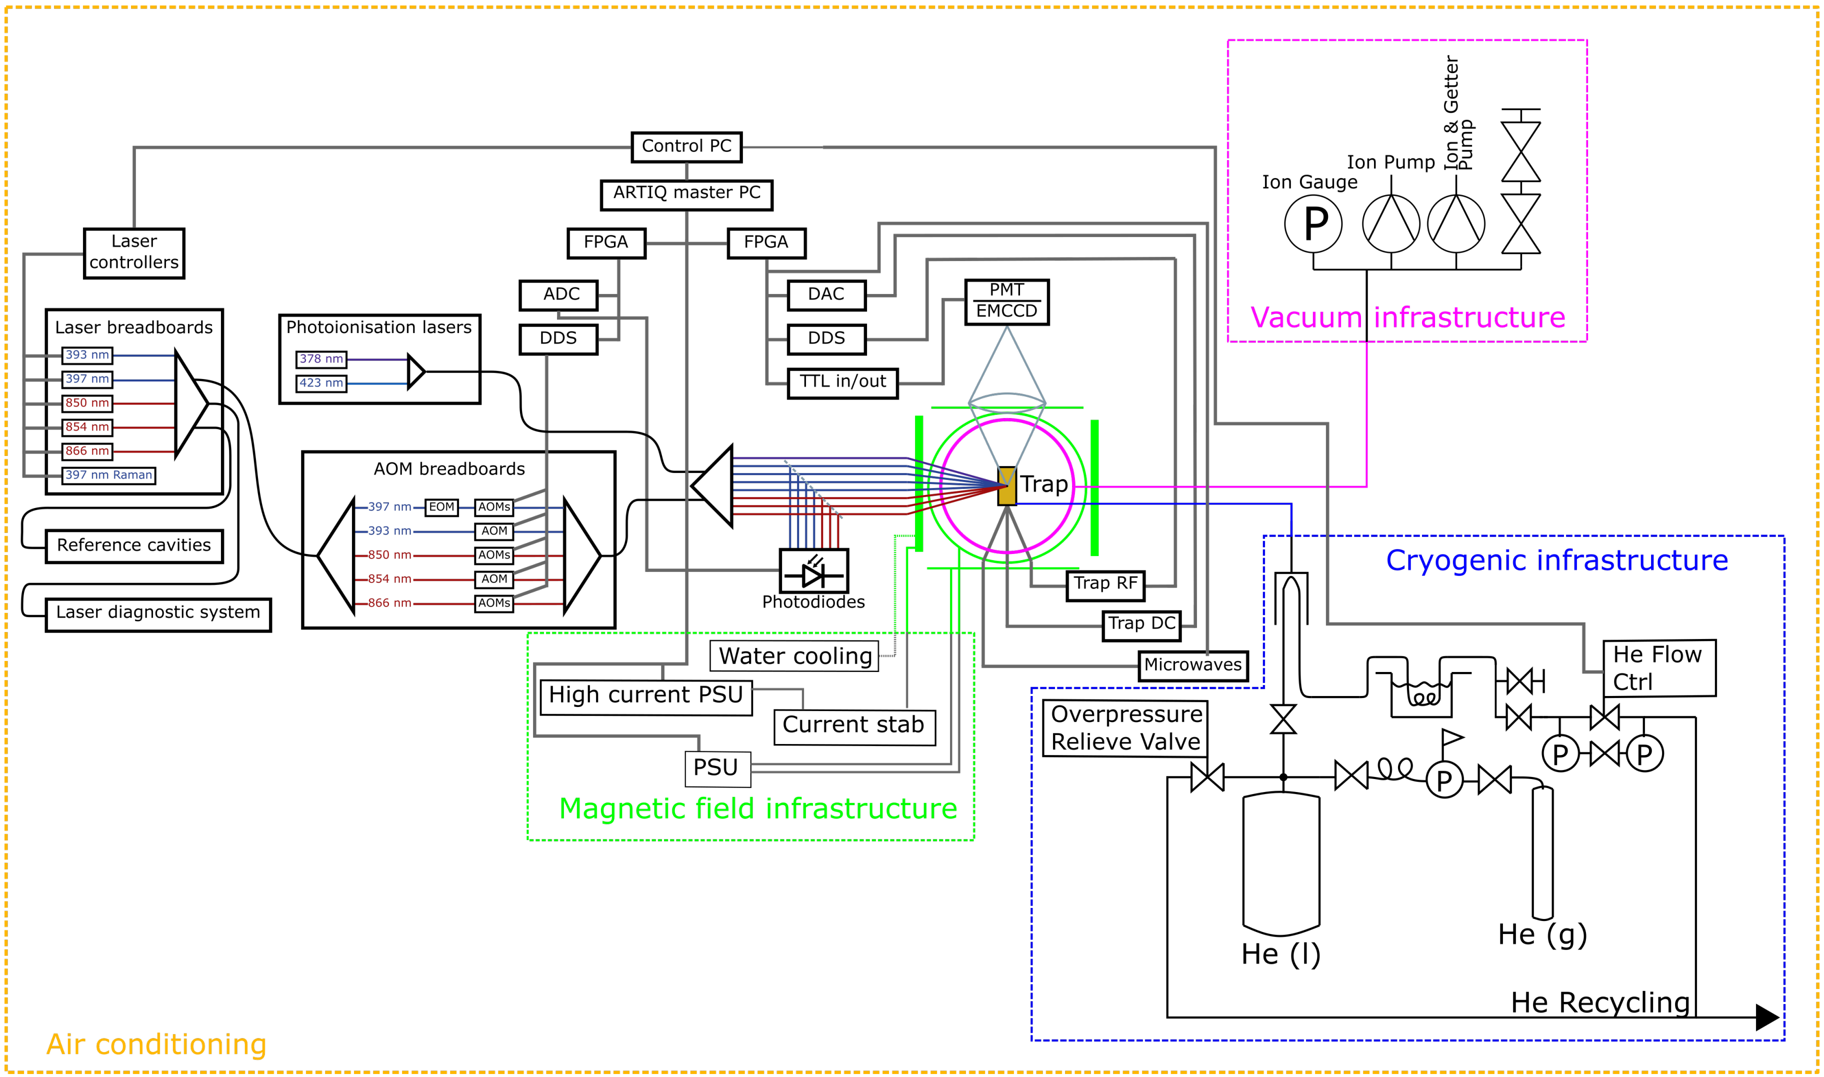

Figure 3.1: Schematic overview of the whole experiment. The experiment control system, see sect. 3.11, can control all parts of the experiment, with the exception of the vacuum system. The air conditioning is also highlighted, as its cycle affects everything else shown; in particular the laser fibres show the same periodic behaviour in polarisation drifts. Modified from [36].

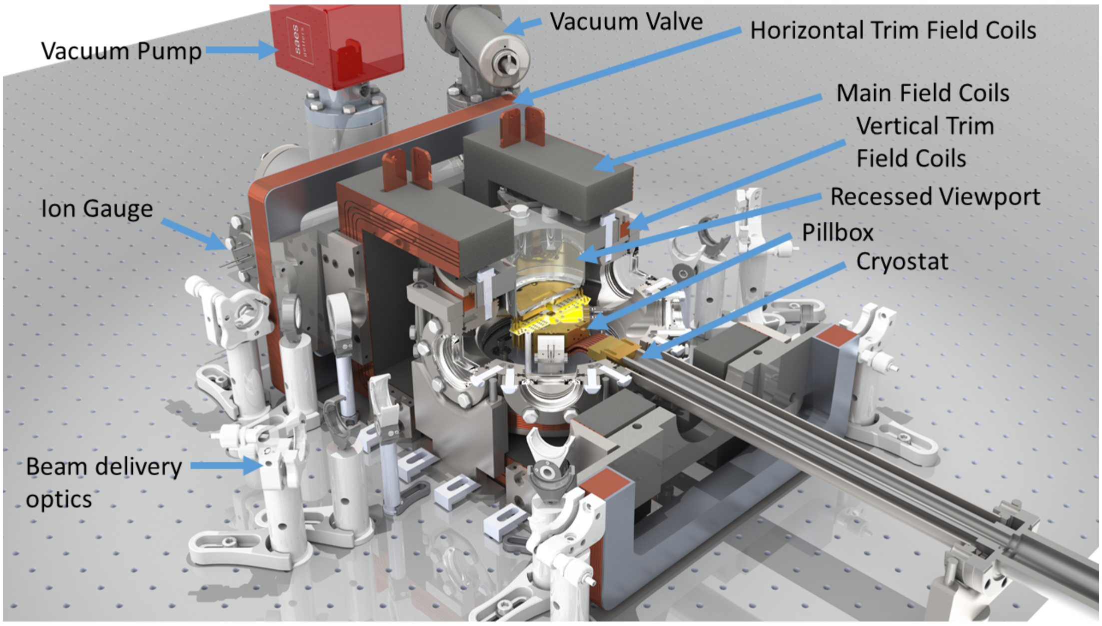

(a)

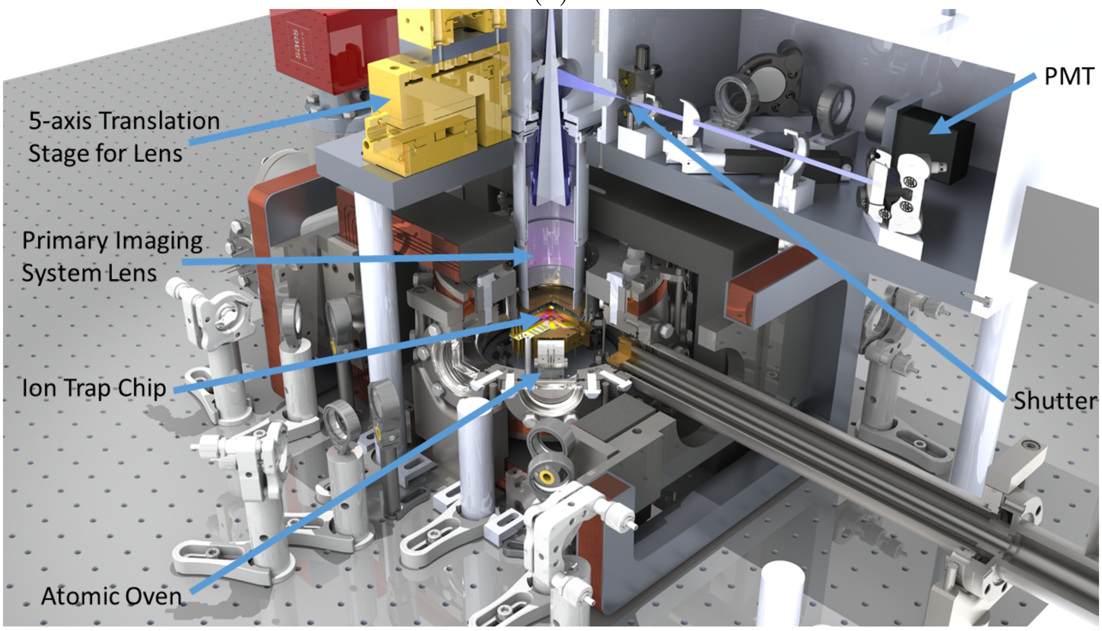

(b)

Figure 3.2: Overview of the whole vacuum system described in this chapter. (a) Seen here is a 1/4cut-out of the vacuum system, including the nominal positions of the surrounding optics, and all three axes of the magnetic field coils. (b) As (a), but with the addition of the imaging system, and with only 1/8-cutout. Rendered in Solidworks Professional 2016.

3.2 Ion Trap Design

The design used for trapping ions is introduced in this section. The design work for the ion trap itself has been done by Dr. Thomas Harty, partially based on previous work [p.118-142, 55]. A model of the trap, and its immediate surroundings, was created using the finite element analysis (FEA) software HFSS from ANSYS Inc. A brief summary of the ion trap design is given in fig. 3.3.

There are 22 DC electrodes in total, 11 per side. They are arranged into three trapping zones, with three DC electrodes per side and zone - except for the central zone which features its central DC electrode further divided into three smaller ones intended to facilitate splitting ion chains. Ions are trapped along the symmetry axis half way between the DC electrodes, see fig. 3.4.

3.2.1 Trapping Fields

For trapping, the “tuning fork” shape in fig. 3.4 is used to create a Paul trap, using frequencies in the radio frequency (RF) range. The RF and DC electrodes are modelled using analytical expressions for “five wire” configurations, the House model [29]. The “wires” are DC-RF-GND-RF-DC, and shown horizontally in fig. 3.4c. The DC “wires” are sectioned into the electrodes introduced above. The central ground (GND) is subdivided into GND-MW-GND, which is not modelled for the RF trapping fields; the GND between the RF and microwave (MW) lines is there to reduce the coupling between them. Notable assumptions for the model as used in this thesis are:

- The electric potential is modelled by a second order Taylor series, i.e. higher order terms are assumed to be small enough to be addressed by a stray field offset term, the micromotion compensation field.
- The effect of the RF potential on the ion can be approximated by a time-averaged pseudo-potential.

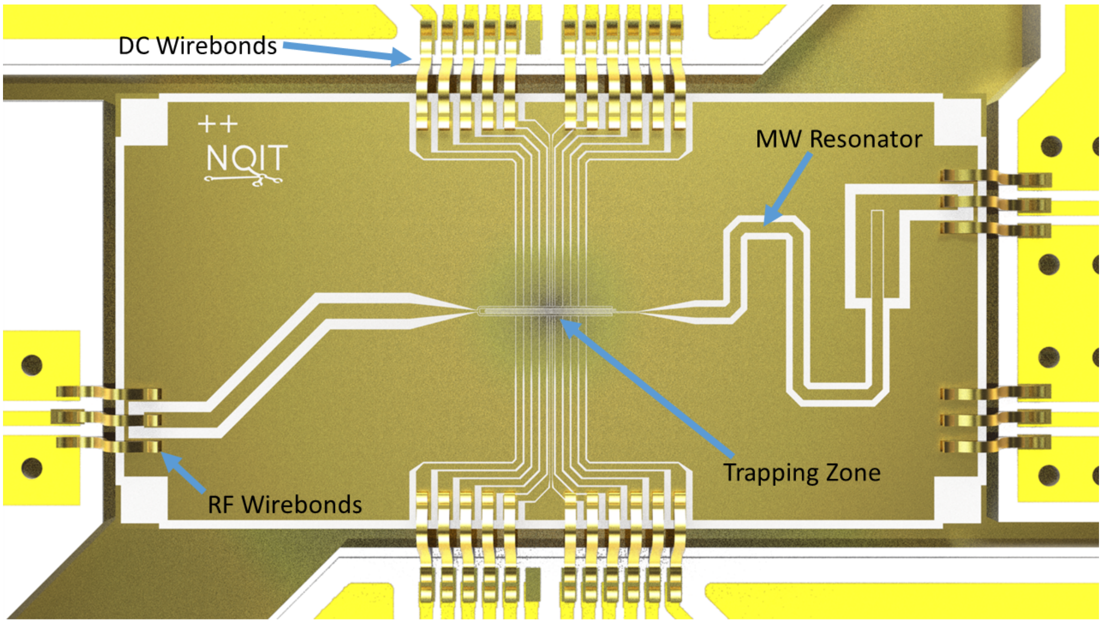

Figure 3.3: Top down view of the ion trap, wire bonded onto the surrounding PCBs. The chip is 10 mm by 5 mm in size. The geometric centre of the image is the ≈ 1 mm long horizontal region where ions can be trapped. Above and below that region, the DC pads can be seen, used for moving the ions along the region, and for potentially swapping their order. To the left, the RF connection is made, creating the pseudo potential Paul trap. To the right, the microwave resonator is located, bent into a serpentine shape to save space. It is coupled to the input microwaves via the small gap between them. The resonator length dictates its resonant frequency. As some discrepancy between the model and measurements is expected, a series of 8 different resonator lengths is fabricated; the one with the closest resonator frequency to the qubit resonance is used for trapping. The individual resonator lengths are marked by a code, going from the shortest, "- - - -", to longest, "+ + + +", resonator. The code is located in the top left corner, "+ +" for this trap, next to the logo of the major funding body for this thesis, NQIT. Rendered with Solidworks 2016 Professional.

- The model assumes all electrodes are on an infinite plane, gaps are neglected[1].
- Except for the DC and RF electrodes, the rest of the plane is assumed to be grounded.
- The DC electrodes are modelled as semi-infinite rectangles, extending outwards from the trap.
- The RF electrodes are modelled as infinite rectangles of equal width.
- The design is symmetric around a symmetry plane half way between the RF electrodes.

The RF trap depth $\psi_{E,RF}$ can be calculated as follows [p. 5, 29]

$$\psi_{E,RF} = \frac{Z^2 e^2 V_{RF}^2}{\pi^2 m \Omega^2} \left[\frac{b}{(a+b)^2 + (a+b)\sqrt{2ab + a^2}} \right]^2 = 34\,\mathrm{meV} \tag{3.1}$$

using the ion's charge $Z = +1$ and mass $m = 43\,\mathrm{amu}$, the elementary charge e, the trapping RF amplitude $V_{RF} = (90 \pm 10)$ V and frequency $\Omega = 2\pi \cdot 63.98\,\mathrm{MHz}$, the width of central GND $a = 60\,\mu\mathrm{m}$ and width of the RF electrodes $b = 23.6\,\mu\mathrm{m}$. The latter two are measured 4.3 values using the above gap-less approximation. The anti-trapping due to the DC potential $\psi_{E,DC}$ will make the total trap depth $\psi_E = \psi_{E,RF} + \psi_{E,DC}$ smaller, so the above RF trap depth $\psi_{E,RF} = 34\,\mathrm{meV}$ is an upper limit.

Due to the RF, the ion is trapped in the symmetry plane between the RF electrodes. Its height above the electrodes can be calculated as [p. 4, 29]

$$z_0 = \frac{\sqrt{a\,(a+2b)}}{2} = 40.1\,\mu\mathrm{m}. \tag{3.2}$$

The escape path of the ion is straight upwards, and it escapes the trap at a height above the electrodes z_E of [p. 5, 29]

[1] Electrodes are symmetrically increased in size until they touch, i.e. half of each gap is modelled as belonging to electrode of either side of the gap.

$$z_E = \sqrt{a^2 + 2ab + 2\,(a+b)\sqrt{\frac{a^2}{2} + 2ab}} = 133\,\mu\mathrm{m}. \tag{3.3}$$

The trap stability parameter q can be calculated to be [p. 5, 29]

$$q = -\frac{16ZeV_{RF}}{\pi m \Omega^2} \frac{b}{(a+b)^2 \sqrt{2ab + a^2}} = 0.27 \tag{3.4}$$

This is in good agreement to the trap stability parameter q calculated based on the motional frequencies of the ion ω_i

$$q = \frac{4\pi}{\Omega} \sqrt{\omega_x^2 + \omega_y^2 + \omega_z^2} = 0.280 \tag{3.5}$$

using the motional frequencies measured in fig. 5.9.

This model is solved for the trap parameters, namely the location of the DC trap potential minimum, the ion axial motional frequency, and the tilt of the ion secular motional frequencies with respect to the trap. This tilt is necessary, as the lasers used for the Doppler cooling scheme of sect. 2.6 need to cool all three axes of motion, but are themselves parallel to the surface; as such they do not have a field component in the vertical direction. To avoid one of the secular motional frequencies being vertical, the set of DC voltages applied to top and bottom electrodes shown in fig. 3.4 are not the same; this induces a tilt in the secular frequency axes, allowing one set of Doppler cooling lasers, arriving from one direction, to cool all three axes of motion.

3.2.2 Microwave Fields

For reasons detailed in sect. 2.9.4, the design goal for the microwave resonator in this trap is to maximize the magnetic field gradient $B'_{MW,\pi} = B'_{MW,y}$, while minimizing the magnetic field $B_{MW,\pi} = B_{MW,y}$ affecting the chosen qubit[2], both

[2]Notably, it is not necessary to cancel out the entire magnetic field, which would require a more complicated trap structure.

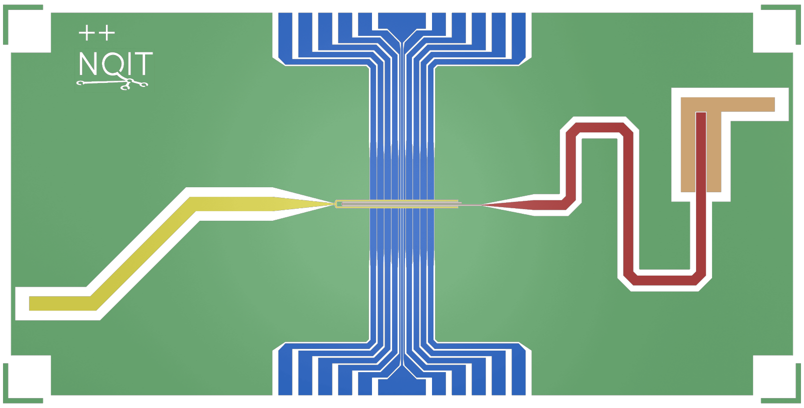

(a)

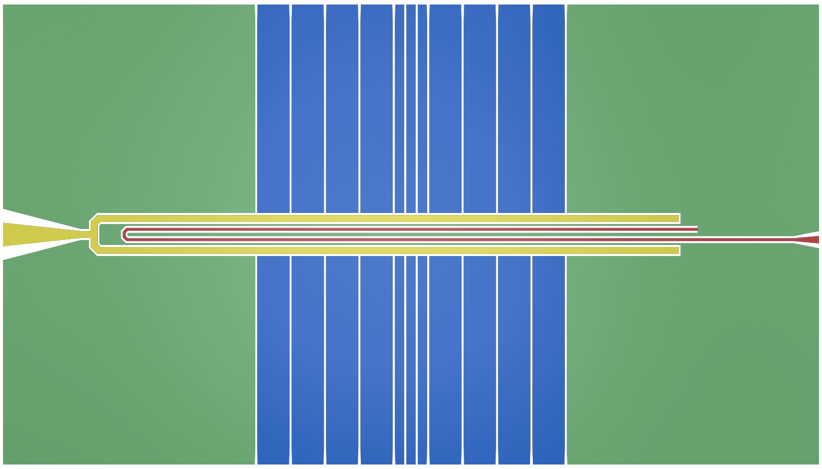

(b)

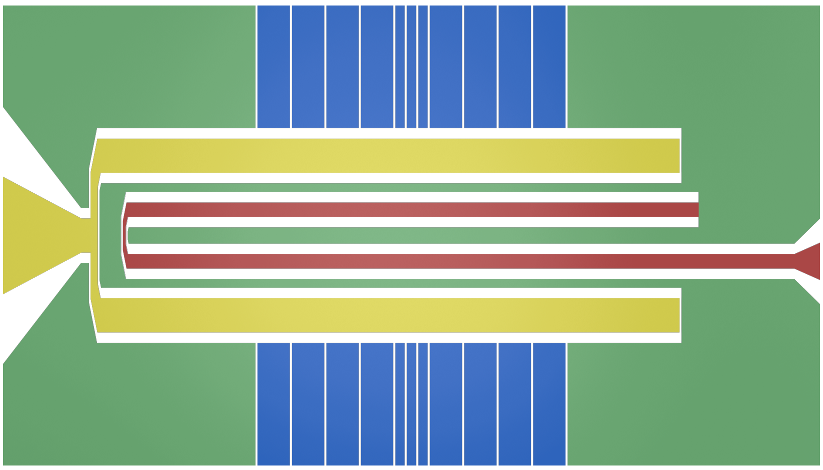

(c)

Figure 3.4: False colour diagrams of the ion trap, also see fig. 3.3. The colours are intended to highlight different functions: white for gaps in the conductive layer, green for ground, yellow for RF, red for the microwave resonator, orange for the incoming microwave line, and blue for DC. (a) Top down overview of the trap. (b) Zoomed in version around the centre. (c) As (b), but the vertical dimension is stretched by a factor of 5 to allow the narrow features to be more clearly distinguished. The microwave resonator is grounded at its end, creating a $\lambda/4$-resonator. Rendered with Solidworks 2016 Professional.

at the position of the ion. To achieve this, the microwave line is bent into a "U" shape, resulting in the field created by one line to be closely cancelled out by the other, the simulated field can be seen in fig. 3.5. Therefore, the overall electrode arrangement is DC-RF-GND- MW-GND-MW -GND-RF-DC. The magnitude of the residual field in the polarization of the qubit is mainly governed by the phase mismatch due to the lengths of the two lines, and thus is a function of the length of the interaction region. A discussion of the simulated versus the measured magnetic field gradient can be found in sect. 5.5.

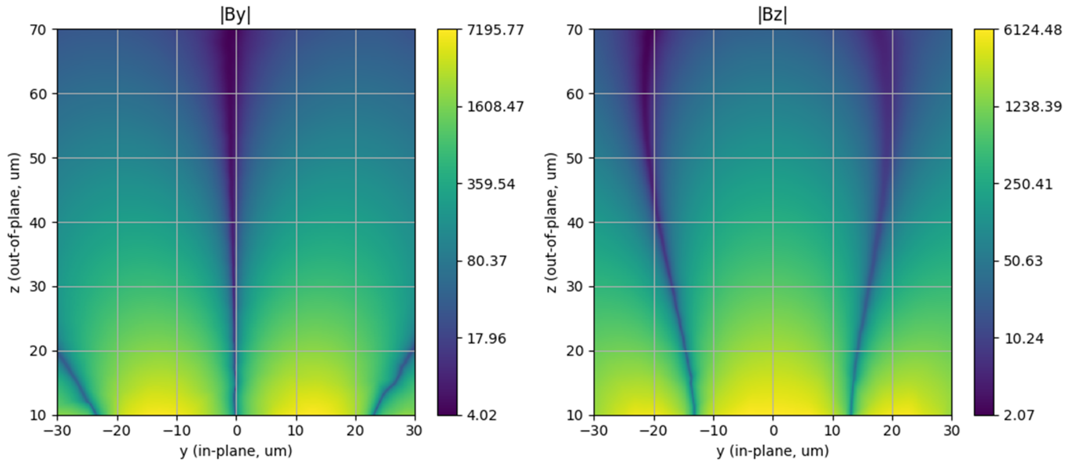

(a) (b)

Figure 3.5: Simulation of the magnetic field created by the microwave currents. The centres of the upper surface of the microwave conductors are located at $y = \pm 13.5\,\mu$m, $z = 0$. The ion is located at $y = 0$, $z = 40\,\mu$m. The magnetic field scale bar is shown in μT. The nominal microwave input power into the on-chip resonator is 1 W. (a) This shows the quality of the nulling of the π-polarised magnetic field $B_{MW,\pi}\,(y = 0, z = 40\,\mu m) = 6\,\mu$T. The deviation of the null from the $y = 0$ axis is due to the finite length of the microwave resonator's "U" shape and the resulting phase shift that causes this small asymmetry, also see fig. 3.4c. As can be seen, there are significant fields $B_{MW,\pi}$ in the $\pm y$ directions, leading to a high magnetic field gradient $B'_{MW,\pi}\,(z = 40\,\mu m) = 28\,{}^{T}\!/\mathrm{m}$. (b) This shows the strength of the σ- polarised magnetic field. Simulated with ANSYS HFSS.

3.2.3 Laser Access

Laser paths have to be taken into account when designing this chip, particularly for the location of the wire bonds. The static magnetic field is orientated along the top to bottom axis of fig. 3.3. As previously discussed in sections 2.6 and 2.7, the laser access needed is at angles of 0°, 45°, and 90° to the axis of the magnetic field to allow for the relevant laser polarisations to be produced. The initially planned vacuum window allowing for in-plane 90° beam axis was swapped for an additional ion pump; to allow for a 90° beam, it can be sent through the imaging system and arriving in the out-of-plane direction.

The chip is designed to allow all three sets of access parallel to the plane of the chip for the middle and left trapping zone as seen in the fig. 3.3. Due to space constraints of the used wire bonds and their pitch, the right trapping zone might have a wire bond blocking access.

3.3 Pillbox

The ion trapping chip is the centrepiece of the whole experiment. Surrounding it lies a box, hereafter referred to as the pillbox. It is made out of OFHC sourced from Montanwerke Brixlegg AG, chosen to increase the thermal conductivity, see sect. 3.6. The design features an octagon symmetry, as shown in fig. 3.6. The ion trapping chip rests on the centre of the pillbox. Immediately surrounding it are PCBs supplying the ion trapping chip with DC, RF, and microwave lines, see sect. 3.7. The pillbox is mechanically and electrically connected to the base flange, see 3.6.3. The base flange itself houses all electrical feedthroughs, as well as a few spare connections, see 3.3.4. Connected to the pillbox is a copper braid, shown in the top left of fig. 3.6. It transports heat away from the pillbox and towards the cryostat, see 3.6.

3.3.1 Black Body Radiation

In order to reduce the amount of thermal radiation absorbed by the pillbox and copper braid, both are gold-plated. This is done by SGA Technologies Ltd, with an average gold thickness of 5.4 μm, which is important both for attaching the ion trap onto the pillbox, see 4.2.1, and for making the surfaces chemically inert. The latter helps to maintain a clean surface for less outgassing and a better vacuum.

3.3.2 Laser Access

As the ions in the trap will be cooled and manipulated with lasers, as well as microwave gradients, optical access is required. This is achieved via 8 holes. The bigger the size of the holes, the greater the vacuum conductivity, which should be minimized, see sect. 3.4.1. The diameters of the holes are designed such that they minimize their vacuum conductivity, but still achieve negligible clipping of the laser beam.

For a Gaussian beam, the intensity profile is

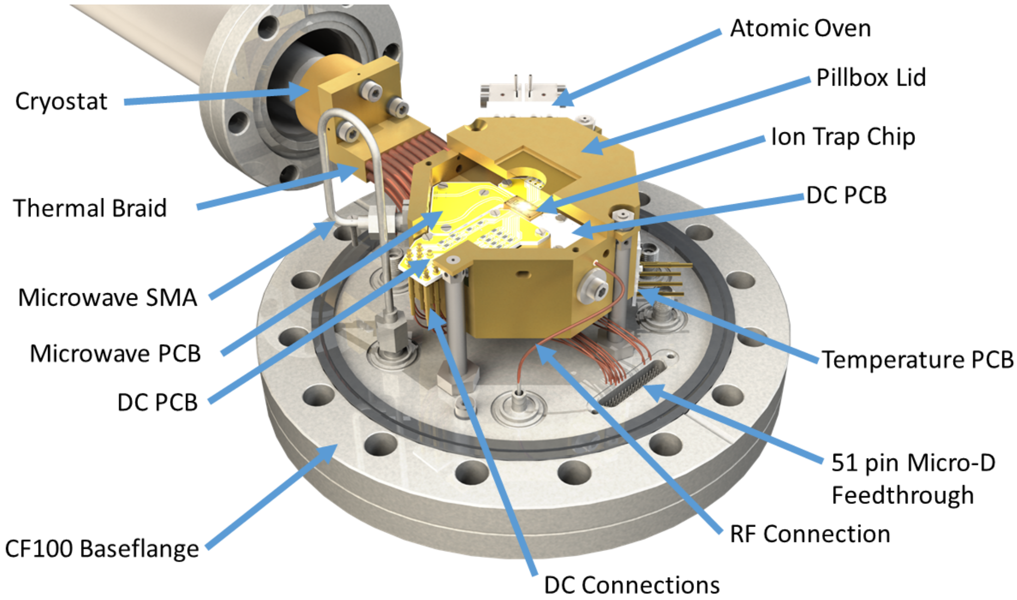

Figure 3.6: Image of pillbox, cryostat, copper braid, and base flange, all of which is explained in the text. Rendered with Solidworks 2016 Professional.

$$I(r) = \exp\left(-2\frac{r^2}{\omega^2}\right) \tag{3.6}$$

where the "spot size" ω is the distance from the axis at which the intensity falls to $1/e^2 = 0.135$, and with the above defined spot size ω and the intensity I as a function of the radius r from the beam centre. The spot size evolves with propagation distance z, as

$$\omega\left(z\right) = \omega_0\sqrt{1+\left(\frac{z}{z_R}\right)^2} \tag{3.7}$$

with $z_R = \frac{\pi\omega_0^2}{\lambda}$, using a spot size at the waist ω_0, and a wavelength of light λ. The amount of power going through a hole of radius r is [20]

$$P(r,z) = P_0\left[1-\exp\left(-2\frac{r^2}{\omega^2(z)}\right)\right] \tag{3.8}$$

with the initial power $P_0 = \frac{1}{2}\pi I_0\omega_0^2$. Given beam waist ω_0, wavelength λ, hole radius r, the distance between the outer edge of the hole and the beam waist of $z = 21$ mm, the ratio of power lost to initial power a can be calculated by combining equs. 3.7 and 3.8:

$$\alpha = \frac{P\left(r,z\right)}{P_0} = 1 - \exp\left(-2\frac{r^2}{\omega_0^2\left(1+\left(\frac{z\lambda}{\pi\omega_0^2}\right)^2\right)}\right) \tag{3.9}$$

Assuming a hole radius of $r = 0.75\,\mathrm{mm}$, and using the variables with the highest expected losses, i.e. beam waist $\omega_0 = 18\,\mu\mathrm{m}$, wavelength $\lambda = 866\,\mathrm{nm}$, the transmitted power ratio α is 99.998%. Only for hole sizes of $r < 0.5\,\mathrm{mm}$, the lost power becomes greater than 1%. However, to facilitate beam alignment, a hole radius of $r = 1.25\,\mathrm{mm}$ is chosen. The clipped power is then negligible if the beam is a Gaussian ($1 - \alpha = 8 \cdot 10^{-14}$). The effect on changing the transmitted modes is of higher concern than raw power; but with such a small change in power this, too, should be negligible.

This calculation assumes a perfect beam alignment. This, along with more practical issues of aligning the beam through such a small hole could mean that a bigger hole is needed. However, in the first iteration of the pillbox design, this is used as a starting point.

3.3.3 Lid

Besides laser access, the pillbox also needs to provide an aperture for light collection from the ions. This will be used for efficient readout of the internal state of the ions, see sect. 3.13. To that end, the cone of light that can be collected with the readout optics is a half angle of $\frac{\alpha}{2} = 13.1°$. Therefore the pillbox lid features such a cut-out, as seen in fig. 3.6. The lid also has a recess to allow for the spot welding of a thin gold mesh for electromagnetic noise screening.

3.3.4 Base flange

The flange that supports the pillbox is a custom CF 100 flange, made of SS 316. It features one RF[3], one 51 pin micro-D, and 4 SMA feedthroughs, parts 15310-03-W, 23053-01-W, and 15263-03-W from Testbourne Ltd, respectively, as well as several mounting holes. In the current iteration of the experiment, 6 pins of the 51 μD, and 3 SMA feedthroughs go unused, and are included for future proofing. The arrangement of the feedthroughs includes a minimum distance between feedthroughs of 1.5 mm, between feedthrough and tapped holes of 2.9 mm, and between feedthrough and copper gasket of 2 mm. These distances are chosen in order to create reliably leak-proof welds. The base flange is manufactured by Testbourne Ltd. To increase the area available on the baseflange, high clearance gaskets – part SG-060X from Duniway Stockroom Corp – are used, i.e. ones with a higher internal radius than standard CF100 gaskets.

3.3.4.1 Legs

The mechanical support of the pillbox is provided via rigid legs attached to the base flange. As the cryostat used to cool the pillbox is both a source of vibration, and expands significantly upon cooling, the attachment between the two is flexible, see sect. 3.6.[4] Their material has to have good mechanical properties as well as low thermal conductivity, as will be discussed in sect. 3.6.3. Three good choices of materials are stainless steel (SS), the machinable ceramic Macor, and the plastic PEEK. Their thermal conductivities vary by two orders of magnitude; their mechanical strengths also vary.

At first, legs made of Macor, part 2255533 from RS Components, were chosen, due to an order of magnitude lower thermal conductivity, see tab. 3.1. However, Macor is prone to cracking, both during machining, and during tightening. Sev-

[3] This is an SMA feedthrough with a bare ⌀1 mm pin on the vacuum side.

[4] For future experiments the pillbox could be rigidly connected to the cryostat, without further mechanical support, if the vibrations of the cryostat are small enough not to interfere with the operation of gates.

eral iterations of Macor legs were made. They need to interface with an M3 tapped hole in the base flange on the one side, and an M3 clearance hole in the pillbox on the other, and need to provide venting holes for any trapped gas. In all designs this is achieved by a ⌀1 mm coaxial hole through the centre of the legs. The bulk of the 40 mm long legs was made of ⌀5 mm Macor, with a ⌀4 mm central cut-out to reduce thermal conductivity.

The first design featured female threads for the pillbox and male Macor threads for the base flange. The latter proved very easy to break, presumably at least in part due to the coaxial venting hole.

The second design had female threads on both sides; the bottom was planned to be fastened using Ag plated grub screws, part MT4 -316-A from U-C Components Inc. While this allowed the legs to be inserted and loosely screwed in, they shattered upon even hand tightening the screws.

The third design features multiple parts; a shorter Macor section to allow room for two SS 304 inserts on the ends. These have male M3 threads on both sides; one that screws into the Macor section. The other screws into the base flange and is tightened by an M3 nut onto the pillbox. As with the Macor part, these inserts have a ⌀1 mm coaxial hole for venting. They also have a hexagonal outer section to allow for counter-torquing with a spanner; a 10 mm spanner for the base flange and a 5 mm spanner for the pillbox insert. The advantage of first screwing the SS inserts into the Macor, and then tightening the inserts is that the inserts can be precisely machined to have a very flat base. The hypothesis is that the second design failed due to a non-flat surface being pressed onto the brittle Macor, causing localised spikes in pressure; hence a flat base is advantageous. Testing the breaking torque with a torque wrench, the Macor legs break at 0.20 N·m, while a typical highest torque applicable by hand is 0.14 N·m. The inserts could be screwed into the Macor section and hand-tightened, without any damage occurring. The legs were thermal shock tested

— repeatedly submerged in liquid nitrogen and held over a hot air gun — and survived this without damage, too. However, during the first test of the assembled system, they broke upon cooling to 4 K, see fig. 3.7. As the break occurred near the bottom of the leg, both the SS insert and the Macor should be close to room temperature, so differential expansion of these two parts is less likely to be the cause of the breakage. Differential expansion of the cooled pillbox to the room temperature base flange pulling the three legs together could be the issue. However, subsequent results obtained in sect. 5.6 point towards the thermal braid being the cause.

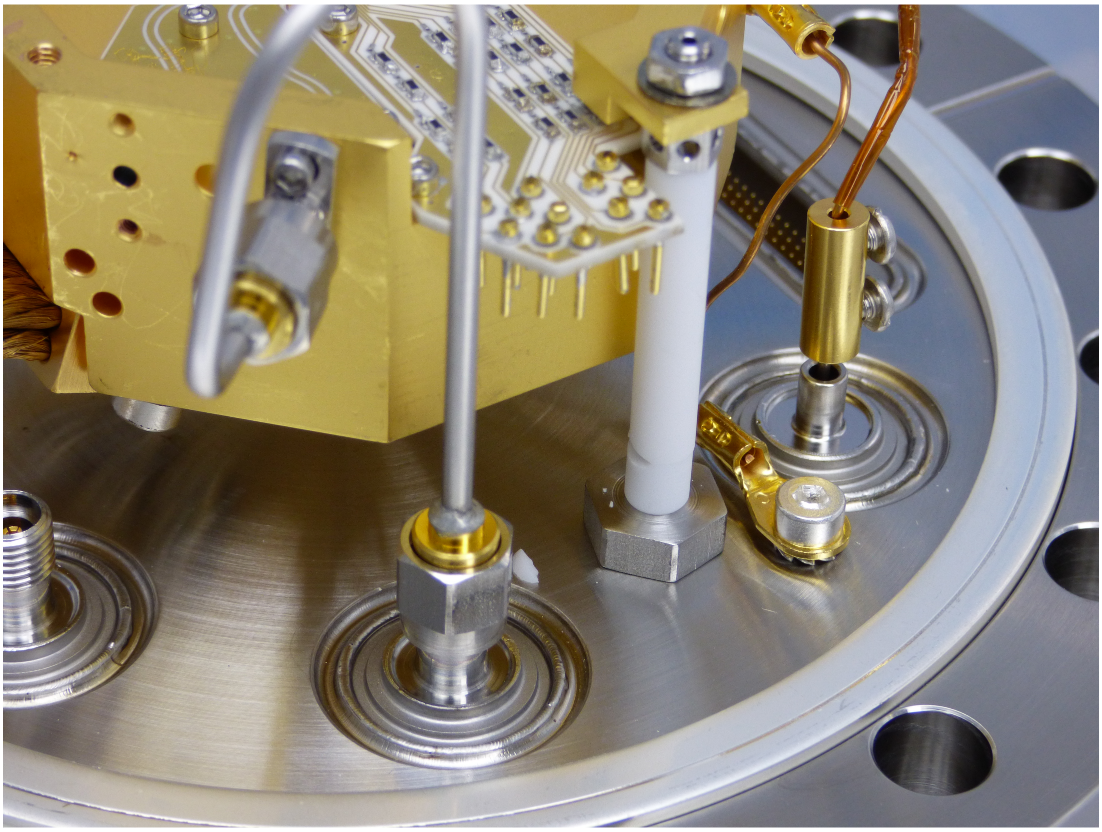

Figure 3.7: Picture of the broken Macor leg (white cylinder, middle of picture). The pillbox was assembled onto the base flange, the vacuum system evacuated and cooled using the cryostat. During the first rapid cool-down, with peak temperature rates of $40\,{}^{K}/_{\mathrm{min}}$, one of the three Macor legs broke; seen here after partial disassembly of the vacuum system. Taken with a Panasonic DMC-FZ72 camera.

Due to time limitations, the cause of the failure of the third design iteration of the Macor legs was not further investigated. Instead, they were replaced using a

SS version. For the long middle section a tube with a wall thickness of 0.24 mm, at an outside diameter of 5.5 mm, is chosen, made of SS 316, part 506-371-68 from Goodfellow Cambridge Ltd. To make them interchangeable with their Macor counterparts, inserts with a diameter slightly larger than the inside diameter of the tubes, are crimped into both ends. To increase friction, the insert surfaces are roughened with a knurling process before insertion. The inserts, also made of SS 316, have female M3 threads. Despite the much thinner wall thickness, 0.24 mm vs 1 mm, the thermal conductivity of the SS legs is about three times that of their Macor counterparts. For a more detailed comparison see sect. 3.6.3. As the SS inserts are useful for counter-torquing, they are re-used for these SS legs.

As the SS legs have a comparatively high heat load, they could be replaced using a future PEEK version of the legs, but their mechanical suitability still needs to be determined.

3.3.4.2 Calcium & Strontium Ovens

In order to obtain trapped calcium ions, a small sample is heated to create a flux of neutral Ca atoms across the ion trap's surface. It consists of a SS tube, crimped shut on one end, with a diameter d = 2.2 mm, in which nuggets of calcium are placed. Subsequently, the tube is crimped shut on the other end. Naturally occurring calcium features 96.941% ^{40}Ca and 0.135% ^{43}Ca. As this work uses both ^{40}Ca^{+} and ^{43}Ca^{+} ions, the loading rate of ^{43}Ca is increased by using isotopically enriched calcium. The calcium is bought from the National Isotope Development Center (NIDC), specified as a mixture of 30.0% ^{43}Ca with 81.12% purity, and 70.0% ^{40}Ca with 99.97% purity. The final mixture is thus 30.0% $\cdot$ 81.12% = 24.3% ^{43}Ca. Assuming the "impurity" of the ^{43}Ca stems entirely from ^{40}Ca, the concentration of ^{40}Ca is 75.6%. Therefore, the overall mixture contains 24.3% ^{43}Ca, and 75.6% ^{40}Ca.

For future experiments not featured in this work, ^{88}Sr can be included in a second SS tube[5], with a purity of $(99.84 \pm 0.02)\%$, part 7440-24-6 also from the NIDC. Impurities are ^{87}Sr at $(0.11 \pm 0.02)\%$, ^{86}Sr at 0.05%, with specified amounts of less than 0.05% from any single source; notably including calcium at <5 ppm, although this should not cause problems.

Previously, the SS tube would be welded between two pins of an electrical power feedthrough on the base flange, completing the electrical circuit between the pins. The SS tube would then heated up through the application of a static current. A small hole in its side, pointing through a hole in the pillbox to the ion trapping region, then allows a small amount of calcium vapour to reach the trapping region. Upon arrival, the calcium encounters a two-step ionisation laser, makes ionisation isotope selective. The Ca comes in 5.4 mg packages[6], which is enough to fill about two oven tubes.

However, in this work, a new oven design is used, shown in fig. [6]. It features the pair of SS tubes being clamped in a construct of SS and Macor. This design has multiple advantages: Its position can be mechanically well defined; as both Ca and Sr oxidise rapidly in air, it is an advantage to be able to quickly insert the SS tubes into the oven without having to align the tubes. Also, the inclusion of a thermal probe allows the oven to be over-driven safely. That is, one may intentionally dissipate more heat than would be safe in steady state, heating the oven up to the target temperature quickly. The applied current is then reduced to a level safe for steady state.

Power is supplied via the 51 pin micro-D feedthrough in the base flange, and there is an option to include a thermal probe spot welded onto the SS tubes. Due to the relatively high peak current requirements, up to 9 A, of this oven design, 3 physical pins are used for each logical connection, thus totalling $3 \cdot 4 = 12$

[5]While assembling the current system, the decision was made to leave out the ^{88}Sr as its rapid oxidation would have caused problems in a hypothetical future opening of the vacuum system to e.g. correct an error made in the assembly. The second SS tube is fully connected and, after breaking vacuum, could be used by simply inserting the ^{88}Sr and crimping the end of the SS tube shut.

[6]Vacuum packaged, $< 5 \cdot 10^{-6}$ Torr, to avoid oxidation

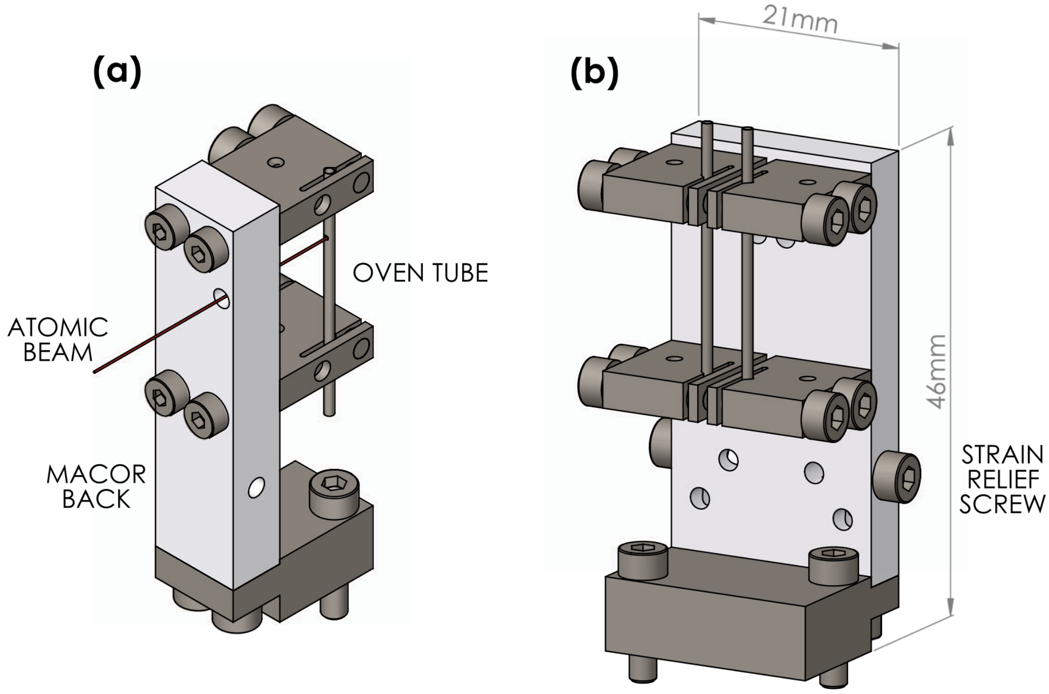

Figure 3.8: Overview of the used type of oven. The oven design can be used for both (a) single species and (b) dual species experiments. Thermocouples (not shown) are spot-welded onto the oven tubes, and current-supplying wires are screwed into the holes next to the oven tubes. This work uses a slightly adapted dual species oven design, with one big hole in the macor back to accommodate both the oven tubes' atomic beams and a laser beam path. Taken from [6].

connections for the oven's current supply; another 4 are used for two pairs of type K thermocouples. This implies that the measured thermocouple voltage only measures the temperature difference between the oven's junction and the feedthrough. The latter's temperature is not directly measured, but just assumed to be close to room temperature. As the temperature of the vacuum chamber, and thus the feedthrough, is stable, this reduces the accuracy of the measured temperature by a small, but constant, amount.

3.4 Ultra High Vacuum

Trapped ions need a vacuum in order to avoid background gas collisions from both knocking them out of the trap and from introducing errors in gates. To that end, a vacuum system is designed. The base-flange with the pillbox and oven on it, see sect. 3.3.4, is attached to an octagonal UHV vacuum chamber, part MCF600-SphOct-F2C8 from Kimball Physics, with a CF100 port on the top and bottom, and eight CF40 ports arranged in an octagonal way around it. For optical access, six CF40 windows and one CF100 window are attached to it, see sect. 3.5, leaving two opposing CF40 ports blank. One of those ports is connected to a cryostat, see sect. 3.6. The other CF40 port is attached to a 5-way cross that houses the experiment's vacuum pumps, ion gauge, and valves, which are explained in this section.

In order to protect any trapped ions from collisions with background gas an Ultra High Vacuum (UHV) is needed, defined as $1 \cdot 10^{-9}$ to $1 \cdot 10^{-12}$ mbar. As a general guideline, a pressure of $< 10^{-11}$ mbar is aimed for; this has achieved two-ion lifetimes of ≈ 1 hour in previous work [p. 105f, 51]. Throughout the entire vacuum system, care is taken to avoid trapping any pockets of gas that could then become a virtual leak. To that end venting holes are used, wherever possible. Otherwise, vented screws are used.

3.4.1 Pump

In order to gauge the effect of different vacuum geometries on the final achievable pressure, the molecular flow model is used [p. 32-49, 21], and introduced in sect. 4.4.2. A schematic overview of the vacuum system is shown in fig. 3.9.

The vacuum is maintained via two pumps. Firstly, a NEXTorr D200-5, part 5H0170, is used, which is a combination of sintered porous getter disks, part St 172, which can pump H_2 at the eponymous 200 $^l/_s$, with a small diode ion pump, which can pump Ar at 6 $^l/_s$, controlled via a power supply, part 3B0408, all from SAES Getters. To increase the pumping speed of gases not effectively pumped with the getter disks, a StarCell Vacion Plus 20 ion pump, connected via a 200°C bake-able cable to a MiniVac Controller, parts 9191145, 9290705, and 9290290, from Agilent Technologies LDA UK Ltd, is used. It features an eponymous 20 $^l/_s$ pumping speed for N_2.

The overall pumping speed is dependent on which molecules need to be pumped, the pressure, temperature, saturation level of the pumps, and vacuum conductivity to the pumps from the source. A discussion of these effects can be found in sect. 4.4.1.

Pressure is measured with a Nude Bayard-Alpert Ion Gauge, part P104145 from Ideal Vacuum Products LLC., which has two iridium filaments, although only one is used at any given time - the other is used as a backup in case the first one breaks.

The pumps that are included in the vacuum system do not have the capacity to create a vacuum from atmospheric pressures, only to maintain and improve it. The initial creation of the vacuum involves several types of pumps, which will be elaborated on in sect. 4.4.4. After the initial vacuum has been created, the system is sealed with ZCR40R all metal valves by Vacgen. In case one such valve were to fail, and to concentrate wear and tear on only one of them, two valves are put in series, with only the outer one being used in normal operation.

Importantly to note, these valves limit the maximum bake temperature to 350°C.

The two pumps, the ion gauge and in-series valves, are connected to a customised 5-way cross, ZFL-5X63/40-JWOX from LewVac Components Ltd, with all connectors being CF 40. To maximise the vacuum conductivity from the octagonal vacuum chamber, the CF40 cross is welded together using CF63 piping, which has the same outer diameter as the CF40 flange's end piece.

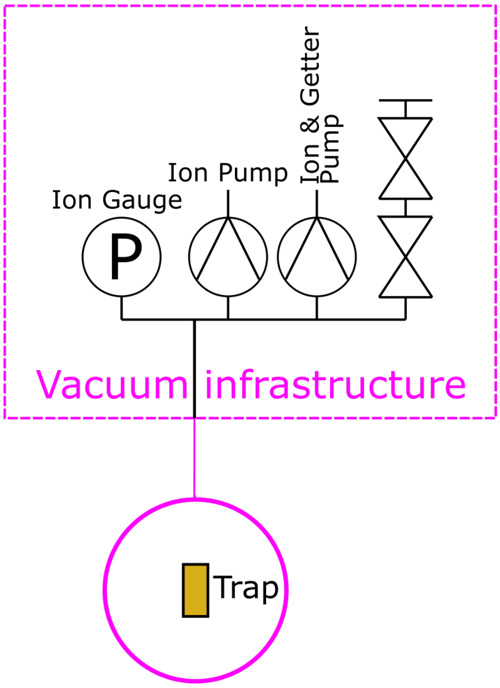

Figure 3.9: Schematic overview of the vacuum system. It consists of an ion gauge, an ion pump, and a combination ion & getter pump, as well as two valves for the initial vacuum creation (placed in series in case of leaks).

3.4.2 Outgassing Rates

In order to achieve UHV conditions, it is imperative that low outgassing materials are used. Fig. 3.10 presents a collection of five different sources of data on outgassing rates. Between the sources, these rates may differ by many orders of magnitude for the same materials. This is in part due to the importance of the treatment of the materials. Baking temperature and duration, and surface cleanliness are all important factors that influence the final outgassing rate.

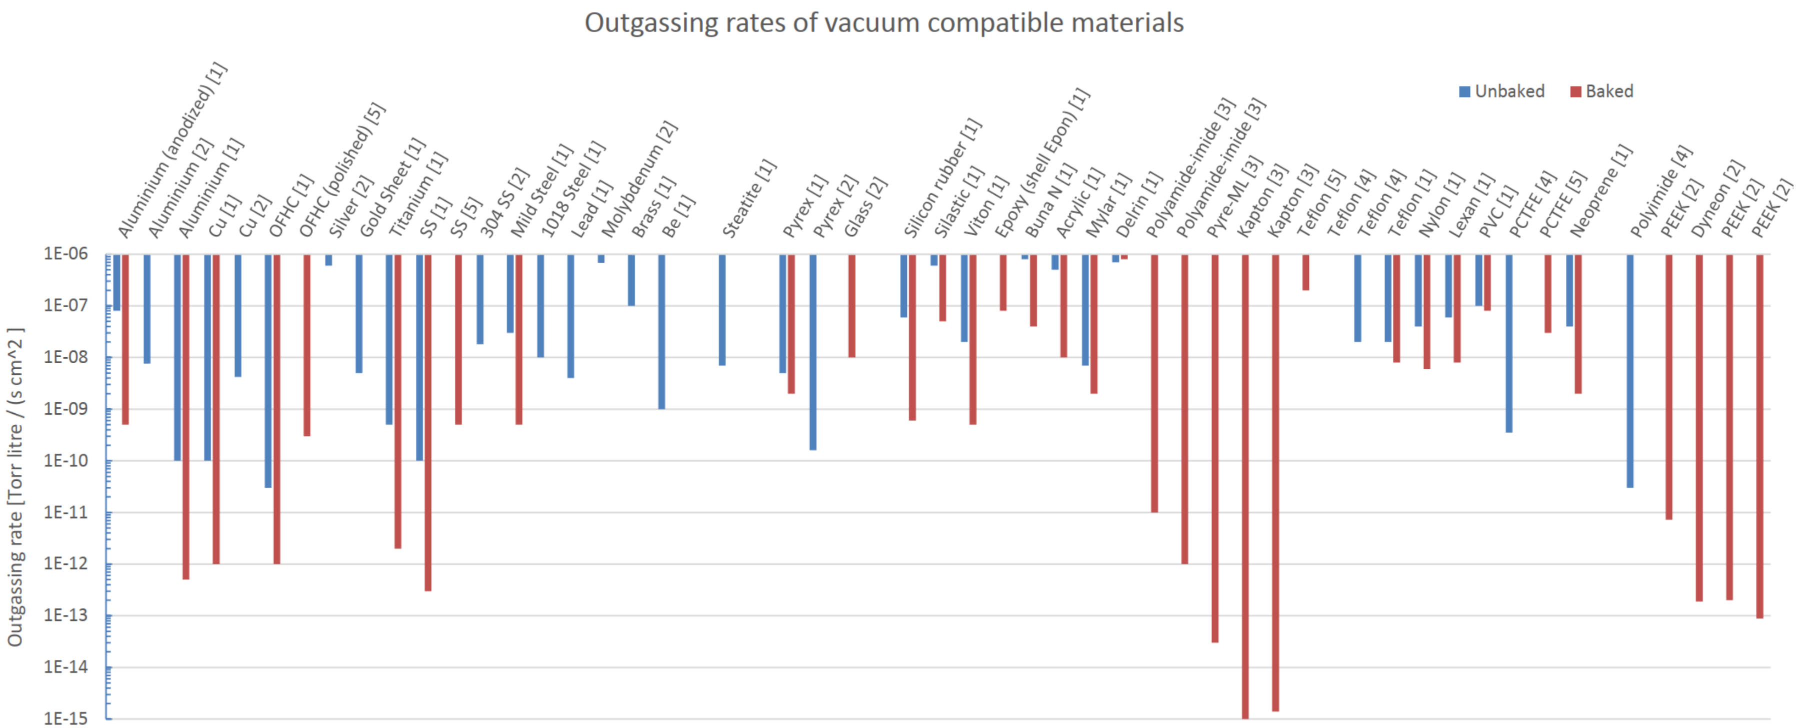

Figure 3.10: Outgassing rates of various materials. These values should be taken as a very rough guide only, and depend on many factors like cleanliness, surface condition, type of bake, and purity. General guidelines in the group are to use metals like SS, Cu (in the form of Oxygen Free High thermal Conductivity copper, OFHC), Ag and Au, as well as ceramics like Macor and to avoid – or sparsely use – plastics, while the only acceptable plastics would be Kapton, PTFE, and PEEK. However, the outgassing rates found here seem to contradict these guidelines; for example steel seems to potentially be outgassing a lot more than Kapton or PEEK, while PTFE (Teflon) seems to be a high outgassing choice. As the group's guidelines have historically resulted in acceptable vacuum pressures, they are continued to be used in this thesis, despite the contradictions. Sources (1) to (5) are [7, 60, 11, 44, 22], respectively.

3.5 Optical Access

For optical access, one CF100 viewport and six CF40 viewports are used.

The CF40 viewports are custom part 190-62056-5 from Allectra Ltd., made of SS 316 and fused silica, part FPFS 7980 from Corning Inc. The surfaces are coated to decrease reflectivities, whose values can be found in fig. 3.11a.

The CF100 viewport is a custom piece from the UK Atomic Energy Authority (UKAEA), made of fused silica, part Spectrosil 2000 from Heraeus Holding GmbH, and SS 316. Its window is recessed by 25 mm[7] into the vacuum chamber, which allows a lens to be brought closer to the ion, and ultimately a greater angle of light to be collected from any ions, see fig. 3.11b. Its anti reflection coating, from Laseroptik GmbH, can also be found in fig. 3.11a.

R[%]	Wavelength[nm]	
	CF40	CF100
0.25	397, 422	
0.35	674, 729	
1	375, 393, 408, 423, 461	360 – 430
5	850, 854, 866, 1004, 1033, 1092	
10	375 – 1092	
20		650 – 1100

(a)

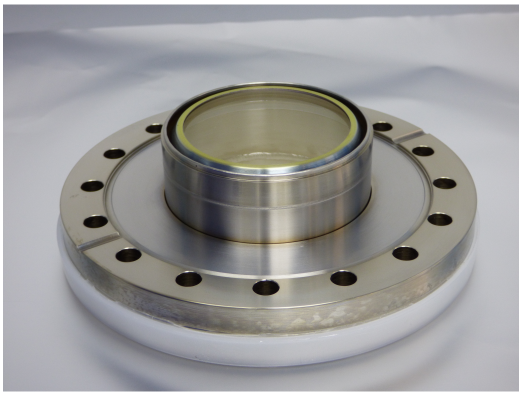

(b)

Figure 3.11: Specified upper limits (a) on reflectivities R per surface of the viewports as a function of wavelength. (b) The recessed CF100 viewport as seen from the vacuum side, taken with a Panasonic DMC-FZ72 camera.

3.6 Cryogenics

In order to cool most of the surface that the ions see to cryogenic temperatures, a cryostat is used. The initial target temperature to be reached was 20 K, however later considerations have lowered this to $\approx$ 16 K, see fig. 4.15. This section uses the "nominal" temperature of 20 K.

[7]as measured from the vacuum side of the flange to the vacuum side of the window

3.6.1 Cryostat

The cryostat used in this experiment is capable of both liquid helium (LHe) and liquid nitrogen (LN_2) use, part ST-400 UHV Supertran from Janis. It is a flow cryostat, i.e. the cryogen flows in through the centre of a double walled pipe, and flows past the inside of the cryostat tip. Upon arriving there, most of the liquid will evaporate, and the gas flows in the outer layer of the double walled pipe, reducing the heat-load on the centre. Further details can be found in fig. 3.12.

The length of the cryostat from the tip to the other end is 485.90 mm, necessitating a custom vessel used for its cleaning, see sect. 4.4.5. According to the manufacturer, during cool down from room temperature to both liquid nitrogen and liquid helium temperatures, it contracts by 0.9 mm. This contraction needs to be compensated to avoid putting strain on the pillbox mounting, which could result in deformation and damage of the weakest link. Additionally, the manufacturer reports that there is anecdotal evidence for vibrations of “a few microns”. In this experiment, both the contraction compensation and vibration isolation is done via a copper braid.

3.6.2 Copper Braid

The copper braid is a P6-503 compact option OFHC Cu braid, part P16-0418 from Technology Applications Inc., with a braid length of $25.4^{+1.3}_{-0.5}$ mm. In order to minimize the thermal resistivity, by maximizing the force pressing surfaces together, it has three M3 through holes to connect to the cryostat, and five M3 through holes to connect to the pillbox.

Initially, a concern was an excessive force applied to the pillbox' mounting due to bending of the braid. It has a spring constant of $\approx$ 0.5 N/mm, interpolated from manufacturer's data. Therefore, the designed 5 mm bend of the braid results in a force of $\approx$2.5 N, which is well below the tolerances of the legs holding up the pillbox. The thermal conductivity is specified to be 0.93 $^{W}/_{K}$, at a temperature

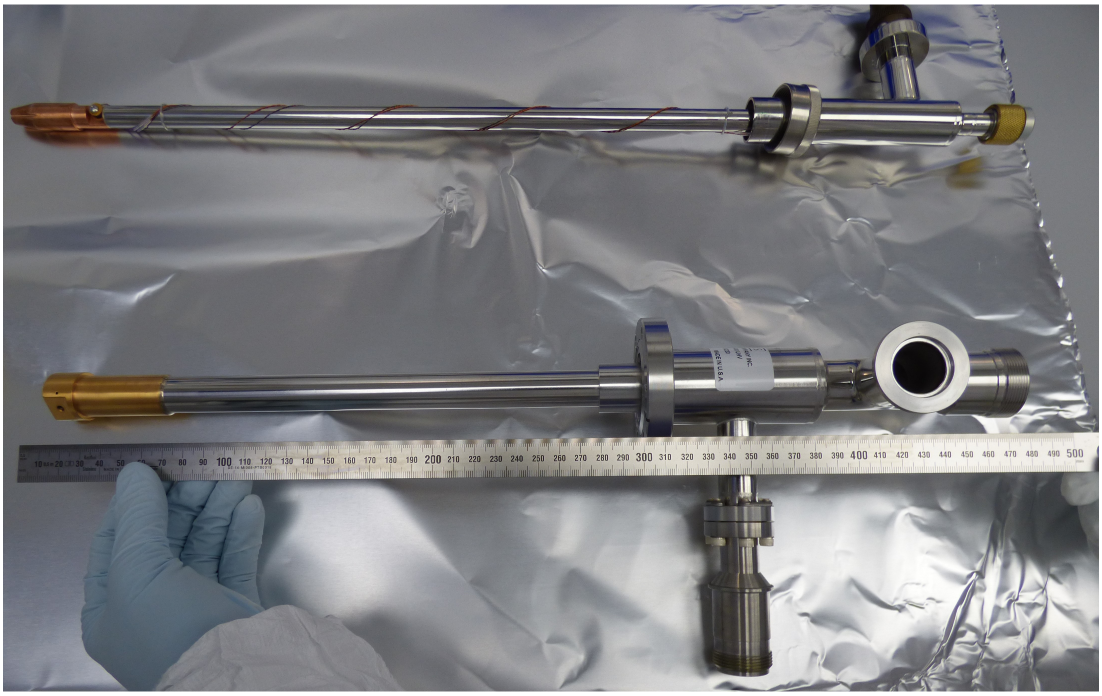

Figure 3.12: Picture of the cryostat. The bottom piece is installed into the vacuum chamber via a CF40 joint (the flange sits slightly to the right of the middle). The left side is exposed to UHV conditions, and the gold plated end interfaces with a copper braid detailed in sect. 3.6.2. Pointing downwards is a CF16 electrical feedthrough that comes with the cryostat; this is later replaced with a CF16 blanking flange as it is not needed. The top piece inserts into the bottom one by sliding it into the end. LHe is injected via a transfer line (not pictured) that slides into the centre of the top piece. The temperature is measured via a diode located in the left tip of the insert. An electrical heater is also included in that tip. Both the diode and the heater are addressed via a connector at the top edge of the image. The ruler is 50 cm long. Taken with a Panasonic DMC-FZ72 camera.

of 20 K and braid length of 100 mm. Therefore, the thermal conductivity κ_{Braid} of the chosen braid is $0.93\,{}^{W}/_{K} \cdot \frac{100\,\text{mm}}{25.4\,\text{mm}} = 3.7\,{}^{W}/_{K}$.

3.6.3 Thermal Budget

It is important to know how much liquid helium will be used during normal operation of the system, both due to the cost of the cryogen, and to appropriately size any storage system. As such, thermal budgeting is used. However, the thermal conductivity of materials changes with temperature, as seen in table 3.1. During operation at $\approx$ 20 K, some of the parts relevant for the heat budget are close to being at a single temperature, e.g. the copper braid, while others are connected to a room temperature reservoir on one side, and $\approx$ 20 K on the other. Therefore, table 3.1 features both the thermal conductivities at single temperatures, e.g. $\kappa\,(20\,\text{K})$, and their integrated values for a temperature difference, e.g. $\int_{20K}^{300K} \kappa \text{dT}$. As the flow cryostat could potentially be run with the more economical option of LN2, the thermal conductivities are also specified for 77 K, the boiling point of LN2 at atmospheric pressure.

The nominal target temperature to be achieved in this thesis is T_{He} = 20 K, is well above the boiling temperature T_b of LHe. This is mainly for economic reasons: the latent heat of vaporization L per litre is nearly two orders of magnitude lower for LHe, $L_{He} = 2.614\,{}^{kJ}/_{l}$, than for LN_2, $L_{N_2} = 160.6\,{}^{kJ}/_{l}$. The basic properties of the two considered cryogens can be found in table 3.2. This, combined with about an order of magnitude increase in price per litre[8], means that it would cost nearly three orders of magnitude more to give the same cooling power at the boiling temperature of LHe than that of LN_2.

At 20 K, the specific heat capacity of the helium vapour is the main source of cooling power Q, not the latent heat of vaporization L:

[8] At the time of writing, the Department of Physics in Oxford charges 0.50 $^{£}/_{l}$ for LN_2, and 5.85 $^{£}/_{l}$ for LHe. For LHe, the price assumes that the resulting gas is recycled.

Material	$\int_{20K}^{300K} \kappa \mathrm{dT}$	$\int_{77K}^{300K} \kappa \mathrm{dT}$	$\kappa\,(20\mathrm{K})$	$\kappa\,(77\mathrm{K})$	$\kappa\,(300\mathrm{K})$
Unit	$\left[\frac{\mathrm{W}}{\mathrm{m}}\right]$	$\left[\frac{\mathrm{W}}{\mathrm{m}}\right]$	$\left[\frac{\mathrm{W}}{\mathrm{m \cdot K}}\right]$	$\left[\frac{\mathrm{W}}{\mathrm{m \cdot K}}\right]$	$\left[\frac{\mathrm{W}}{\mathrm{m \cdot K}}\right]$
Polyimide (Kapton) [47]	42.9	37.6	0.0478	0.126	0.192
PTFE (Teflon) [47]	69.8	58.3	0.142	0.232	0.273
Macor [50] and [p. 179, 52]	330.	270.	0.6	1.6	1.4
SS 304 & 316 [47]	3 010.	2 700.	2.17	7.92	15.3
Be Cu [47] [p. 12-205, 12]	8 150.	6 740.	10.66	36.3	8.9
Alumina [p.12-221, 12]			23.	150.	33.
Mo [p.12-206, 12]	50 350.	32 180.	287.	216.	139.
Au [p.12-218, 12]	105 000.	70 000.	1 580.	340.	317.
OFHC Cu (RRR=50) [47]	147 000.	91 700.	1 368.	515.	392.
OFHC Cu (RRR=100) [47]	167 000.	93 800.	2 423.	547.	396.
Ag [p.12-218, 12]	170 000.	94 000.	5 100.	480.	429.
OFHC Cu (RRR=150) [47]	178 000.	94 400.	3 245.	556.	398.
OFHC Cu (RRR=300) [47]	197 000.	95 400.	5 052.	572.	398.
Sapphire (36° to c-axis) [p.12-221, 12]			3 500.	1 100.	
OFHC Cu (RRR=500) [47]	211 000.	95 900.	6 575.	579.	402.

Table 3.1: Comparison of thermal conductivities of materials used in the experiment, with thermal conductivities κ at a specified temperature, as well as their integrals over a specific temperature difference. RRR is the residual resistance ratio, a measure for the purity of a metal. It is defined as the ratio between the resistivity at 300 K to that at 0 K (although it is often measured at 300 K and 4.2 K for experimental simplicity), meaning the higher the RRR, the purer the sample. Some values are interpolated as their sources do not specify the above values directly.

$$Q = L \cdot \rho + c\,(l) \cdot \rho \cdot (T_{He} - T_b) = 21.01\,\mathrm{kJ/l} \tag{3.10}$$

This increases the cooling power of the LHe option by about an order of magnitude, making it more affordable, albeit still two orders of magnitude more expensive than LN_2 if used to dissipate the same power[9].

Tab. 3.3 shows the heat budget for the pillbox and its environment. The contributions of the support legs, DC & RF wires, co-axial microwave lines, as well as the black body radiation are considered. The contribution from the dissipation of microwave power on the ion trapping chip is to be determined, and the current figure of 2 W is just a tentative estimate.

[9] To decrease the long-term costs of cooling to 20 K, a closed-cycle system can be used. The above approach assumes the LHe is liquefied using room temperature He, recycled from the cryostat via uninsulated piping. An alternative would be to use a closed-cycle system. The advantage is that the output of the cryostat, still being only slightly above 20 K, is then directly used as the starting point of the cooling system, as opposed to 300 K. Thus, much less energy needs to be dispelled, making this a more economical option in the long-term.

Property	**Helium**	**Nitrogen**	**Source**
Triple Point	2.1768 K	63.151 K	[p. 6-163, 12]
Normal boiling point	4.222 K	77.355 K	[p. 6-163 12]
Latent heat of vaporization L	20.91 $\frac{\text{kJ}}{\text{kg}}$	199.2 $\frac{\text{kJ}}{\text{kg}}$	[p. 6-163, 12]
Molar mass	4.0026 $\frac{\text{kg}}{\text{kmol}}$	28.0134 $\frac{\text{kg}}{\text{kmol}}$	[p. 6-163, 12]
Density ρ of liquid at normal boiling point	0.1250 $\frac{\text{kg}}{\text{l}}$	0.8061 $\frac{\text{kg}}{\text{l}}$	[p. 6-163, 12]
Expansion ratio, T_b to $T = 21°\text{C}$	1 : 754	1 : 696	[46]
c_v of liquid at normal boiling point	≈2.6 $\frac{\text{kJ}}{\text{kg K}}$	1.085 $\frac{\text{kJ}}{\text{kg K}}$	[p. 6-22 & 6-28, 12]
c_v of vapour at normal boiling point	≈3.12 $\frac{\text{kJ}}{\text{kg K}}$	0.7710 $\frac{\text{kJ}}{\text{kg K}}$	[p. 6-22 & 6-28, 12]
c_p of liquid at normal boiling point	5.105 $\frac{\text{kJ}}{\text{kg K}}$	2.041 $\frac{\text{kJ}}{\text{kg K}}$	[p. 6-163, 12]
c_p of vapour at normal boiling point	9.327 $\frac{\text{kJ}}{\text{kg K}}$	1.124 $\frac{\text{kJ}}{\text{kg K}}$	[p. 6-163, 12]

Table 3.2: Properties of two cryogens. Normal conditions are a pressure of P_n=1 atm=101.325 kPa. [p. 220, 15] and [p. 6-22 & 6-28, 12]

	Material	d_{out} mm	d_{in} mm	l mm	$\int_{20K}^{300K} \kappa \mathrm{dT}$ $^W/_\mathrm{m}$	n	P W
Support legs	SS 316	5.5	5.02	24	3010.	3	1.43
DC wires	BeCu	0.321	/	100	8150.	28	0.18
RF & GND wire	BeCu	0.81	/	100	8150.	2	0.08
Coax	SS 304	2.19	1.64	100	3010.	1	0.13
	PTFE	1.64	0.51		69.8		
	Ag	0.51	0.5		170000.		
	OFHC Cu (RRR=100)	0.5	0.45		167000.		
	SS 304	0.45	/		3010.		
Radiative							0.15
Active							2.0
					Total:		4.0

Table 3.3: Heat budget between the pillbox and its environment for LHe. All of the parts feature a cylindrical symmetry. Their individual heat loads are calculated using their outer diameters d_{out}, inner diameters d_{in}, lengths l, integrated thermal conductivity $\int_{20K}^{300K} \kappa \mathrm{dT}$, and quantity n. For the legs, only the section without insert is counted towards the length. The coaxial cable is described in sect. 4.3.3. The total heat load of 4.0 W implies a liquid helium consumption of $\dot{V} = \frac{P}{Q} = 0.7\,^l/_\mathrm{hr}$ with active microwave power, and $0.3\,^l/_\mathrm{hr}$ without. This also implies the temperature difference with active microwave power across the copper braid to be $\Delta T = \frac{P}{\kappa_\mathrm{Braid}} = 1.1\,\mathrm{K}$.

3.6.4 Thermal Expansion

The effect of temperature on the volume of materials is an important consideration in this experiment due to the wide range of temperatures experienced; the parts get assembled at $T_{room} \approx 300\,K$, they will be baked at $T_{air} \approx 450°C$, $T_{hard} \approx 350°C$, and $T_{soft} \approx 190°C$, and subsequently cooled down to $T_{cryo} \approx 20\,K$, i.e. the nominal temperature used for LHe cryogenics in this experiment. As will be explained in sect. 4.4.9, the baking temperatures correspond to an air bake, to create an outgassing reducing layer of oxide on the metal surfaces, a "hard" bake which consists of the bulk of the vacuum chamber (with the exception of the windows and the baseflange) to get the majority of the system outgassed, and a final "soft" bake, which includes all vacuum components.

The effects of thermal expansion are particularly important with the parts in the base flange assembly which experience T_{room}, T_{soft},and T_{cryo}. While the temperature change between T_{room} and T_{soft} is only experienced once for these components, that between T_{room} and T_{cryo} is repeated multiple times; this can lead to e.g. bolts becoming looser over multiple thermal cycles. In turn this increases the hysteresis of regaining the ion's position after a thermal cycle, as well as decreasing the thermal conductivity to the cryostat, leading to higher liquid helium consumption. Furthermore, different rates of thermal expansions can lead to high stress in materials, potentially damaging them.

This section outlines the calculations done to minimise problems due to thermal expansion. Table 3.4 shows the difference in linear thermal expansions in various materials. Generally, loosening of screws upon cooling down from T_{room} to T_{cryo} can be avoided by choosing different materials. Alternatively, washers made of an appropriate material can be applied. In most cases for this experiment, screws made of SS 316 are used to fasten a piece of OFHC Cu. Comparing the values from table 3.4, this means that by default the screw's applied force would become lower when cooling down; if it becomes too low, the screw might slip, and repeated thermal cycling might loosen the screw too much. This can be combated by using a material of low thermal expansion, e.g. by using Molybdenum. A convenient way of introducing Molybdenum is by using M3 washers made out of the material, part 02-39-021 from CMR-Direct, with a thickness of 0.5 mm.

Material	$\int_{20K}^{300K} \alpha_L \mathrm{dT}$
Unit	$10^{-5}\ {}^{m}/_{m}$
Invar	40.2
Epoxy (Stycast 2850 FT)	45.69
Alumina [p. 68, 52]	60.
Sapphire	83.73
Mo	95.8
SS 304 & 316	310.9
Be Cu	328.8
OFHC Cu	337.
PTFE (Teflon)	2281.

Table 3.4: Integrals of thermal expansion coefficient α_L of various materials at 300 K and 20 K. Calculated from [p.12-210ff 12]

For example, the OFHC Cu braid between the pillbox to the cryostat needs a high thermal conductivity to efficiently transfer the heat. This is achieved by pressing the surfaces together; to first order the thermal conductivity is linearly dependent on the force any screws exert. Upon cooling down, contractions in the materials can be calculated via

$$\Delta L = L \int_{20K}^{300K} \alpha_L \mathrm{dT} \tag{3.11}$$

Using this formula, the OFHC Cu braid's screw hole length change ΔL_{Cu}, the SS 316 screw length change ΔL_{SS}, and the Mo washer thickness change ΔL_{Mo} can all be calculated. The important figure in the end is the differential expansion of the three quantities combined, ΔL_{total}; more specifically its polarity:

$$\Delta L_{total} = (\Delta L_{SS} + \Delta L_{Mo}) - \Delta L_{Cu} =$$

$$= \left(L_{SS} \int_{20K}^{300K} \alpha_{L,SS} \mathrm{dT} + L_{Mo} \int_{20K}^{300K} \alpha_{L,Mo} \mathrm{dT} \right) - L_{Cu} \int_{20K}^{300K} \alpha_{L,Cu} \mathrm{dT} \approx +0.2\,\mu\mathrm{m}$$

In this case, two washers of 0.5 mm thickness, yielding $L_{Mo} = 1\,\mathrm{mm}$, are added to get the overall differential result to be a slight tightening upon cooling. Using one washer only would result in an expansion of -9 μm, while no washers would result in -20 μm, both of which would mean a slight loosening of the screws. The same method of adding Molybdenum washers is used for screwing the OFHC Cu braid onto the cryostat, and the pillbox onto its legs. Notably, the SS 316 screws that hold the alumina PCBs onto the pillbox tighten by 2 μm without the use of washers.

3.6.5 Helium Storage

Liquid helium is stored in a Helistor 250 dewar, part 9914335 from Statebourne, with an eponymous nominal capacity of $V = 250\,\mathrm{l}$ of liquid helium. It is specified to have an empty weight of 168 kg, a static evaporation[10] of 0.9 % per 24 hours, i.e. $\dot{V}_{evap} = 250\,\mathrm{l} \cdot 0.09\,\%/\mathrm{day} = 2.3\,\mathrm{l}/\mathrm{day}$, and a maximum overpressure of 0.5 bar. It has certified lifting points used for transporting it to and from the lab. Based on the LHe flow rate of $\dot{\mathrm{V}} = 0.7\,\mathrm{l}/\mathrm{hr}$ calculated in sect. 3.6.3, one dewar should last for about $t = \frac{V}{\dot{V} + \dot{V}_{evap}} = 13$ days. As filling a dewar takes several hours, two are used to be swapped out when needed. This dewar's height, added to the height of a top entry transfer line, is higher than the laboratory's ceiling. While it could be inserted outside of the laboratory, this could result in problems e.g. when parts of the cryogenic system ice up. Therefore, a side entry transfer line is used for which a mating port on the dewar is needed, part 0700122 from AS Scientific Products Ltd. This side entry mating port is already included in the

[10] While the experiment is not running, the static evaporation is exhausted via a dedicated line to be recycled. A ball valve is used to close this line otherwise.

dewars.

Fluid level is measured with a helium depth indicator, model HDI-AS from Twickenham Scientific. It uses a resistive measurement across a superconducting probe inserted into the dewar; only parts submerged in the liquid become superconducting. This gives a height reading, which can be converted into a volume reading using the dewar manufacturer's conversion chart. However, the pressure inside the dewar changes this reading, so an empirical model could be used to convert readings to an equivalent reading at 1 atm. The output of the helium depth indicator is logged using its RS232 interface.

3.6.6 From LHe To Exhaust

This section explains the liquid helium cooling system, see fig. 3.13. As previously mentioned, the LHe is stored in a dewar, with the manufacturer's overpressure release valve set to 0.5 bar[11].

The liquid helium is transported from the dewar to the cryostat using a flexible, vacuum insulated, custom transfer line from AS Scientific Products Ltd. Inside the cryostat, the liquid helium boils off, and increases in temperature, removing the heat from the experiment. The helium is then transported via a custom 2 meter long flexible, vacuum insulated, $\varnothing$10 mm inner diameter KF25 hose, also from AS Scientific Products Ltd. At this stage, the now gaseous helium is still very cold, presumably only a few degrees above the cryostat's set-point. A rough calculation based on the calculated LHe flow results in a heating requirement of $\approx$ 50 W to bring the helium back to room temperature for recycling. Therefore, it is subsequently heated by passing through $\approx$ 7 m of coiled $\varnothing$10 mm Cu piping, from Yanik Brewing Supplies, submerged in tap water inside a water boiler, part MFCT1010 from Burco. The water boiler has a rated power of 1.6 kW, i.e much higher than the heating requirement, but can be set to keep the

[11]Note that this section follows the convention that the pressures are taken relative to 1 atm, i.e. a relative pressure of 0.5 bar relates to an absolute pressure of 0.5 bar + 1 atm$\approx$1.5 bar.

water at $\approx 30°C$, to keep the water from freezing. Both to reduce corrosion, and to limit bacterial growth, a "central heating protector" is used, part 77218 from Screwfix Direct.

After the heating system, a junction splits the path; firstly to a KF16 vacuum valve, which is usually closed; its use is explained later in this section. Secondly, to another KF16 vacuum valve, which is usually open, and to a check valve. The helium takes the second path, subsequently passing through a flow controller, part 4850ABB2T1B1K3C00A from M J Wilson Group Ltd, which measures the flow rate via the increase in temperature of passing gas due to a small internal heater, and knowing which gas is used. It also features a sphincter valve to regulate the flow. The controller's RS232 interface is used to set the flow to either fully closed or an arbitrary set-point in the range of $0.6 - 30\,{}^{l}/_{\mathrm{min}}$[12]. The helium is then, via a another check valve, and a gas meter, returned to the building's helium infrastructure for re-liquefaction. All bits of KF plumbing are made using LewVac Components Ltd.'s range, e.g. FL-CRSV40KF. Where overpressure is expected, overpressure rings are used.

The LHe consumption rate can thus be tracked in two ways. Firstly, via the fluid level measurement from sect. 3.6.5. Secondly, via the gas meter. This can then be compared to the calculated heat budget, see table 3.3. The process for starting the LHe flow is explained in the following paragraphs.

Firstly, the dewar's needle valve and the KF16 vacuum valve leading to the flow controller are both closed. Then the second KF16 vacuum valve is opened, and connected to a diaphragm backed turbomolecular pump, parts A74602983 and EXT 70H, regulated by a controller, part EXC 120, next to which a gauge is used to measure the pressure, part D14701000, all from Edwards. The system is then evacuated for ≈ 30 min to both remove any air and moisture, which could otherwise cause a blockage upon contact with LHe. The pump is then disconnected, and the two KF16 vacuum valves opened and closed, respectively.

[12]This is entered in conditions of a temperature of 0°C and pressure of 1 atm.

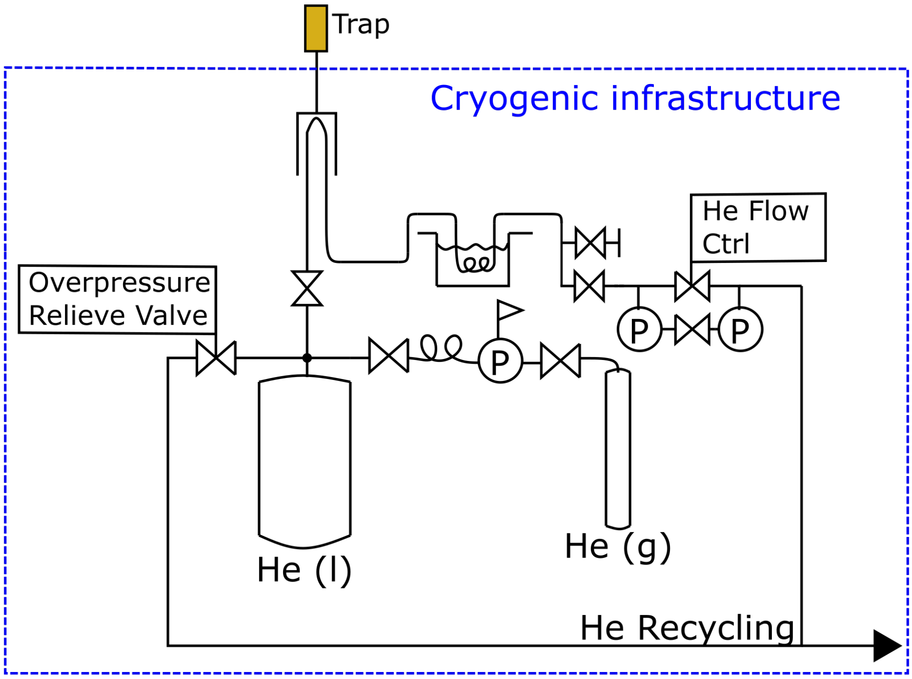

Figure 3.13: Schematic overview of the liquid helium cooling system. The cryogenic infrastructure consists of a liquid helium dewar, “He (l)”, that can be overpressurised with the help of a gaseous & compressed helium cylinder, “He (g)”. The “bow ties” denote valves, and the encircled “P”s denote pressure gauges. The liquid helium is led along a heat exchanger, cooling the trap, after which it is warmed up - the pipe containing the helium is led through a water bath. Now at room temperature, the helium flow is restricted by a flow controller, which creates a back pressure, ultimately restricting the amount of liquid helium exiting the dewar. The helium is then recycled.

Secondly, initial overpressure in the LHe dewar is generated via a compressed helium cylinder, regulated via part 3960002626 from Express Weldcare Services Ltd, with a maximum outlet pressure of 1.5 bar. The regulator is set to ≈ 1 bar, and a ball valve on the inlet to the dewar is opened until the pressure in the dewar reaches the desired pressure of e.g. 0.3 bar, as read out by an analogue pressure gauge on the dewar. Once this pressure is reached, the valve and regulator are turned off again, as the subsequent running of the experiment creates enough heat load to keep the dewar pressurised.

Finally, the dewar's needle valve is fully opened, and the flow controller set to its maximum flow. The pillbox' rate of temperature change, as read out by the cryostat's diode, is monitored; the flow controller's set-point is used to adjust this rate.

The temperature can be influenced in a few ways: Firstly, the cryostat's heater can be used for "fast" stabilisation. Secondly, the flow controller can be used for "slow" stabilisation. And finally, the dewar's pressure can be manually changed.

Due to the nature of the cryostat's design, its surfaces that are exposed to air get cold enough to grow layers of snow. This becomes problematic when some of the snow falls off, and when the cryostat is warmed up again; water damage can occur in the vicinity of the cryostat. To avoid this, polyethylene foam, part CF09X15X1M from Ryan Air Conditioning Spares Ltd, is used as a first layer, followed by aluminium foil, part UHV-001500 from All-foils Inc, sealed with sticky aluminium tape, part 1517CW-H746 72MMX50m from 3M.

3.7 Electric Connections

The base flange beneath the pillbox houses all the electrical connection of the ion trapping chip to outside the vacuum chamber. As a base, it uses a Testbourne CF100 flange, onto which the electrical feedthroughs are welded. DC connections are made via a 51 micro-D connector, part 23053-01-W, RF connection is made via an SMA to pin feedthrough, part 15310-03-W, microwave connections are made via four SMA feedthroughs, part 15263-03-W, see 3.14.

From outside the vacuum chamber, the microwaves are transmitted via an SMA cable, through the SMA feedthrough. In vacuum, a semi-rigid coaxial cable transmits the microwaves from the feedthrough to the outside of the pillbox. Therefore, the materials are chosen to minimise heat conductivity while minimising electrical resistivity, see 3.6.3. The chosen cable is a custom SMK one from Allectra, part 390-SMK-MM-100-MCOA; by diameter, the centre conductor is 90% SS, 10% Cu, with a thin Ag-plated finish, while the outer conductor is SS, with PTFE in between. High quality SMA and 2.92 mm connectors, aka. SMK, are inter-mateable [45].

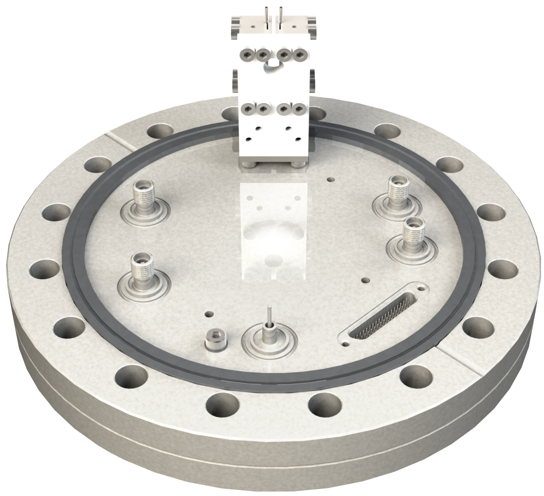

Figure 3.14: Base flange providing electrical connections and mechanical support to the pillbox. The atomic oven is also shown here, highlighted by its white MACOR ceramic. It has a hole in the centre for allowing beam access; the oven's tubes are slightly offset. Rendered with Solidworks Professional 2016.

3.7.1 Reverse Engineering of SMK Connector

The cable transmits signals onto an SMK connector, part 1014-25SF from Southwest Microwave. Internally, these connectors have two plastic parts, made of

PCTFE and ULTEM 1000, which is problematic because their outgassing rates are not known to be low enough for ultra high vacuum usage. They are therefore replaced with PEEK versions using the following method:

In a beaker, the Ultem 1000 is removed with the use of N-Methyl-2-pyrrolidone (NMP) at 50°C. The inner pin is then loose and can be removed. The PCTFE ring is still attached, but can be cut off with a scalpel. They are then replaced with PEEK versions of the same dimensions. Using a series of custom designed tools, the metal plug of the SMK connector is pushed out, the inner pin with PEEK disks inserted, and the metal plug pushed back in.

As the above method destroys the plastic disks that are to be removed, they could not be measured to get the dimensions for the PEEK disks. This was done via microscope images of an unmodified connector, partial information from the manufacturer's diagram, and measurement of a connector cut into two halves. To make sure this replacement of dielectric materials did not adversely affect the performance of the SMK connector, its S-parameters are measured, see fig. 3.15.

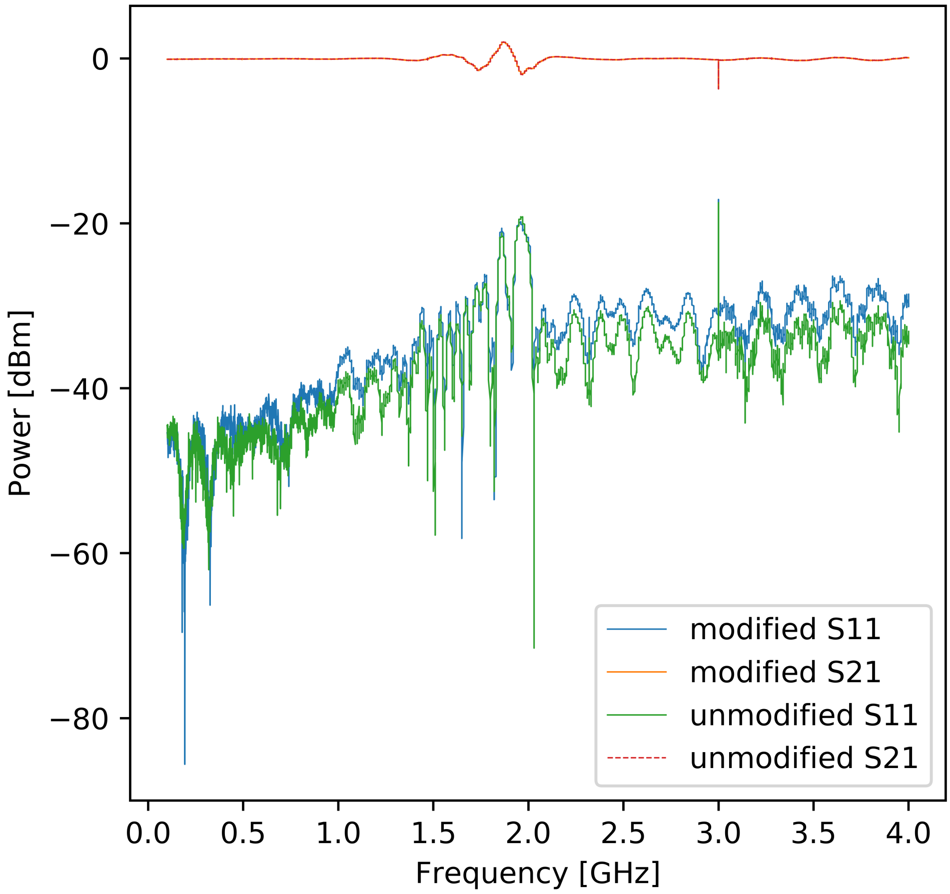

Figure 3.15: SMK connector S11 & S21 parameters. The modified connector shows slightly higher S11 parameters, although both show acceptable performance of ≈−30 dB S11 reflections for the region of interest between 3 and 3.5 GHz. The S21 parameters of the modified vs unmodified versions are indistinguishable on this scale. The spike at 3.0 GHz and the “wobble” of the S21 parameters between 1.5 and 2.0 GHz are assumed to be a measurement artefacts. Taken with a Rohde & Schwarz ZVL13 Vector Network Analyser.

3.8 PCB Designs

Inside the vacuum chamber, four types of PCBs are placed. They have to be low outgassing and resistant to the UHV cleaning process, see sect. 4.4.5, capable of withstanding a wide temperature range of 4K to 570K, and able to transmit microwaves without too much loss. This severely limits the options of materials. All the following PCB designs are manufactured by CMS Circuit Solutions, with 1.02 mm thick, 96% purity Alumina. It has a dielectric constant of ≈9, and a loss tangent of 0.0002 [10].

All resistors used have a resistivity of $R = 1.00 \pm 0.01\,\mathrm{k}\Omega$, part CHR 0603-HT TK50 HT from SRT Resistor Technology GmbH, while the capacitors have a capacitance of $C = 1\,\mathrm{nF}$. They are soldered outside a cleanroom using SAC 305[13], part 24-7068-7603 from Kester. Unfortunately, ordinary leaded soldering alloys, e.g. 60% Sn, 40% Pb, are not UHV compatible due to Pb outgassing during the bakes necessary to achieve UHV conditions, see sect. 4.4.9. While the used SAC 305 includes a flux, the soldering process is still difficult compared to ordinary solder. Also, the used Alumina has a high thermal conductivity one to two orders of magnitude higher than a typical PCB material like FR4, making localised heating more difficult.

To facilitate soldering, a hot plate is used, part SP88857205 from Fisher Scientific, to elevate the PCB's temperature to just below the solder's melting point, $T_m = 220°\mathrm{C}$, for example to $T_{HP} = 200°\mathrm{C}$. Subsequently, a combination of a hot air gun, part 2311402 from Tenma, and a soldering iron, parts PU 81 & WS 81 from Weller, are used, for which two people are required. One person further preheats a component with the hot air, and the second person holds a component in place with tweezers and applies the SAC 305, melting it with the soldering iron. Subsequently, the flux is cleaned off, and the usual cleaning steps as described in sect. 4.4.5, are followed, although without ultrasonic agitation to avoid damaging

[13] An alloy of 96.5% Sn, 3% Ag, 0.5% Cu. All alloy concentrations are given as mass fractions, not volume.

any components.

They are wire bonded onto the ion trap chip, see sect. 4.3.1, and fastened with M1.6 x 6mm screws made of SS 316, part SSCF-M1.6-6-A4 from AccuGroup.

For all PCBs, the conductor layers are Au by default, part C5756 from Heraeus, which is Au wire bond-able, but not SAC 305 solder-able; for those pads AuPtPd is used, part C6029A from Heraeus. To decrease the cost, the bottom ground layers on the RF, μW, and T PCBs are made of Ag instead of Au, part 6111 from Koartan. The schematics of these PCBs can be found in fig. 3.16, while their routing can be found in fig. 3.17.

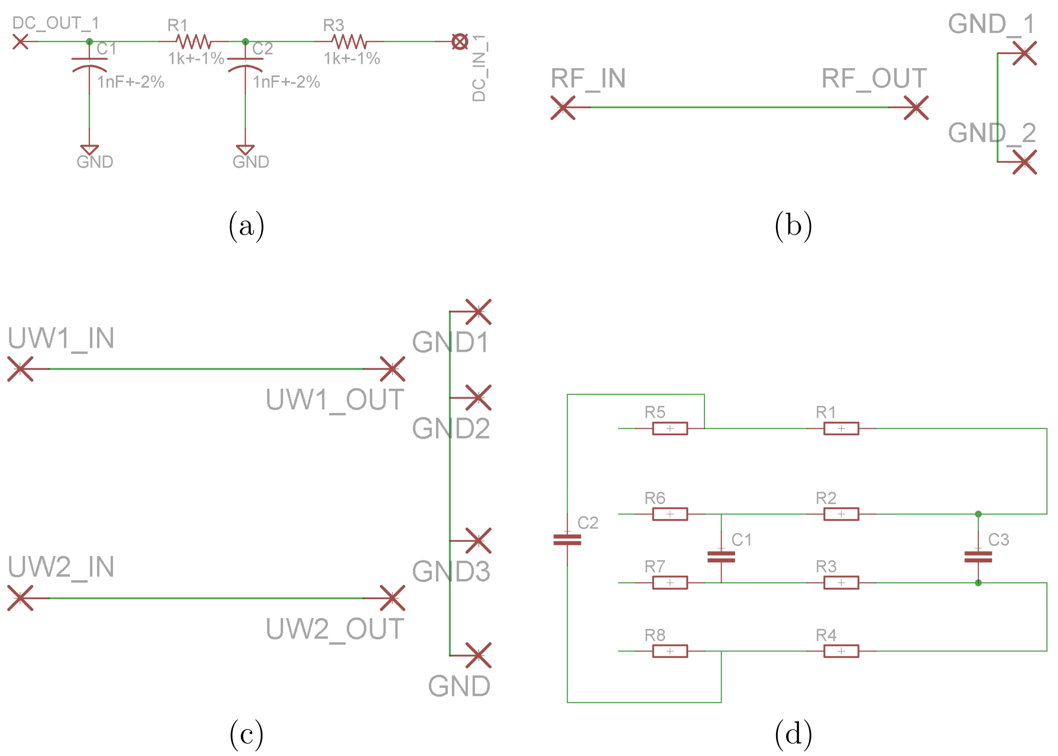

Figure 3.16: Schematic overviews of the in-vacuum PCBs. (a) For the DC PCB, all 11 channels have the same schematic, and only one channel is shown. (b, c) The ground pads on the right are the pads used to wire bond the grounding onto. (c) The μW PCB features a spare μW-line, which was initially needed for some gate schemes under consideration, but the final trap design and corresponding gate scheme only needs one μW-line. (d) The diode in the T PCB is not shown, but connected to both sides of the capacitor C3. Created with EAGLE 6.6.0.

DC PCB The DC PCBs are used to transfer the DC signals to the ion trapping chip. To reduce the noise arriving at the DC pads, this PCB features

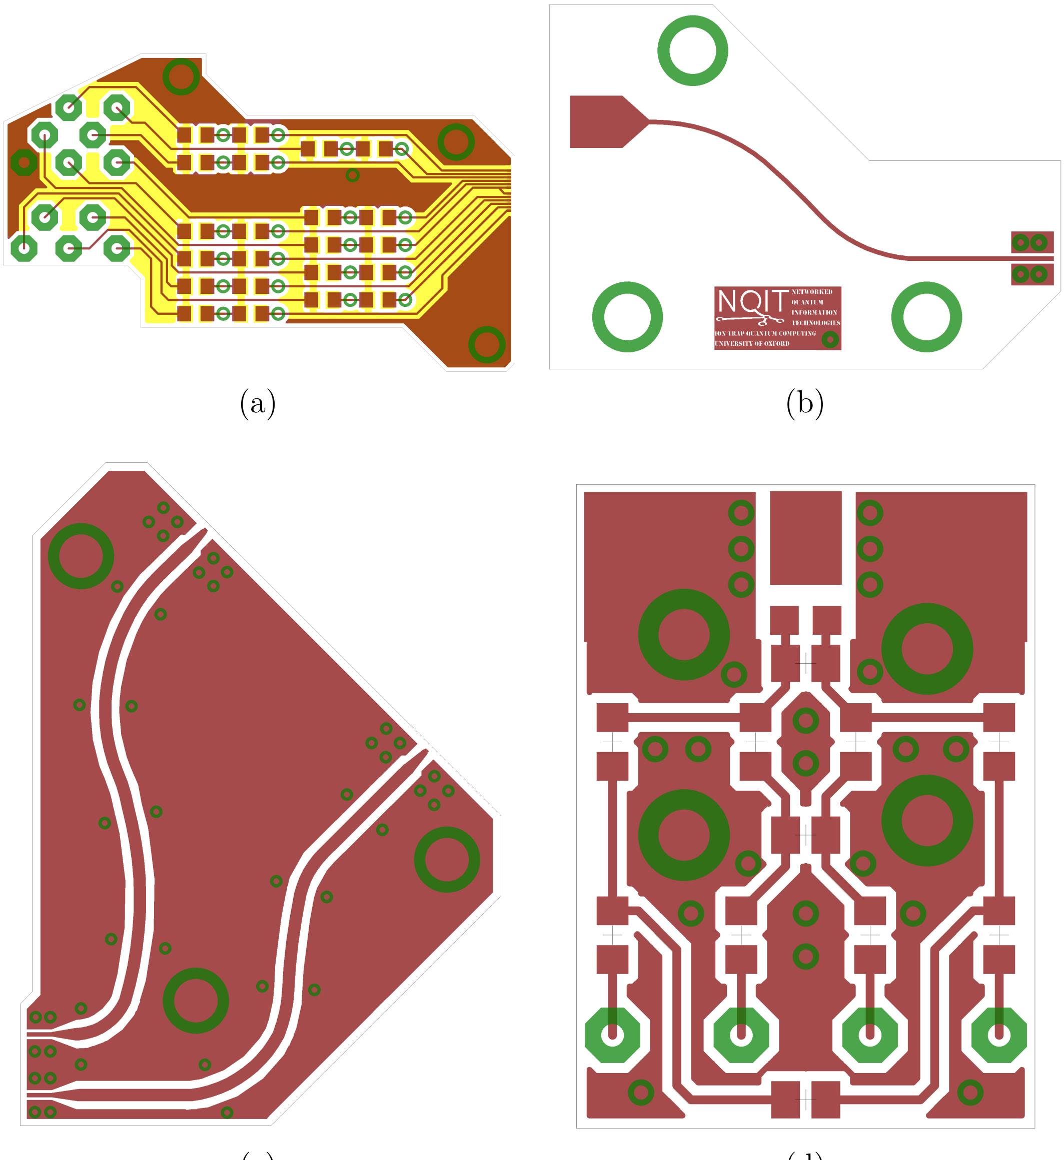

Figure 3.17: Routing overviews of the in-vacuum PCBs. The layers are: red (top conductive layer), yellow (bottom conductive layer), and green (outlines of vias). The bottom conductive layer is only shown for (a), the DC PCB; the other PCBs have a simple grounding layer comprising the entire bottom. For photographs of these PCBs on the pillbox, see fig. 4.19c. Created with EAGLE 6.6.0.

an 11-channel 2$^{\text{nd}}$-order low pass filter, made of two R-C low pass filters in series. The frequency of a 3 dB drop in transmitted power is $f_{3\text{dB}} = (1/2\pi RC) \cdot \sqrt{\sqrt{2}-1} = 100\,\text{kHz}$ [5], above which the attenuation closely matches $-40\,{}^{dB}/_{\text{decade}}$. In particular, the motional frequencies of $f_{\text{ion}} = 1\ldots10\,\text{MHz}$ should therefore be well filtered. To achieve the electrical connections, sub-D high density crimp pins are used, parts 212-PINFHD-26 and 212-PINMHD-26 from Allectra. The male versions are placed with their pins facing downwards, through corresponding holes in the PCBs, and soldered on. As their upward facing crimp end is not needed, it is removed with side cutters for increased optical access. Their female counterparts are crimped onto the end of the DC lines, see sect. 4.3.2, using a crimp tool, part 214-CTOOL-HQ from Allectra. The two ends are then mated. Optionally, they can be insulated from each other using Kapton tubing, but this is found to be unnecessary as the mechanical stability of the mate is sufficient to deter movement, even in a subsequent bake and cryogenics.

RF PCB This PCB transmits the RF signals to the ion trap. A microstrip transmission line is used here, i.e. the ground for the line is on the bottom side of the PCB. The optimal width for a characteristic impedance of 50 Ω at a frequency of 64 MHz is calculated using TX-LINE, a transmission line calculator [32]. After production of the PCBs, further simulations of the ion trap chip indicated that more fixed boundary conditions on the RF line are needed. Therefore, an air[14] core inductor is formed out of a Cu wire, which should be attached as close to the ion trap chip as possible. Unfortunately, the transmission line is made of pure Au, i.e. soldering with SAC 305 is not possible. Therefore, a eutectic solder alloy of 80% Au, 20% Sn is used instead, see sect. 4.2.2.

[14]Or more precisely, a vacuum core

Due to the higher melting temperature of this solder, $T_m = 278°\mathrm{C}$, soldering is a bit more difficult than with SAC 305; however as this alloy does not need flux, it can be soldered in-situ, without any further cleaning. The eutectic solder is also used for attaching the incoming RF wire onto the PCB, which would be hard to remove for cleaning.

μW PCB This PCB is used for transmitting microwave signals from the edge of the pillbox to the ion trap chip. It has two channels for microwaves – they were intended for driving single and two-qubit gates as earlier designs of the ion trap chip needed the second channel. However, only one is now unused, and grounded to avoid any floating potentials. Using the dielectric constant and loss tangent of Alumina, a coplanar waveguide, with ground on the opposite side of the PCB and a characteristic impedance of 50 Ω for 3.2 GHz is created. The PCB includes many plated through holes to ensure good grounding and avoid any parasitic other modes, e.g. slotline modes. All calculations for transmission lines on PCBs are made with TX-LINE.

T PCB For measuring the temperature of the pillbox, this PCB is intended to be used. To be useful in all situations, temperature range that should be measurable is $T = 4 \ldots 320\,\mathrm{K}$, while it needs to be able to withstand $T = 4 \ldots 460\,\mathrm{K}$. This, along with the usual UHV material restrictions, limits the choice of temperature sensors[15]. Diodes as temperature sensors are generally not recommended to be used in a magnetic field or when RF signals are present; the latter due to their tendency to rectify the signals, creating a systematic error. However, the magnetic field used in this work, 28.8 mT, should be small enough

[15] A typical choice of temperature sensor for similar purposes are platinum resistive temperature devices (RTDs). While their sensitivity is good at LN_2-temperatures, it drops significantly below 30 K, and becoming unusable below 15 K; they would therefore only be marginally acceptable for measuring the intended pillbox temperature of $\approx 16\,\mathrm{K}$.

to be negligible, particularly as the diode's axis is arranged in a 45° angle to it, reducing the effective magnetic field. To counteract the RF signals, a 2nd-order low pass filter is used, similar to the DC PCBs. The chosen diode, part DT-670-SD-1.4L from Elliot Scientific Ltd, is contacted in a four-wire arrangement, to be read out by a temperature controller, Model 335 from Lakeshore, which would feed back onto a heater in the cryostat's head. Here the low pass filter's capacitors are not grounded on one side, but rather connected pairwise to the current, and voltage carrying wires. Also, any systematic offset due to the RF signal and magnetic field could be measured by briefly turning them off; subsequently the measured offset could be subtracted. Unfortunately, this PCB was stored in air for ≈ 1 month before attempted vacuum assembly, at which point it failed electrical testing it passed before storage; presumably due to oxidation. As measuring the pillbox's temperature can also be done more indirectly using the cryostat's temperature diode, assembly proceeded without this PCB. At the next opportunity, a T PCB is intended to be inserted; this is likely to be taken from the second batch of PCBs as explained below.

Only a small batch of PCBs, 6 DC, and 3 each of the others, are ordered before assembly takes place, and the entire batch is used up in testing and assembly. If for any reason the ion trap chip needs to be switched out, the wire bonds would need to be removed, which could damage the PCBs' wire bonding pads. To allow for a change of trap, a new, bigger batch of PCBs, 40 DC, and 20 of each of the others, was ordered. These include some design changes to incorporate the lessons learned from the first batch:

- Molten SAC 305 solder consumes pure Au on contact, leaving no conductive layer in its place. In particular, any small droplets created from soldering in one location can splash over to another, breaking a connection. This increased the difficulty of soldering components to the PCBs. To counteract this, all areas that are not intended to be solder-able are covered in a thin layer of Alumina dielectric. As the spare Au layer was used for training on the wire bonder before bonding onto the chip, a small area on the μW PCB is left without covering dielectric.

- Given the general difficulty of soldering SAC 305, the two pole filter of the DC PCBs[16] is switched out for a single pole to reduce the number of components to fit onto one side.

- As the RF PCB needs an inductor, a corresponding solder pad is introduced; the new dielectric layer would make it difficult to introduce one otherwise.

- The RF PCB is extended, and its wire connection is swapped out for a crimp connection analogous to the DC PCBs. This requires a change in the Pillbox design, namely a wall to be removed.

- The μW PCB's unnecessary second microwave line is removed, and the launch pin solder pads enlarged for increased solder-ability.

The schematics of these version 2 PCBs can be found in fig. 7.3, while their routing can be found in fig. 7.4.

3.9 Magnetic Field Coils

3.9.1 Main Field Coils

The chosen qubit needs a magnetic field of 28.8 mT. It is also useful to have additional magnetic field coils in the other two axes to cancel out stray magnetic fields, including the earth's magnetic field, as well as correct small misalignments

[16] Due to difficulty in obtaining more of the same capacitors, a different model is chosen, part GRM1885C1H102FA01J from Murata Electronic

of the main coils with respect to the σ_+-polarised laser beams needed for SPAM, see sect. 2.7.

The vast majority of the magnetic field is supplied via the main field coils. They are capable of creating both 14.6 mT and 28.8 mT for use with $^{43}Ca^+$ ions. They are foil wound copper coils, and are specified to create 14.6 mT with a voltage of 1.5 V, and a current of 70 A, while 28.8 mT takes 3 V and 137 A, made by Stangenes Industries Inc. This current is supplied by an Agilent 6671A power supply, capable of a voltage of 8V at a current of 220A, connected via an insulated cable with a $70\,\mathrm{mm}^2$ copper cross section, part 51105 from Foster & Company Weldtech Ltd.

At 28.8 mT, the coils should then generate a heat load of $P = I \cdot V = 400\,\mathrm{W}$. Concerning water cooling of the coils, the manufacturer recommends a flow rate of $Q = 1.9\,{}^{l}\!/\mathrm{min}$, so the change in the temperature of the water should be $\Delta T = \frac{P}{c_p \cdot Q \cdot \rho} = 3°\mathrm{C}$, using the specific heat capacity c and density ρ of water, and assuming the coils only heat up the cooling water. This increase in water temperature is small enough not to cause problems.

Using a flow rate of $Q = 2.0\,{}^{l}\!/\mathrm{min}$, and measuring the temperature of the side facing away from the water cooling channel, which is expected to be the hottest part of the coils, a final coil temperature of $T_0 = 45°\mathrm{C}$ is measured. Starting from a thermalised water cooled state, and then turning the power supply on, the temperature of the coils increases as $T\,(t) = T_1 \mathrm{e}^{-\lambda \cdot t} + T_0$, with a decay constant of $\tau = {}^{1}\!/14\,\mathrm{min}$. Using a flow rate of $Q' = 3.5\,{}^{l}\!/\mathrm{min}$, the final coil temperature of $T_0' = 42°\mathrm{C}$, with a decay constant of $\tau' = {}^{1}\!/11\,\mathrm{min}$. As the additional vibrations likely have a negligible impact on the performance of the ion, the higher flow rate is used.

The field along the z-axis is specified to change by $0.35\,\mathrm{nT}$ over a distance of 20 μm when the field is 28.8 mT.

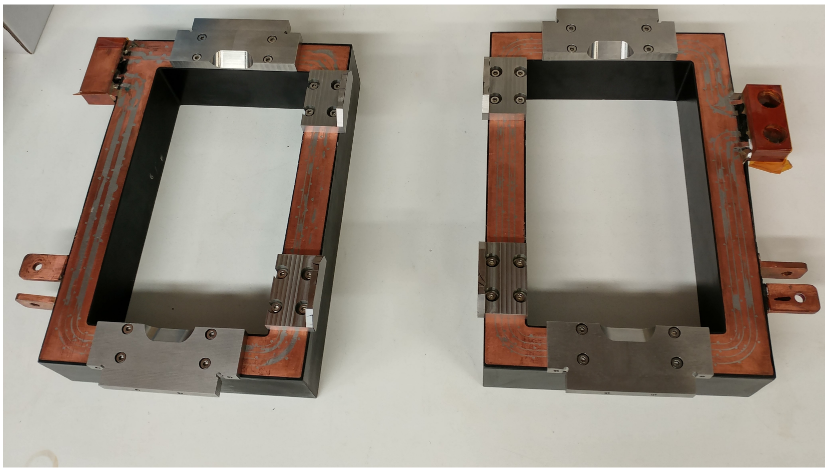

Figure 3.18: Main field coils. The custom SS mounting brackets are attached. They have female parallel 3/8" British Standard Plumbing (BSP) connections for water cooling. Electrical insulation is given by the black resin. Their size is maximised given the geometry of the rest of the experiment. Taken with an LG G5 camera.

3.9.2 Trim Coils

Through geometric considerations, for every 1° of misalignment of the main field coils at 28.8 mT with respect to the lasers and trapping fields, the resulting magnetic field is $\approx$ 0.5 mT. One degree of alignment precision is realistic to achieve for the main field coils. Therefore, to cancel out any extra field in the two perpendicular directions, two extra pairs of magnetic field coils are used, called trim coils.

For the first pair, a rectangular design is chosen, that is big enough to fit over the vacuum chamber. This means that the rectangular trim coils do not need to be able to withstand the bake-out process, sect. 4.4.9. In order to choose the type of wire and number of windings, and to get the resulting cross sectional area, a simple model is used.

At a distance of $d = 140\,\mathrm{mm}$, a single rectangular wire section, $l = 226\,\mathrm{mm}$ by $h = 0.137\,\mathrm{mm}$, creates a current and winding dependent field strength of $0.02316\,\mathrm{Gauss}/(\mathrm{Amp}\cdot\mathrm{turn})$. A cross sectional area $A = 5.25\,\mathrm{cm}^2$ is given by constraints on the vacuum system design. A combined power $P = 10\,\mathrm{W}$ is chosen to be dissipated in these coils to minimize the impact on the temperature of the vacuum system and the laser optics surrounding it. Using round wires, the best packing arrangement is hexagonal, with a density of $\eta_1 = 90.69\%$. Assuming that the insulation surrounding wires decreases the copper inside the wire to $\eta_2 = 75\%$, and using a wire outer radius $r = 1.25\,\mathrm{mm}$,

$$n = \frac{\eta_1 \eta_2 A}{\pi r^2} \tag{3.12}$$

n =73 windings are possible. At a resistivity of electrical grade copper at $\rho_1 = 16.8 \cdot 10^{-9}\ \Omega\mathrm{m}$, this results in a resistance per length of $\rho_2 = 13.7\,\Omega/\mathrm{km}$, and thus a total resistance

$$R_{tot} = 2n\,(2l + 2h)\,\rho_2 \tag{3.13}$$

$R_{tot} = 1.45\,\Omega$. A field of $B = 0.89$ mT can be created using a current of $I = 2.63\,\mathrm{A}$, and a voltage of $U = 3.81$ V, for which a QL335TP power supply from Thurlby Thandar Instruments Ltd. is used, specified to $I_{max} = 5$ A, and $U_{max} = 15$ V.

(a) (b)

Figure 3.19: Horizontal (a) and vertical (b) trim coils. They are terminated using plastic terminations rated to a current of 5 A. Taken with an LG G5 camera.

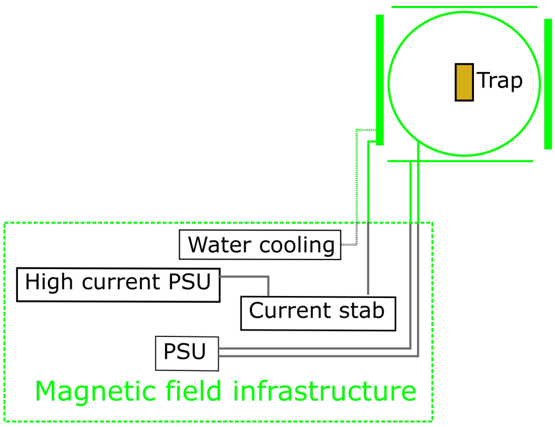

Figure 3.20: Schematic overview of the magnetic field coils. The thick lines surrounding the trap indicated the main coils, while the thin lines and the circle indicate the trim coils. A stabilisation circuit is planned to lower the noise of the high current power supply, although this is not yet implemented in this thesis. A water cooling system removes the heat from the main coils; this is not necessary for the trim coils. An interlock system (not shown) cuts off the high current power supply in case the water cooling fails.

3.10 Lasers

As explained in sect. 2.5, the lasers used in this work fall into two narrow wavelength ranges; [378 nm . . . 423 nm] and [850 nm . . . 866 nm], hereafter referred to as blue and red lasers, respectively, for simplicity. A schematic overview of the laser system is shown in fig. 3.21, and a more detailed explanation of it will be available in as yet unpublished work [36].

Except for the photoionisation lasers, all lasers are Toptica DL Pros, with integrated 30 dB isolators. Toptica controllers are used to stabilise the current through the diodes, voltage at the piezos, and temperature of the systems. They are mounted on 40 cm by 40 cm breadboards, which in turn are mounted in drawers of a 19 inch rack, greatly reducing the footprint of the lasers. On the boards, $^1/_2$ inch optics are used to distribute the lasers to up to 5 fibre ports each; a nearly arbitrary splitting ratio between these is achieved via a combination of $^\lambda/_2$-waveplates and PBSs. For fibre ports where the polarisation alignment is critical, i.e. fibres going to the experiment and to the PDH lock, an extra $^\lambda/_2$-waveplate is used to align to the slow axis of the fibre.

Generally, 3 fibres are used for each laser; one to go to the AOM breadboards, one to go to the cavities for a PDH lock, and one for diagnostics.

The diagnostics consist of a fibre switch which multiplexes all lasers into a wavemeter, WS/7 433 from HighFinesse, and two optical spectrum analysers (OSA); one for the blue and one for the red lasers. The traces that the OSAs produce can be used to detect whether a laser is lasing at a single frequency, "single mode", or needs adjusting of the diode current, piezo voltage, housing temperature, or grating angles. The wavemeter is used as a rough guide of the laser's frequency and for initial trapping; the more precise guide is the PDH lock to the cavities. The wavemeter can be calibrated using a HeNe reference for short-term accuracy.

The cavity arm consists of mode matching into the cavity, and a combination of a PBS and $\lambda/4$-waveplate to steer the reflection off the cavity onto a photodiode. For the blue lasers, phase modulation is achieved by using an EOM in this arm. For the red lasers the current is modulated. Using a Pound–Drever–Hall (PDH) lock module, Toptica PDD 110, the output of the photodiode is converted into an error signal, showing the deviation of the laser's frequency from the cavity's resonance. Using the lasers' controllers, this error signal is processed through a digital proportional–integral–derivative (PID) controller which uses the diode's current as a quick adjustment (but with a low travel range) of the laser's frequency, and the laser's piezo voltage as a slow adjustment (but with a large travel range). While tuning the PID loop, unwanted feedback onto the laser was found, and another set of 30 dB isolators were installed in series with the existing ones; no further unwanted feedback was detected afterwards.

The AOM breadboards feature a combination of single- and double-passed acousto-optic modulators (AOM), and one EOM. These are used both for switching the lasers on and off during e.g. a gate sequence, and to change the lasers' frequencies. The AOM breadboards are also packaged onto rack drawers. A further set of fibres takes the light from these to the experiment table, where they are combined and focused onto the ion.

In this thesis, laser intensities are given as multiples of a saturation intensity I_0, defines as the intensity where the rates of spontaneous and stimulated emission are equal.

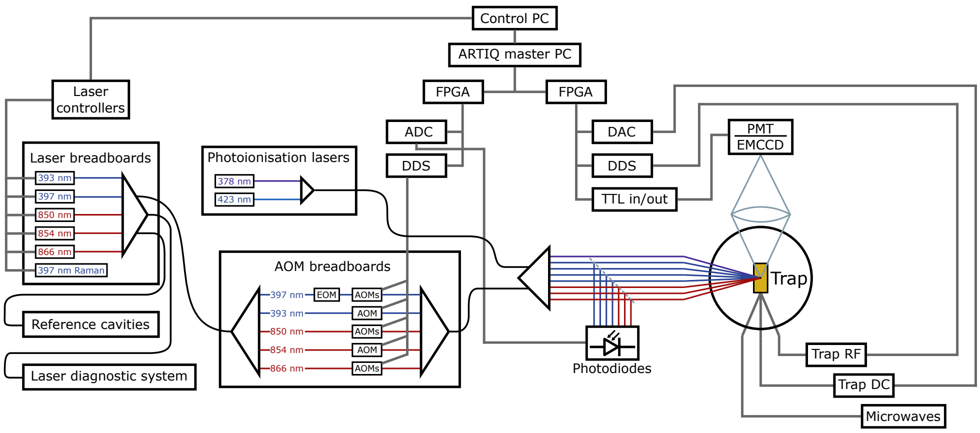

Figure 3.21: Schematic overview of the laser system. Reading from the photodiodes are used in a digital PID loop – comprised by the ADC and DDS – and fed back onto the AOMs to stabilise the lasers' amplitude. Taken from [36].

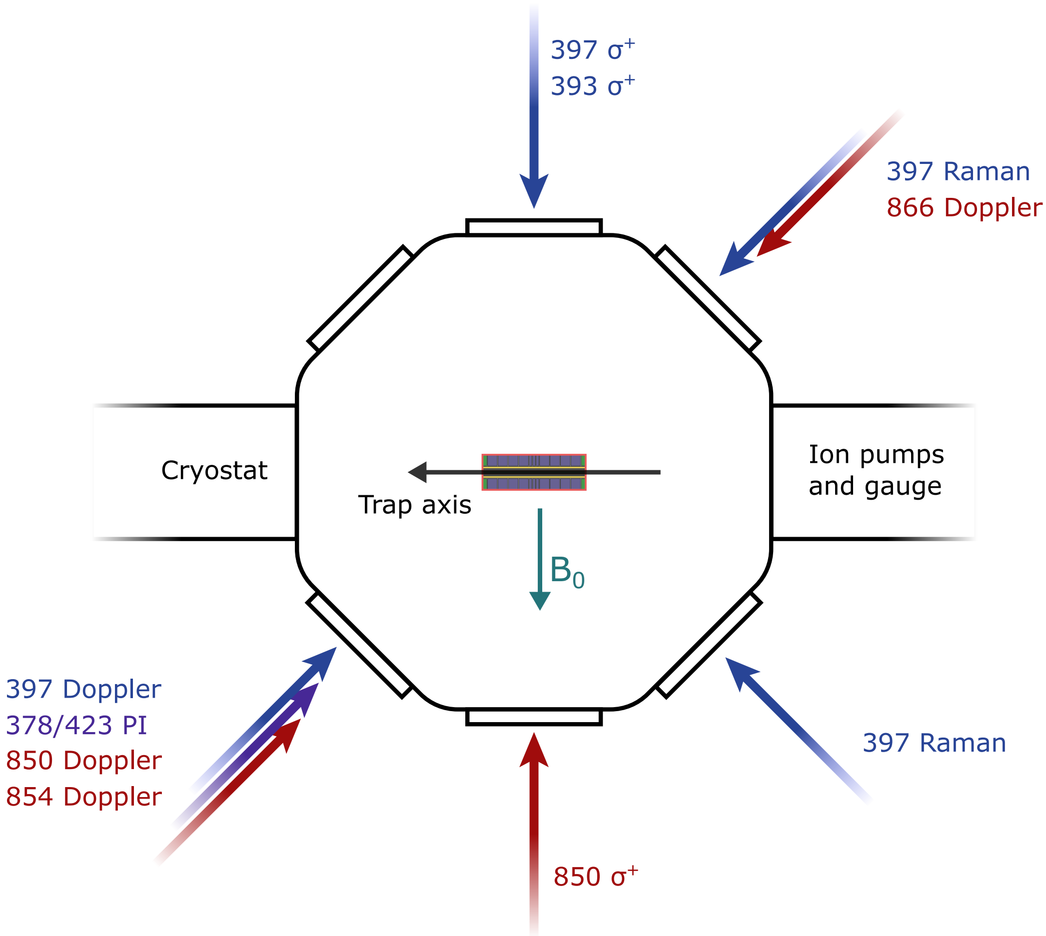

Figure 3.22: Schematic overview of the laser beam directions. At the time of writing, the σ^+-polarised 850 nm and the 397 nm Raman lasers have not yet been implemented. Taken from [36].

3.11 Experiment Control

The experiment is controlled by the control system infrastructure Advanced Real-Time Infrastructure for Quantum physics (ARTIQ). It is a Python based software package that can be used together with the Sinara hardware line to program field-programmable gate arrays (FPGAs) with nanosecond time resolution, allowing in- and output of a variety of signals.

A typical experiment is written in Python and has multiple parts; a *build* method that creates a user interface, a *prepare* method that gets run when the experiment is next in line to be executed, a *run* method that ensures this main part has exclusive access to the hardware, and finally an *analyze* function for any analysis that can be run while the next experiment is running. The *run* method has the option of running on the FPGA, which has real-time access to its attached hardware, but can also control other hardware over Ethernet. While only a subset of Python is available to be compiled to run on the FPGA, experiments written in this thesis are not found to be limited by the size of the subset, although sometimes a workaround is required.

A single Linux server runs the "ARTIQ Master". It executes any non-realtime operations of the experiments, and schedules the order of experiments. Any number of other computers, both Windows and Linux, can connect via "ARTIQ Dashboards", so multiple users can work in parallel; care is taken to ensure that one user's experiments do not influence another's.

The hardware that is intended to be run with ARTIQ is the Sinara family [8], which are PCBs using the Eurocard Extension Module form factor (EEM). These are then mounted in 19 inch racks. Used in this thesis are:

FPGAs Two "Kaslis", running in a master/slave configuration, linked by digital fibre optics. They distribute various clocks, all based on a single source - a synthesizer, part 8656B from Hewlett Packard, creating 100 MHz from a Rb frequency reference, model FS725 from Stanford Research

Systems, with an output of 10 MHz.

DAC A 32-channel "Zotino", which supplies the 22 DC voltages used for trapping; the remaining channels are unused. Notably, this limits available DC voltages to $\pm 10\,\mathrm{V}$, although this is found to be sufficient.

DDS These 4-channel "Urukuls" are used to create the $1 \ldots 400\,\mathrm{MHz}$ signals necessary for driving AOMs, which the lasers are routed through. As the DDSs are not powerful enough to directly drive the AOMs, an 8-channel RF amplifier, "Booster", is used. Notably, this RF amplifier has programmable power interlocks, greatly reducing the chance of accidentally damaging the AOMs. It also has power meters, facilitating diagnostics.

ADC One of these 8-channel "Samplers" is used to read the voltages of photodiodes in the lasers' paths. Two DDSs are combined with this ADC to create a digital servo controller, a "Sampler-Urukul Servo" or "SU Servo". This servo stabilises the laser's intensity by taking the photodiode's voltage and applying an error signal to the DDSs.

TTL An 8-channel "DIO_BNC" is used for any Transistor-Transistor Logic (TTL) devices, namely for the input of a counted photon from the imaging system, sect. 3.13, and the output of triggering microwave switches.

Other devices that are not of the Sinara family can still be used; the piezo controllers, part KPZ101 from Thorlabs Inc., used in this thesis can be plugged into USB hubs connected to the ARTIQ Master. Devices connected via Ethernet are also straightforward to integrate. Older communication standards, namely RS-232, RS-485 and GPIB, are adapted to Ethernet, using parts ES-246 and ES-320 from Brainboxes and gpib-eth from Prologix LLC, respectively.

Any results of run experiments are automatically archived, and can be retrieved using the experiment's run ID (RID). They can also be displayed on the ARTIQ dashboards. To help with diagnostics, miscellaneous data is also logged to a time-stamped InfluxDB database, which can be viewed on a browser via the Grafana dashboard. For example, SU Servo intensities are logged here, which show the effects of air conditioning affecting the fibres.

3.12 Microwave Drive Chain

In this work, microwaves are initially created using two synthesizers.

The two synthesizers, a Keysight N5181B MXG Analog Signal Generator and a Rohde & Schwarz SMA 100 A, are connected to two switches, both Mini-Circuits ZASWA2-50DR-FT+; the first to switch between the synthesizers, the second to switch the microwaves on and off[17]. This is then connected to a Mini-circuits ZRL-3500+ amplifier, specified for a gain of $g = 16 \pm 3$ across the frequencies of the transitions. The output of the amplifier is then led to the ion trap chip via the chain of feedthrough, semi-rigid cabling, and PCB introduced in sect. 3.7 & 3.8. To calibrate the output power of the synthesizers, the amplifier output is connected through a 30 dB attenuator to an RF power meter, Agilent V3500A. As the Rohde & Schwarz synthesizers can only supply a maximum of $p_1 = 15.00\,\mathrm{dBm}$, the Keysight synthesizer is used to balance the powers, and set to $p_0 = 15.83\,\mathrm{dBm}$, leading to a reading of $p' = -8.30\,\mathrm{dBm}$ on the power meter, i.e. $p = p' + 30\,\mathrm{dB} = 21.7\,\mathrm{dBm}$ without the attenuator. The synthesizers are also connected to the Rb frequency reference from sect. 3.11.

This setup is limited to only two microwave frequencies, which is sufficient for initial probing of the $4\mathrm{S}_{1/2}|4, +2 \ldots +4\rangle \to 4\mathrm{S}_{1/2}|3, +3\rangle$ transitions. As state-prep leaves the ion in $|4, +4\rangle$, the first - and sometimes only - transition probed in any

[17] While in principle, a single synthesizer could be used to both change its frequency, and switch itself, this Ethernet based communication would be orders of magnitude slower than the TTL signals that can be used for the switches.

case is $|4,+4\rangle \rightarrow |3,+3\rangle$. Therefore, the Keysight synthesizer is permanently set to $f_{4\rightarrow 3} = 3.198\,658\,286\,4\,\mathrm{GHz}$, and the magnetic field is changed to make this synthesizer's frequency resonant[18]. The second synthesizer's frequency is calibrated to be resonant to the next transition, which is influenced by the AC Zeeman effect. However, more frequency components are needed for reaching the qubit transition, see sect. 2.8.

To this end, a setup based on two channels of a direct digital synthesizer (DDS) and an IQ mixer, Eclipse IQ2040MP4, is used. The IQ mixer takes the input from the DDS channels and adds or subtracts their frequency $f_{\mathrm{DDS}} = 1\ldots400\,\mathrm{MHz}$ from a fixed, so called local oscillator frequency $f_{LO} \sim 3\,\mathrm{GHz}$. The proportion of addition to subtraction is governed by the phase relationship of the DDS channels; at a phase difference of $\phi = \pm 90°$, the IQ mixer adds (subtracts), resulting in an achievable frequency range of $f_{IQ} = f_{LO} \pm f_{DDS}$. As shown in fig. 2.7, the relevant transition frequencies are $f_{4\rightarrow 3} = 2.55\,\mathrm{GHz}$, $f_{3\rightarrow 3} = 2.67\,\mathrm{GHz}$, $f_{2\rightarrow 3} = 2.79\,\mathrm{GHz}$, $f_{2\rightarrow 2} = 2.91\,\mathrm{GHz}$, $f_{2\rightarrow 1} = 3.02\,\mathrm{GHz}$, $f_{1\rightarrow 1} = 3.12\,\mathrm{GHz}$, and $f_{0\rightarrow 1} = 3.22\,\mathrm{GHz}$[19]. To generate power at all the used frequencies with a constant local oscillator frequency, it needs to be $f_{LO} > f_{1\rightarrow 1} - f_{LO}$ and $f_{LO} < f_{4\rightarrow 3} + f_{LO}$; therefore, the range is $2.72\,\mathrm{GHz} < f_{LO} < 2.95\,\mathrm{GHz}$.

Three more considerations further narrow this range; in order to suppress anharmonicities in the DDS, it should not be set to a frequency approaching its maximum. Secondly, the IQ mixer has a non-negligible amount of power leakage at the local oscillator frequency; therefore this frequency needs to be ideally far away from any resonances. Thirdly, when generating power at a frequency $f_{IQ} = f_{LO} \pm f_{DDS}$, power at the opposite frequency $f_{opp} = f_{LO} \mp f_{DDS}$ is also generated with a finite extinction ratio between the corresponding powers $^{P_{OP}}/_{P_{IQ}}$, potentially driving further unwanted resonances.

[18] This also calibrates out the AC Zeeman shift induced on this transition, i.e. there is a small systematic offset in the magnetic field away from the hyperfine qubit's FOFI field. However, for the results achieved in this thesis, this effect is negligible.

[19] The last transition frequency is not intended to be used, but included in this list as it could accidentally be driven.

To satisfy the first two criteria, the acceptable range of local oscillator frequencies is narrowed to $2.79\,\text{GHz} \ll f_{LO} \ll 2.91\,\text{GHz}$, due to the resonances at $f_{2\to3}$ and $f_{2\to2}$. The spacing between the resonances is nearly constant, $\approx 110\,\text{MHz}$, so placing the local oscillator frequency exactly halfway between two resonances would mean that another resonance would be excited, too. To satisfy the third criterion $f_{LO} \neq \frac{f_{2\to3}+f_{2\to2}}{2} = 2.85\,\text{GHz}$. For future experiments, off resonant excitation could be modelled for typical gate sequences and an optimal local oscillator frequency chosen.

The final local oscillator frequency is chosen to be $f_{LO} = 2.875\,\text{GHz}$. While it is initially generated by a synthesizer, this frequency choice is allows the synthesizer to be replaced by a more economical frequency multiplier, using the $f_K = 125\,\text{MHz}$ clock from a Kasli master FPGA, to create the local oscillator frequency directly, $f_{LO} = 23 \cdot f_K$.

This results in an achievable frequency range of $f_{IQ} = f_{LO} \pm f_{DDS} = [2.475, 3.275]$ GHz[20].

The IQ mixer output is then run through a switch, Mini-Circuits ZASWA2-50DR-FT+, a 5 dB attenuator[21], and finally through two amplifiers in series, Mini-Circuits ZRL-3500+. The local oscillator is supplied by a MXG N5181B synthesizer, set to a power of $p = 17\,\text{dBm}$, connected to the IQ mixer via a 2m SMA, and then a 6 dB attenuator[22].

To calibrate out any phase and amplitude mismatches, a Rohde & Schwarz ZVL network analyser is connected to the output of the amplifiers (through a 10 dB attenuator). An algorithm is then used to vary the amplitude and phase of one of the DDS's channels at a particular frequency, and the power in the opposite frequency component P_{OP} is minimized. Typical extinction ratios of $^{P_{OP}}/_{P_{IQ}}$ are $50 \ldots 70\,\text{dB}$; higher extinction ratios would likely just require more time. Using ≈ 30 frequencies across $f_{IQ} = [2.475, 3.275]$ GHz, a function $P_{fit}(f)$

[20] Technically, the range of $f = [2.874, 2.876]$ GHz needs to be excluded from this range, as the IQ mixer works decreasingly well at lower frequencies. However this range is well away from any resonances and unused.

[21] This attenuator reduces reflections back into the IQ mixer which cause lower performance.

[22] Again, this attenuator reduces reflections.

can then be fitted to the corrective amplitudes and phases. A symmetric linear fit, $P_{fit}(f) = P_0 \cdot \mathrm{abs}(f - f_{LO}) + P_1$, produces extinction ratios of $30 \ldots 70\,\mathrm{dB}$ across the frequency range; as this is $\approx 20\,\mathrm{dB}$ lower than the extinction ratios achieved at the individual frequencies, a higher order fit function could therefore improve the extinction ratios, at the cost of additional complexity. However, an extinction ratio of $> 30\,\mathrm{dB}$ is sufficient for the work presented in this thesis.

3.13 Imaging system

Fluorescence from the ion is collected with a numerical aperture (NA) of $NA = 0.6$, i.e. a half-angle of $\frac{\alpha}{2} = \sin^{-1}(NA) = 36.9°$ and a solid angle of $\Omega = 4\pi \cdot \left(1 - \cos\frac{\alpha}{2}\right) = 0.40\pi$. Therefore, the fraction of light collected at this stage is $\eta_\Omega = {}^{\Omega}/_{4\pi} = 10\%$.

A lens system, part 16580 from Photon Gear Inc., is used to collect this light, referred to as the primary lens. It is factory customised to take the effect of the glass of recessed viewport, sect. 3.5 into account, as well as optimised for operation at a wavelength of $\lambda\left(4\mathrm{S}_{1/2} \rightarrow 4\mathrm{P}_{1/2}\right) = 397\,\mathrm{nm}$. A secondary lens, part AFL12-20-S-U from Asphericon, is used to collimate the light onto either a photomultiplier tube (PMT) or a camera, chosen by a movable mirror[23]. The PMT, part H10682-210 from Hamamatsu, has a specified quantum efficiency of $\eta_{PMT} = 31\%$. The combined loss from various mirrors and lenses is expected to be $\eta_R \approx 95\%$. To reduce the effect of any stray light hitting the PMT, a narrow bandpass filter is used, part FF01-400/12-25 from Laser 2000 UK Ltd, with a transmission efficiency of $\eta_{BP} = 96\%$. Therefore, the overall detection efficiency of a photon emitted by the ion is $\eta = \eta_\Omega \cdot \eta_R \cdot \eta_{BP} \cdot \eta_{PMT} = 2.8\%$, assuming perfect alignments of the optics.

The expected count rate at the detector Θ_q is

[23]For convenience, this mirror can be moved by an actuator signalled by an Ethernet-controlled Beaglebone computer.

$$\Theta_q = \eta \cdot A \cdot n \tag{3.14}$$

with the Einstein coefficient A and the state population n. As described in sect. 2.7, the $4S_{1/2} \rightarrow 4P_{1/2}$ is used for fluorescence detection, therefore $A = 132\,\text{MHz}$. Using an experimentally achievable $4P_{1/2}$ state population of $n = 7.6\%$, the count rate Θ_q should be $\Theta_q = \eta \cdot A \cdot n = 280\,\text{kHz}$.

The primary lens has its focus before the secondary lens; an adjustable slit, part SP 40 from Pro-Lite Technology, is used here to reduce the amount of surface scattered light arriving at the PMT.

Both lenses are on different types of 5 axis translation stages, the primary lens on parts M-GON65-L, M-GON65-U, and M-562-XYZ, and the secondary on part 9081-M, all from Newport Spectra Physics Ltd.

Chapter 4

Fabrication

The previous chapter introduced the design considerations of the experiment. This chapter is about the implementation of the design. The major pieces of work for this are the process development of ion trap chip fabrication and bonding processes, process development for wire bonding the chip onto the PCBs, and creating an infrastructure for baking out the vacuum system.

4.1 Ion Trap Chip

The wafers used as a basis for ion trapping chips are custom monocrystalline 99.996 % sapphire sourced from University Wafer Inc., with a diameter of 76 ± 0.3 mm, and a thickness of $430 \pm 25\,\mu$m. Their surface roughness on both sides is specified to be ≤ 0.3 nm.

In order to gold-plate wafers, an electroplating setup is designed, based on an unpublished design by NIST [53]. The modified version used here is shown in fig. 4.1.

It was difficult to get the electroplating right. In total, 38 wafers were processed in-house in the search for a suitable parameter set, plated with solutions from three different suppliers, the details of which can be found in the appendix, sect. 7.1; the final parameter set is explained in the following section.

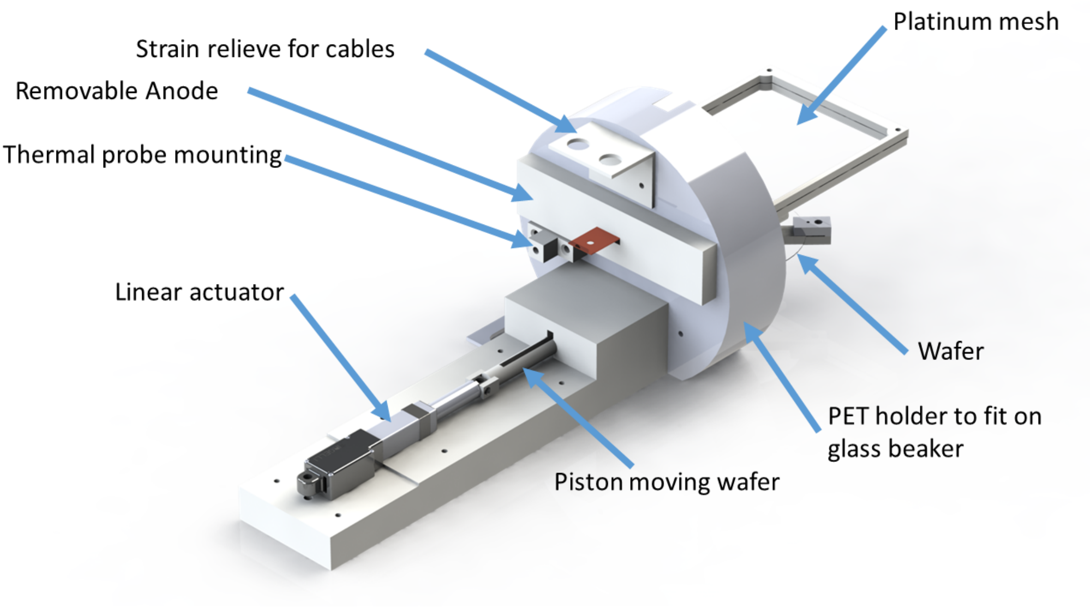

Figure 4.1: Electroplating setup. The bulk of the setup is made of PET (white). Attached is a linear actuator (black/grey), which moves a rod connected to the holder for the wafer (just visible on the right). In operation, this setup rests on a 1 litre glass beaker (not shown), which would attach from the right. A thermal probe (not shown) can be attached through the lid. Rendered in Solidworks Professional 2016.

4.1.1 Final Electroplating Process

The electroplating solution used in the final process is ECF64D, part 400005130 from Metalor Technologies SA, with a gold concentration of $15\,{}^{g}\!/_{l}$. The plating parameters adopted are a current density of $2\,{}^{mA}\!/_{\mathrm{cm}^2}$, applied for 31.6 min, at a temperature of $(50 \pm 0.1)^\circ$C, with a stirring speed of 400 rpm. This should yield an additional gold thickness of 4 μm on the wafer[1], based on the manufacturer's quoted plating rate of $1\,{}^{\mu m}\!/_{7.9\,\mathrm{min}}$ when using the above current density. As an anode, the manufacturer recommends platinised titanium. The exposed area of the wafer after patterning with photoresist is calculated to be 24.1 cm^2, therefore the current is set to $2\,{}^{\mathrm{mA}}\!/_{\mathrm{cm}^2} \cdot 24.1\,\mathrm{cm}^2 = 48.2\,\mathrm{mA}$.

The final electroplating setup[2] uses a constant current supply, part 2401 from Keithley Instruments, positive polarity connected to a titanium anode, plated

[1]The wafer is already coated with a 0.1 μm layer of gold, see sect. 4.1.2
[2]In order to be chemically inert, the bulk of the setup is made of PP, and PTFE

with 1 μm platinum, part TiPTA-01-MS-010 from William Gregor Ltd. Between the current supply and the anode, a 50 Ω ballast resistor is placed to hold the circuit resistance, as seen by the current supply, nearly constant. The other polarity is connected via a PTFE coated silver plated copper wire, part 2844/7 BK005 from Alpha Wire, which is crimped onto a gold plated ring[3], part A-RT18-22 from LewVac Components Ltd., which is in turn pressed onto the gold coated wafer via a clamp, screwed down via a hand tightened M5 screw made of PP. To minimise exposure of the uninsulated part of the crimped end of the wire, it is covered in a PTFE shrink tube series of increasing size, parts TFER-5/64-X , TFER-1/4-X, and TFER-5/8-X-STK from TE Connectivity, with initial diameters of 2 mm, 6.4 mm, and 15.9 mm, covering part of the ring crimp as well. Before insertion, the current supply is switched to a constant voltage of 0.01 V to create well-defined electrical conditions for the circuit upon wafer submersion. The current supply is then directly switched back to apply the plating current. Also before insertion, the wafer is cleaned using an O_2-plasma, although the details of the parameters used are confidential, see sect. 4.1.2. For all wafers processed in-house, an O_2 plasma barrel etcher, part PT7160 RF from Quorum Technologies, is used, at a pressure of $7 \cdot 10^{-1}$ mbar, power of 100 W, and applied for 1 min.

The temperature of the electroplating solution is controlled by an X2 Digital PID, part X2-240JT from BriskHeat Corp., connected to a silicon rubber beaker heater, part GBH1000-2 from BriskHeat Corp., wrapped around a 1 litre glass beaker. The temperature is read out by a PTFE coated type J thermocouple, part P2936-1EA from Sigma-Aldrich Ltd, connected to the temperature controller via two extension cables; one is used to strain relief the thermocouple to the electroplating setup, the other to connect the strain relief extension to the X2 Digital PID. To increase both the temperature and the plating uniformity, the

[3]As this metal part is already gold plated, further gold plating should not contaminate the bath

solution is also agitated using a reciprocal stirrer, part SI-0301 from Scientific Industries, which is set to reverse the direction of stirring every 60 s, which is fast compared to the plating time, to remove the effect of stirring direction on the plating thickness, using a PTFE coated stirrer bar sized 6 mm by 30 mm.

To smooth out any localised effects of geometry on the electroplating, the wafer is constantly moved, using a linear actuator, part L12-50-100-06-I from Firgelli Technologies Inc. It is powered and controlled via an Uno SMD Rev3 microcontroller board, part A000073 from Arduino, programmed to continuously move the attached wafer up and down. Additionally, it can be set to minimal and maximal extensions, useful for initial setup to make sure the wafer neither hits the bottom of the beaker nor is raised outside the liquid level, depending on the position of two switches. The board and switches are mounted in a box, and connected to the electroplating setup via a removable 9 pin D-sub cable. Photos of the finished setup can be seen in fig. 7.1.

4.1.2 Final Ion Trap Chip Fabrication Process

After establishing the parameters used for electroplating, the setup was demonstrated to ECS Partners Ltd, the consultancy arm of the University of Southampton, who then took over the whole ion trap chip fabrication process, but made the details of the chemicals and machines used confidential. The process, illustrated in fig. 4.2, is partially based on an unpublished recipe by NIST [53]. The final step of dicing is performed by the Optoelectronics Research Centre, University of Southampton.

4.1.3 Fabrication Bias

The ion trap fabrication process was anticipated to have some fabrication bias, which is discussed in this section.

Fig. 4.3 shows microscope images of the central region of a sample trap. The

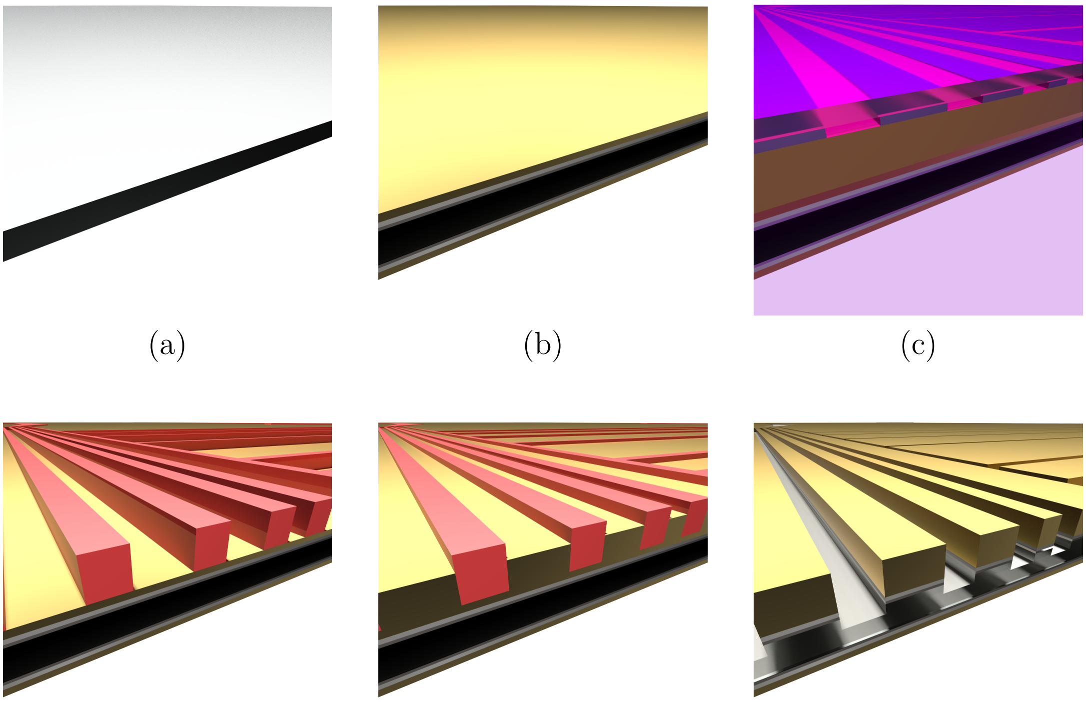

(a) (b) (c)

(d) (e) (f)

Figure 4.2: Ion trap fabrication process, lateral dimensions to scale, thicknesses not to scale. Shown here are cross sections in the sequential stages of fabrication. It starts with a cleaned Ø3 inch, 430 μm thick sapphire wafer (a), onto which metal layers, 40 nm Ti, 60 nm Pt, and 100 nm Au, are sequentially evaporated onto **both** sides using an e-beam (b). A 6 μm thick layer of photoresist is applied to the top side, and then selectively exposed to UV light through a mask (c). The photoresist is developed, washing away the parts previously exposed to the UV light (d). Electroplating is used to grow a 4 μm layer of gold in the gaps of the photoresist (e). The photoresist is washed away, and using argon ion milling, the metal in the gaps is removed, revealing the finished trap. (f). Subsequently, the wafer is cut into 5 mm by 10 mm pieces. Rendered with Solidworks Professional 2016.

scale of the image is derived through the use of a known magnification while working in focus, and is calculated by the microscope software. Measurements are taken to gauge the width of the various gaps and tracks, to compare them to the intended size. The edge of a track is taken as the line where the light intensity is half way between that of the gap and the track. In the HFSS simulation, the gaps in this region are intended to be 4.5 μm. A fabrication bias was estimated to make the gaps $\approx 1\,\mu$m smaller, so to counteract this, the gaps are printed onto the mask as 5.5 μm. However, the fabrication bias is subsequently measured to be only about $\approx 0.3\,\mu$m, i.e. the gaps are measured to be 0.7 μm bigger than intended.

The gold plating process introduces another parameter; the thickness of the plating, intended to be 4 μm; this would imply an overall layer thickness of 4.2 μm combined with the initial underlying seed layers[4]. Surface profiler measurements carried out by the Southampton Nanofabrication Centre measure the overall thickness of the layers to be thicker than anticipated, see fig. 4.4.

In summary, tracks on the traps are $\approx 0.7\,\mu$m narrower, and $\approx 0.7\,\mu$m thicker than designed, with peaks at the edge of tracks that are ≈ 1.9 μm thicker. As the trap design is symmetric around a central line, it should be fairly resistant against both of these kinds of fabrication bias.

[4]The details of the seed layers are explained in sect. 4.2.2

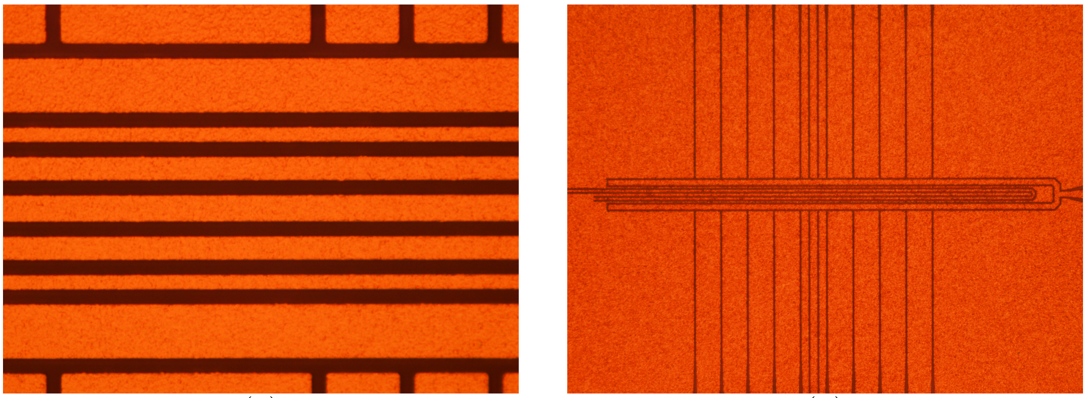

(a) (b)

CAD w/o Bias μm	CAD w/ Bias μm	uScope μm	uScope - w/o Bias μm	Description
4.5	5.5	5.2	+0.7	Gap
5.5	4.5	4.8	-0.7	OuterGND
9.5	8.5	8.8	-0.7	InnerGND
8.5	7.5	7.8	-0.7	μW
18.8	17.8	18.4	-0.4	RF
25.5	24.5	25.0	-0.5	DC 1/3
85.5	84.5	85.3	-0.2	DC 1/1

(c)

Figure 4.3: Microscope images of an ion trap. (a) Image of the central region, this image is used for measurements. Magnification is 50x, at an image width of 175.8 μm. (b) Overview of the whole trapping region, magnification is 5x, at an image width of 1770 μm. Both taken with an Olympus BX61. (c) Measurements of the main features of the ion trap. "CAD w/o Bias" are the lengths used in the simulation of the trap, i.e. the intended lengths, while "CAD w/ Bias" are the ones printed onto the photolithography masks, taking an expected fabrication bias into account. "uScope" are the length measurements taken here. Finally, "uScope - w/o Bias" are the difference between intended and measured lengths.

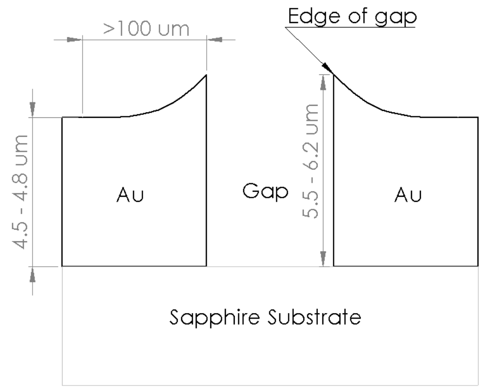

Figure 4.4: Schematic sketch of the electroplating thickness variation, showing a cut-away view from the side, e.g. the x-axis. Apart from the gaps, the Au was planned to be plated to $4.0\,\mu$m thickness. However, after plating, the general thickness was measured as $4.5\ldots4.8\,\mu$m. Close to gaps, current crowding resulted in thicker plating, $5.5\ldots6.2\,\mu$m, which falls off to the general thickness after a distance of $\sim 100\,\mu$m.

4.2 Chip Bonding

O_2-Plasma cleaning is an effective way of cleaning surfaces. However, due to some concerns about unintended consequences of the cleaning process on both the ion trap chip itself and on the pillbox, it is not used. In particular, the concerns are

- Oxidation may occur on any areas of the copper pillbox which remain not gold-plated, e.g. any holes
- Oxidation may occur on the eutectic solder[5] due to its tin content of 20 %
- High voltages have the potential to cause electric arcs, potentially damaging small structures like the ion trap's microwave lines

4.2.1 Gluing Vs Soldering

Initially, the plan was to glue the ion trap chip onto the pillbox using UHV compatible epoxy, as previously performed by NIST. However, thermal conductivity κ at room temperature of epoxy is $0.1 - 1$ $^{W}/_{\mathrm{m}\cdot\mathrm{K}}$ for most epoxies, and goes up

[5] For details on the solder, see sect. 4.2.2

to ≈30 $^{W}/_{\mathrm{m\cdot K}}$ for special high thermally & electrically conductive epoxies[6] with very specific curing requirements [Tip 30, 56]. Upon cooling to a temperature of ≈ 20 K, the thermal conductivity is lowered by a factor of $a = 5$ for conductive, and $a = 2$ for non-conductive glue [p. 26, 4]. This is partially counter-intuitive, as the thermal conductivity of the common filler silver *increases* by an order of magnitude upon cooling to ≈20 K [12-218, 12]. The power P conducted through a layer is

$$P = \frac{A}{d}\int_{T_C}^{T_H} \kappa \mathrm{dT} \approx \frac{A \cdot \kappa \cdot \Delta T}{d} \tag{4.1}$$

The chip has an area of $A = 5\,\mathrm{mm} \cdot 10\,\mathrm{mm} = 50\ \mathrm{mm}^2$, and made of monocrystalline sapphire, with a thermal conductivity of $\kappa\,(20\,\mathrm{K}) = 3500\,^{W}/_{\mathrm{m\cdot K}}$ [p.12-221, 12]. Its thermal conductivity is thus very high in comparison to even thermally conductive glue; therefore even a localised source of heat heats up the entire chip, before being dissipated through the glue; the entire area of the chip can be used in this calculation.

The thinnest recommended bond line thickness for thermally conductive epoxy is d = 25 μm, and d = 12 μm for silver filled epoxy [Tip 6, 56]. However, thermally conductive epoxies are also mechanically weakened by the addition of the filler, and are less likely to survive thermal cycling. In fact, one *recommended* method for removing epoxy is thermal shock [Tip 8, 56]. Thicker bond lines are associated with a higher chances of surviving thermal shock. Gluing a 10 mm x 5 mm chip onto gold plated brass, using the popular thermally conductive epoxy H20E, from Epotek, which has a room temperature thermal conductivity κ of 29 $^{W}/_{\mathrm{m\cdot K}}$, resulted in the chip breaking by cryogenic thermal shock [p. 450, 13] - this is despite bond line thicknesses of (35 ± 15) μm and (100 ± 50) μm. This is cause for concern as that situation is very close to those intended to be used in this experiment.

[6]These are usually filled with flakes of silver

Assuming a highest acceptable temperature difference of $\Delta T = 20$ K, and using the thermally conductive epoxy H20E, from Epotek, with a bond line thickness of 100 μm, a cooling factor of $a = 5$, the resulting permissible power is $P = \frac{50\,\mathrm{mm}^2 \cdot \frac{1}{5} \cdot 29\,{}^{W}\!/_{\mathrm{m \cdot K}} \cdot 20\,\mathrm{K}}{100\,\mu\mathrm{m}} = 60$ W, which would be more than sufficient. As mentioned above however, this is likely to result in breaking by thermal shock. Using a more mechanically stable, unfilled epoxy like H70E, from Epotek, with a thermal conductivity κ of $0.9\,{}^{W}\!/_{\mathrm{m \cdot K}}$, the resulting permissible power is $P = \frac{50\,\mathrm{mm}^2 \cdot \frac{1}{2} \cdot 0.9\,{}^{W}\!/_{\mathrm{m \cdot K}} \cdot 20\,\mathrm{K}}{100\,\mu\mathrm{m}} = 5$ W, which is still more than the ≈ 2 W, but only by a factor of ≈ 2.5. If there were any other effects that lower the effective thermal conductivity, e.g. ones that introduced thermal resistances at glue to metal interfaces, the highest acceptable temperature difference were lower, or the power dissipated in the trap were higher, then the epoxy would limit the permissible power below that intended to be dissipated. This would lower the performance characteristics of the trap. Combined with the uncertain mechanical strength, epoxy presents a non-ideal choice of adhesive.

For the above reasons a different method was chosen, using soldering instead of gluing.

4.2.2 Eutectic Soldering

Soldering directly onto the sapphire presents some difficulty; to achieve good adhesion, both sides to be soldered would ideally be metallised. This way, the thermal conductivity is a lot higher, as metals generally have thermal conductivities several orders of magnitude better than epoxy, especially at temperatures of ≈ 20 K. A layer of titanium on the sapphire assures good adhesion on it, while a final layer of gold assures both a chemically passive surface, and a stable layer to solder onto. However, at elevated temperatures above 250°C, the titanium and gold can diffuse into each other [38]. To prevent this from happening, a barrier layer of platinum can be used [61, 48, 16]. The thicknesses of the Ti and Pt, 40 nm and 60 nm, are based on [35], although with a 100 nm Au layer to bring the layer stack in line with previous electroplating tests. The final layer stack, before electroplating, is thus 40 nm Ti, 60 nm Pt, and 100 nm Au on both the bottom and the top[7].

The Cu pillbox is plated with ≈5 μm Au, therefore both surfaces to be soldered onto are gold. A fluxless UHV compatible solder is chosen, consisting of Au and Sn. This is a form of binary eutectic solder, having a local minimum in its melting point at 80 % Au, and 20% Sn (by weight), at 278 °C, see fig. 4.5. The eutectic mixture comes as a (25 ± 5) μm thick foil, part 304-078-24 from GoodFellow Cambridge Ltd. Once heated above the eutectic temperature, the mixture melts, and starts to consume the Au on both other surfaces. This increases the melting temperature of the mixture, either until it reaches the ambient temperature and freezes, or all the gold is consumed. An overview of the process can be seen in fig. 4.6.

On the ion chip side, the stack is sapphire/Ti/Pt/Au, with respective thicknesses of 430 μm, 40 nm, 60 nm, and 100 nm. On the pillbox side, the stack is simply 5 μm Au, then OFHC Cu. Out of these, relevant thicknesses for eutectic soldering are the 100 nm Au of the Sapphire, the 5 μm Au of the Pillbox, and the 25 μm thick Au/Sn foil. Assuming the entire Au platings of the pillbox and ion trap chip are consumed, the resulting mixture will have a higher melting temperature, which can be calculated using the densities of AuSn, $\rho_{AuSn} = 14.51$ g/cm^3 and of Au, $\rho_{Au} = 19.30$ g/cm^3. $R_1 = 0.80$ is taken to be the ratio of the mass of gold, m_{Au}, to the mass of tin, m_{Sn}, before melting. After melting, a mass of gold, m_{Au+}, is added, leading to a new ratio of gold to tin, R_2:

$$R_2 = \frac{R_1 \cdot \rho_{AuSn} d_{AuSn} + \rho_{Au} d_{Au}}{\rho_{AuSn} d_{AuSn} + \rho_{Au} d_{Au}} = (84.3 \pm 0.9)\ \% \tag{4.2}$$

Interpolating the melting temperatures on fig. 4.5, the melting temperatures

[7]Note that the choice to metallise the sapphire's surface does not rule out using epoxy later on, so it increases the available options

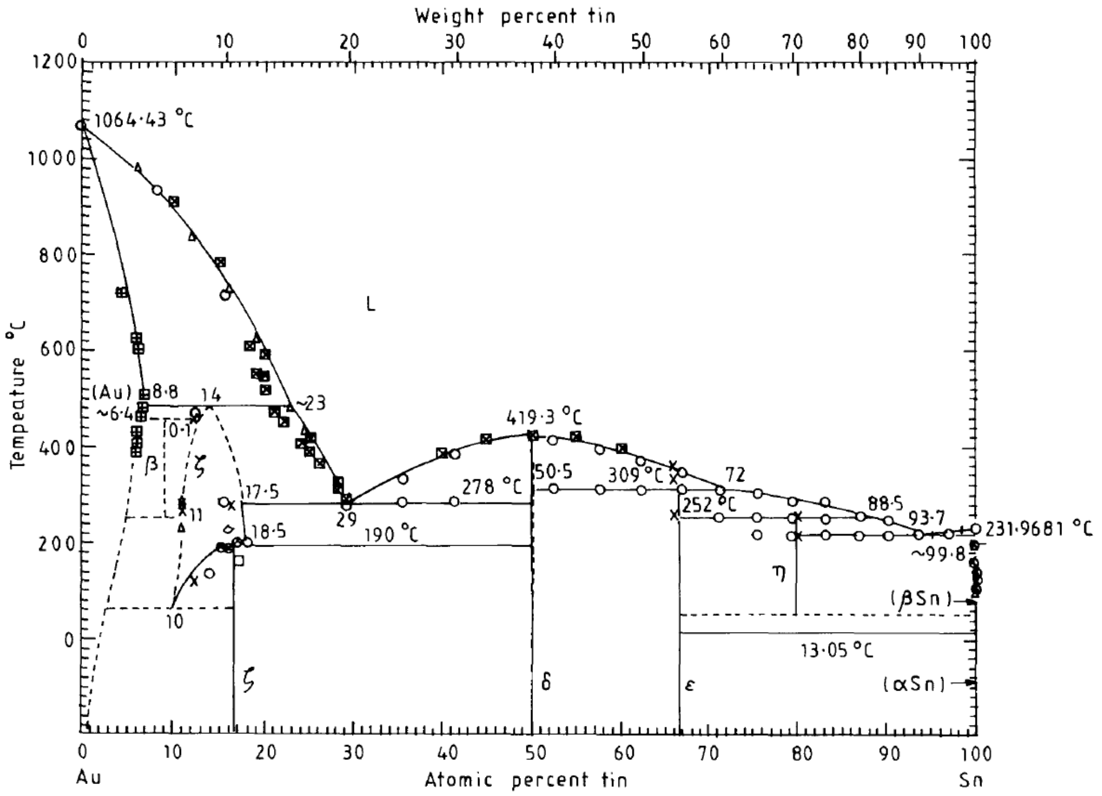

Figure 4.5: Gold-Tin eutectic phase diagram. The local minimum in the melting point is at 278 °C for 80 % Au, and 20% Sn (by weight). For increasing gold content, the melting point increases by ≈ 50 °C/%. [39]

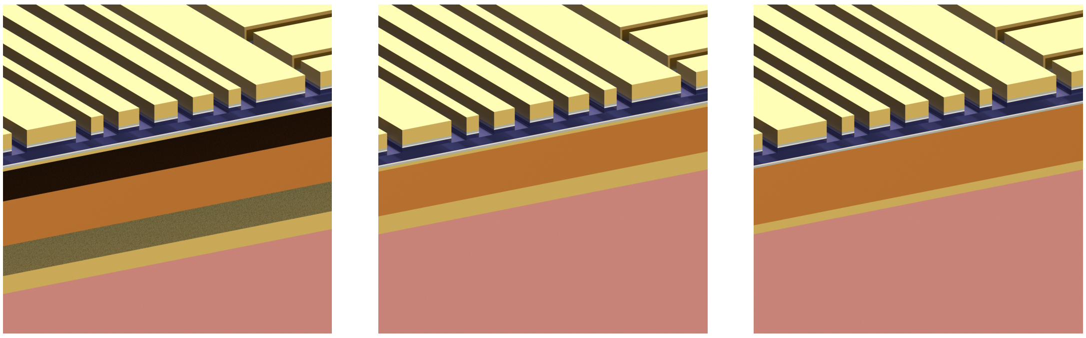

(a) (b) (c)

Figure 4.6: Process diagram for eutectic soldering of the ion trap chip onto the pillbox. (a) A 25 μm Au-Sn foil is placed on the pillbox, and the chip on top of that. (b) A calibrated force is applied to the corners of the chip, closing the gaps. The surrounding air is pumped out, creating a high vacuum. (c) The pillbox is heated to 245°C for 15 min to thermalise everything, then to 305°C for 15 min to melt the solder; the solder freezes again after consuming the entire 100 nm gold layer on the chip and $\approx 2\,\mu$m of the available 5 μm of the gold plating of the pillbox. The pillbox is allowed to cool; after reaching room temperature the system is vented to atmosphere, and the force is taken away. The chip/pillbox combination is is then taken out of the eutectic bonding setup, ready for further processing.

have a roughly linear dependence on Au weight content between 10 % and 20 %, with melting temperatures of 800°C and 278°C, respectively. This implies a $\approx$ 50 °C/% slope, leading to an approximate melting temperature of $(500 \pm 50)\,$°C at the gold to tin ratio of $R_2 = (84.3 \pm 0.9)$ %. This is much higher than the process temperatures of $<$ 350°C used in this thesis, therefore the continuous process of gold being absorbed into the AuSn will cease when the melting temperature reaches the process temperature. Again using the interpolation shown in fig. 4.5, the resulting AuSn ratio is likely $\approx$ 81 %, as opposed to the potentially available $\approx$ 84 %.

As explained above, the temperature needed to solder is $\sim$ 300°C. At these temperatures, there is some concern about room air depositing residue onto the trap's gold surface, potentially leading to a higher heating rate. Furthermore, any holes in the pillbox would be oxidised, which could degrade electrical conductivity of some contacts. To address these concerns, the soldering process is done under vacuum.

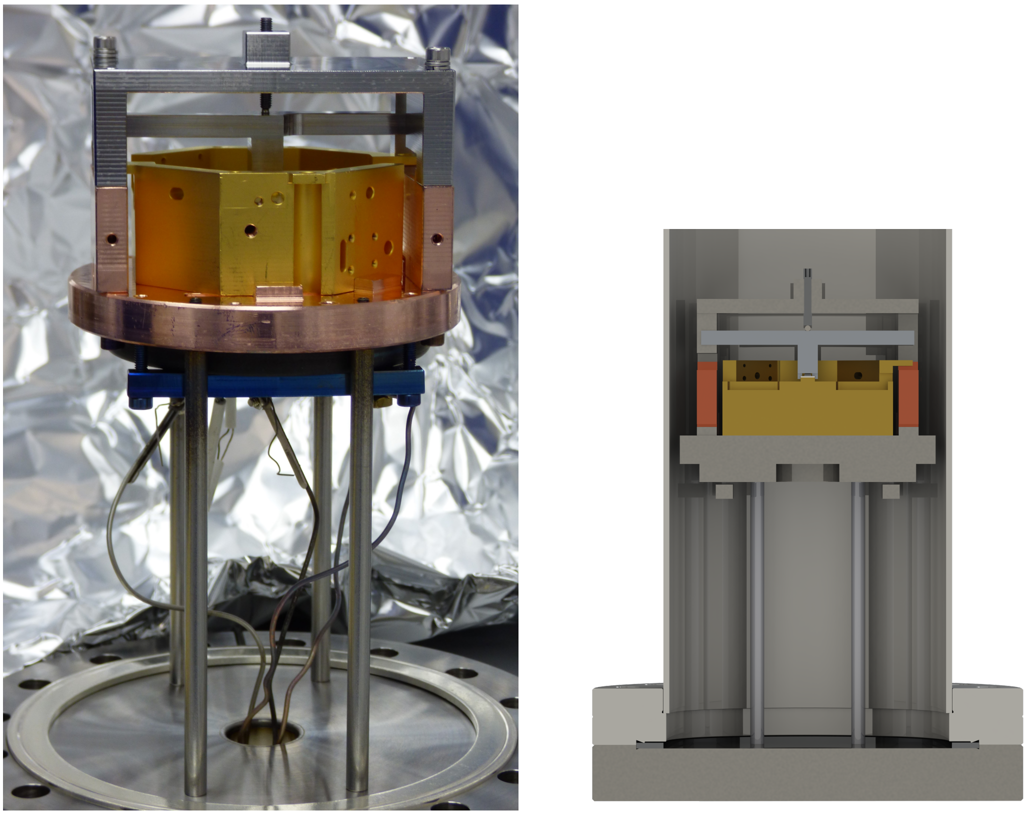

(a) (b)

Figure 4.7: Vacuum oven used for eutectic soldering. (a) Taken with a Panasonic DMC-FZ72 camera vs (b) model cross section rendered with Solidworks 2016 Professional. The oven is enclosed in a CF100 nipple (only shown in (b)), does not touch the walls, and is held up on four solid 10 cm long, Ø5 mm SS 304 legs. This reduces the thermal load via its mounting to a few Watts at 300°C, which is negligible compared to the power of the heater, 300 W. The main heat loss is radiative, experimentally determined to be 10's of Watts based on the heater's duty cycle. The temperature is measured next to the pillbox via a type J thermocouple. Pressure on the ion trap chip is applied via a screw, whose pressure is evenly distributed on the four corners of the chip. The thermal expansion profile is designed such that the pressure slightly increases upon heating to the bonding temperature. During the brief heating to 600°C, the heater's cross bars and screws turned blue.

The eutectic bonding setup is shown in fig. 4.7. The pillbox is attached to a base plate made of OFHC Cu, onto which a 300 W heater is attached, part 3682K22 from McMaster-Carr Supply Company, with a maximum current of 2.5 A, voltage of 120 V, and temperature of 650°C. Using MoS_2-coated M3 screws and cross bars, both made of SS 304, the heater is pressed against the base plate to give a reasonably good thermal contact. The temperature is read by a type J thermocouple, part 409-069 from TC Direct, attached on the top of the base plate. Again using MoS_2-coated M3 screws, four solid 10 cm long, Ø5 mm SS legs, are mounted to the base plate on a CF100 flange with a 25 A high current and thermocouple combination feedthrough, part ZFHTCJ-1-25C-100CF-M3X4 from LewVac Components Ltd. The temperature is read out and adjusted by an on-off temperature controller, part SDC240JC-SAC from Briskheat. As the temperature controller outputs the full 240 V mains voltage, and the heater is only specified for half that amount, a variable transformer is used, part 8902819 from RS Components, limited to 10 A. The inclusion of a variable transformer also has the beneficial side effect of being able to adjust the gain of the temperature controller.

All of the above is enclosed in a CF100 nipple, 270 mm long, part FL-SC100CF from LewVac Components Ltd. Vacuum is generated and maintained using a diaphragm backed turbomolecular pump, parts A74602983 and EXT 70H, regulated by a controller, part EXC 120, next to which a gauge is used to measure the pressure, part D14701000, all from Edwards. Using a ≈1 m section of existing flexible KF25 hosing, the pump is attached to the nipple, through a CF100 to CF40 adapter, part FL-SR100CF-40CF-A from LewVac Components Ltd, and an existing KF25 to CF40 adapter. In an effort to reduce the amount of outgassing that the heater produces, it is incrementally run at higher temperatures. Firstly, the heater on its own is run in air to 200°C for three hours. Afterwards, it is assembled onto the bonding setup, evacuated, and incrementally heated to 400°C, and finally briefly to 600°C, at which point a pressure of 10^{-1} mbar is measured; this should severely limit the outgassing happening at 300°C, which could deposit residue on the trap's surface, potentially worsening heating rates. Typical pressures achieved with the heater receiving 110 V are $2 \cdot 10^{-4}$ mbar at 300°C, and $2 \cdot 10^{-5}$ mbar at 300°C and $5 \cdot 10^{-7}$ mbar at room temperature, the latter two without voltage being applied to the heater.

Pressure on the chip is required while soldering to achieve good thermal conductivity, and proximity of the surfaces to be soldered, while too much pressure breaks the sapphire. Recommended bonding pressure are $0.2 - 0.5$ MPa, corresponding to $10 - 25$ N on the 5 mm x 10 mm chip size [p. 217, 33]. For the subsequent bonds, a force of 25 N is chosen. In order to calibrate the overall force applied to the ion trap, a force sensor is used, part FC2231-0000-0025-L from TE Connectivity. The torque required to turn the M3 grub screw to achieve the target pressure is noted, and subsequently reproduced on the ion trap. However, as this torque is applied by hand, the accuracy is only 25 ± 7 N. On one occasion, the chip cracked, likely caused by overpressure, and/or an uneven surface[8]. Both of these factors were taken into account and no further cracked chips were observed.

The trap itself needs to be well aligned to the pillbox for bonding. To achieve this, a holder is designed that registers itself to the corners of the little "pedestal" that the ion trap rests on; the ion trap is registered via its own corners to the holder, this is also where the force is applied. Both the pedestal and the trap are designed to be 10.0 mm by 5.0 mm; however the trap's size is measured as 10.1 mm by 5.1 mm after the dicing of the wafers. The holder is adjusted to match. The positional and angular alignment tolerances are thus higher than anticipated, although still acceptable.

To minimize the amount of angular movement, the screw's applied torque on the ion trap is limited by three measures, see fig. 4.8. Firstly, the screw, part MT4-320-C from U-C Components Inc., does not press down directly onto the holder, but rather onto a $\varnothing$1 mm SS 316 ball bearing from Simply Bearings Ltd; this minimizes the leverage[9], encouraging slippage, reducing the torque applied to the holder. Secondly, the holder has arms which are registered to the side

[8] In this case, the item bonded to is not a pillbox, but just a small piece of copper, plated to the same gold thickness. It is not machined to as high precision as the pillbox.

[9] In an ideal model, a ball bearing pressing down onto a flat surface gives a point contact, and thus no leverage arm. However, a small indentation in the holder is needed to keep the bearing in place before tightening, resulting in a small leverage arm.

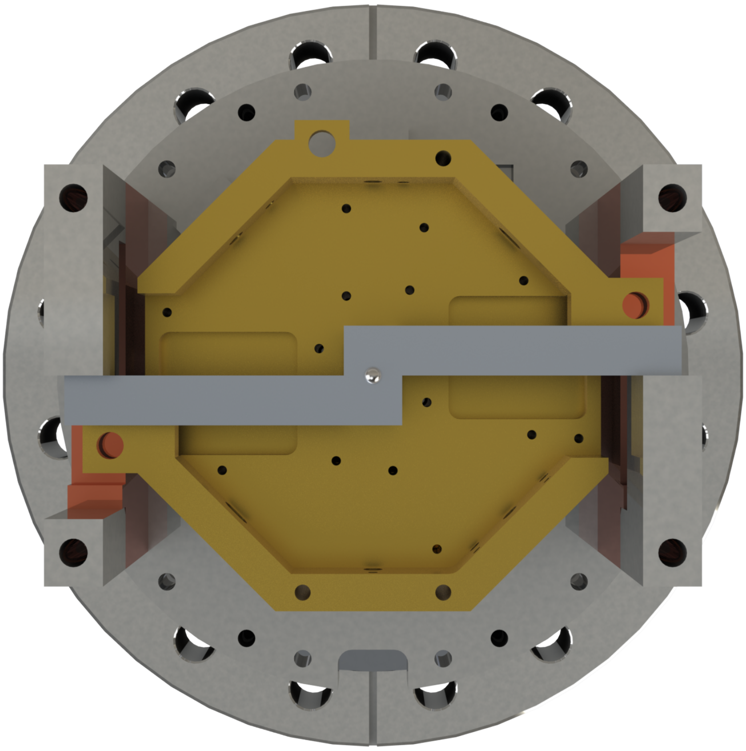

Figure 4.8: Top down view of the vacuum oven system. The grub screw used to apply torque is screwed in a clockwise direction. The effect of the torque of the screw on the ion trap is minimized by using a ball bearing (centre of image). Any remaining torque is on the holder is braced against the side walls. The effect of the torque on the ion trap should be minimal.

walls of the bonder, therefore any torque applied should mostly press against these. Lastly, the holder itself is registered to the corners of the pillbox' pedestal, although as just mentioned this capability is partially negated by the increased trap size; however the first two measures minimize the angular movement of the ion trap sufficiently.

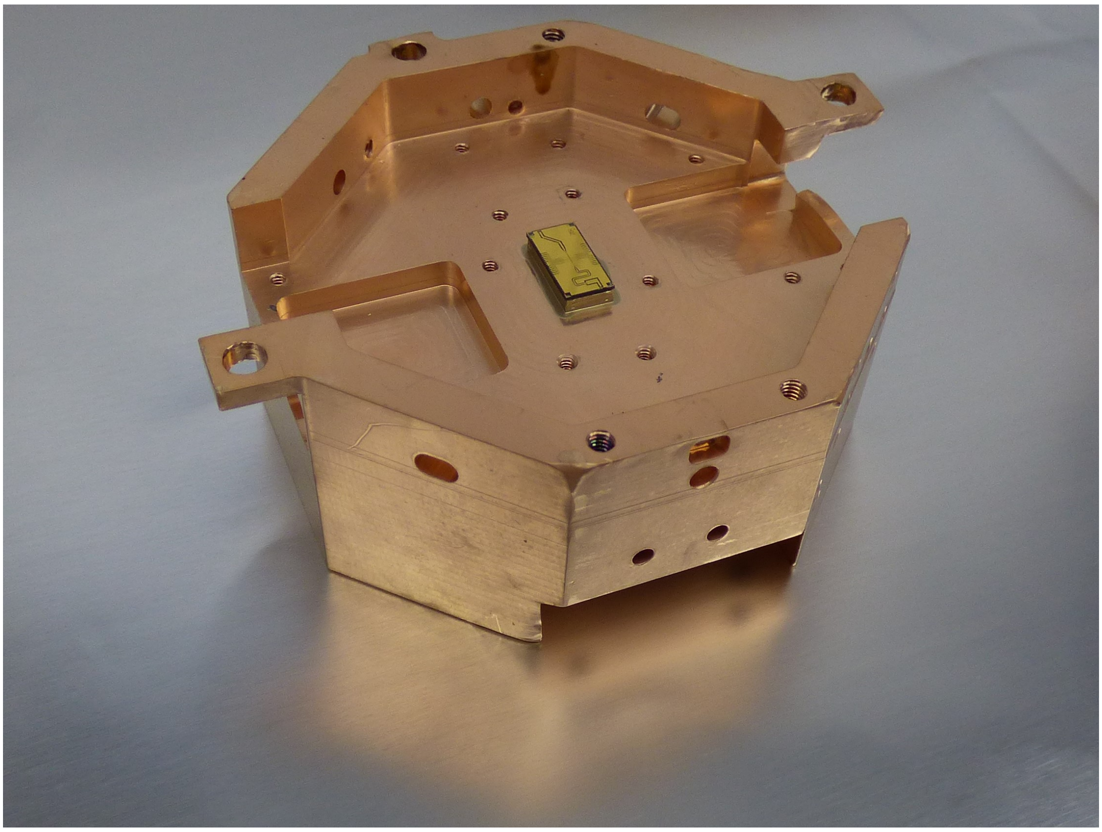

Figure 4.9: Successfully bonded trap. The area immediately surrounding the trap is covered in eutectic solder where the solder foil was previously sticking out. For this test, the hold time at >300°C is 1 hour, which causes the gold plating to migrate into the bulk copper. Loss of gold plating is not observed for subsequent hold times of 15 min, even if followed by a 190°C bake for several days.

In order to achieve a constant force on the chip during the heat up to the process temperature of ≈ 300°C, the thermal expansion of the parts is modelled. Fig. 4.7b shows a cross section of the eutectic bonding setup. To keep the force on the chip constant, the thermal expansion of the central path L_C needs to be the same as that of the outer paths L_O. Therefore,

$$L_C = L_1 + L_2 + L_3 = L_4 + L_5 + L_6 = L_O \tag{4.3}$$

The size, and thus the thermal expansions, of some of these components is already fixed, namely that of the pillbox L_6, ion trapping chip L_5, ball bearing L_3, and aluminium holder L_4. However, the remaining two degrees of freedom can then be used to balance equ. 4.3.

In order to test the performance upon thermal shock, a trap was bonded onto a gold plated block of OFHC Cu. For this test, the volume of OFHC Cu is ≈ 0.4 % of the pillbox, thus giving an equally reduced thermal mass and steeper rates of thermal changes (in otherwise equal circumstances). This is then submerged in liquid nitrogen for ≈ 20 s, until the liquid is no longer evaporating on the block's surface, i.e. reaching thermal equilibrium of ≈ 77 K. It is then transferred out of the liquid, and held under a hot air gun for ≈ 2 min, until ≈10 s after the frost that formed on the cold surfaces upon entering air has evaporated, i.e. ≈ 100°C. This process is then repeated 20 times. The steepest rate of temperature[10] change is $_\Delta\dot{T}_{max} = \frac{373\,\mathrm{K}-77\,\mathrm{K}}{20\,\mathrm{s}} = 15\,\frac{K}{s}$; this rate is more than an order of magnitude higher than the trap will undergo in the experiment[11], therefore if the joint holds here, then it should withstand the comparatively lighter stresses of the experiment, too. After this test, the trap was still found to be attached, resistant to being removed by bare hands, and even pliers were found to be only chipping off the sapphire, not breaking the eutectic bond.

Removing a bonded trap, e.g. due to subsequent damage to it, from a pillbox is thus not trivial. The method used here is a "hammer and chisel", demonstrated in fig. 4.10. The colours of the surface of both the pillbox and the underside of the trap pieces are found to be the same kind of grey as the Au-Sn foil, i.e. the adhesion of the solder to either side was better than the internal adhesion

[10] Assuming a linear temperature slope, which is only a first approximation

[11] Without active throttling of the helium exhaust, the steepest rate of temperature change achieved is $\approx 40\,{}^{K}/_{\mathrm{min}} < 1\,{}^{K}/_{\mathrm{s}} \ll 15\,{}^{K}/_{\mathrm{s}}$

of the solder. This mode of failure is preferable to either side breaking, as it indicates the strongest bond possible for the given adhesive [Tip 13, 56]. Notably, a trap could successfully be bonded onto the pedestal shown in fig. 4.10d, despite the lack of visible pristine gold plating on the pedestal. This could indicate that a pristine gold layer is not critical to the bonding process. Alternatively, it could indicate that there is enough remaining pristine gold underneath the visible eutectic mixture to halt the melting process in this subsequent bond. In any case, attaching an ion trap to its mounting using eutectic soldering is found to form a mechanically strong and reliable bond.

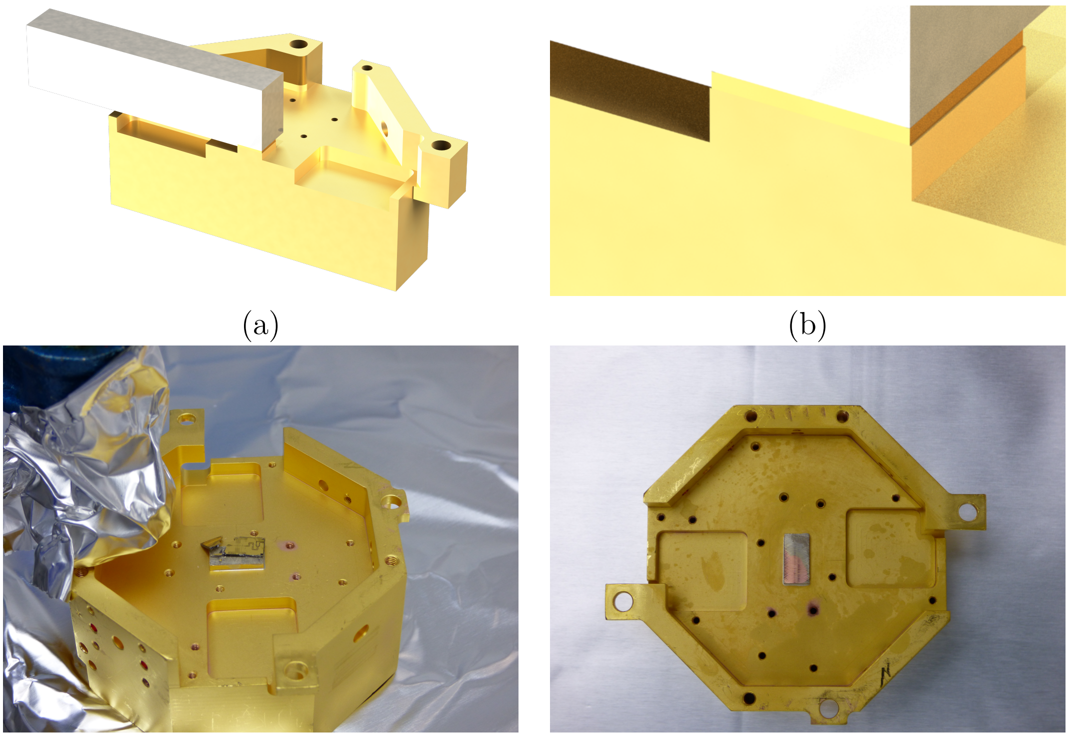

(a) (b) (c) (d)

Figure 4.10: De-bonding procedure of a eutectically bonded trap. (a) A cross section of the custom chisel, made of SS, ready to debond the trap from the pillbox. Force is intended to be applied from the left end. (b) The chisel tip has a slight step in it, touching the trap on the top and long side, while avoiding touching the pillbox. (c) The pillbox is clamped down onto a table (for cleanliness the clamp is covered in aluminium foil, top left). Multiple light hammer taps, using a 0.3 kg hammer, part 200C.26 from Facom, onto the chisel break off pieces of sapphire and debond them; shown here are the remaining pieces of the trap. (d) Using a workshop mill bit, the surface of the pedestal is smoothed, ready for cleaning and the next trap to be bonded. (a) and (b) rendered with Solidworks 2016 Professional, (c) and (d) taken with a Panasonic DMC-FZ72 camera.

4.3 Electrical Connections

4.3.1 Wire Bonding

In order to connect the ion trap chip to its surrounding PCBs, the method of wire bonding is chosen. The initial "wire" choice in this case is a 125 μm x 50 μm, 99.99 % purity, gold ribbon from Inseto (UK) Ltd. Therefore, all surfaces to be bonded are gold - the PCBs, wire/ribbon, and ion trap chip. They are bonded using an F S Bondtec, part 53xxBPA. The ribbon is guided using wedges, part 4WRF4-D440-T7G-M00. The bonding procedure is performed inside a class 1000 clean room.

The best set of parameters to bond the 125 μm x 50 μm ribbon found in this thesis only has a successful bond rate of ≈50 %, despite an extensive parameter space search. Furthermore, the failed bonds sometimes damage the pads to be bonded on, so cannot be retried indefinitely. Combined with the low successful bond rate, this brings the chance of successfully bonding all 22 DC, 3 μW, and 3 RF line[12] bonding pads to near 0, even with retrying failed bonds.

As the above ribbon is the thickest and widest available from this supplier[13], the process was retried with a 75 μm x 25 μm ribbon, 99.99 % gold, part AuR-75x25A50 from Tanaka. This thinner ribbon worked a lot better; a set of parameters with a successful bond rate of ≈ 95 % was found when bonding the μW PCB and RF PCB. However, bonding onto the slightly narrower pads of the DC PCBs results in de-laminating ≈ 50 % of the gold traces from the underlying Alumina.

For the DC PCBs, the thickness of the wire is not as critical as for the microwave and RF current carrying wires. Therefore the wire cross section can be reduced further. A 25 μm diameter wire, 99.99 % gold, part GSA-25A100 from Tanaka, is used, and the process set up with a successful bond rate of 100 %, with

[12] One signal, and two ground bonds

[13] Thus the theory is that the ribbon is only borderline bond-able, because it is too thick

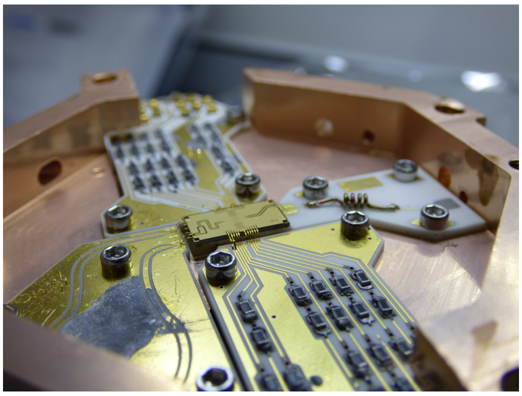

(a)

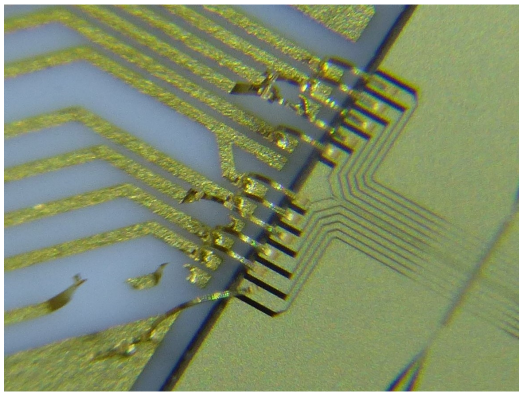

(b)

Figure 4.11: Failed wire bonds. (a) shows an overview, while (b) is a close up of the DC pads. Some of the DC pads de-laminated after bonding with the 75 μm x 25 μm ribbon. While most of the bonds are electrically still connected, they are likely too mechanically weak to be reliable. Images taken with a Panasonic DMC-FZ72 camera, (b) through an additional magnifying glass.

no cases of de-lamination of tracks on the DC PCBs.

All the above successful bond rates are only achieved after doing a few practice bonds, potentially fine tuning the bond parameters used. It is thus useful to allow for an area on the PCBs that this practice bonding can be performed on without creating problems with beam access, or electrical pathways. In this experiment, the spare microwave line on the μW PCB is used for this purpose.

To conclude, the overall process is to use 75 μm x 25 μm ribbon for the μW and RF PCBs, and 25 μm diameter wire for the DC PCBs.

4.3.2 μD Connector

The majority of electrical signals are transmitted using a 51 pin μD-connector. It is used to transmit five types of signals: the 22 DC lines connecting to the trap, 2 DC grounds, a 4 wire temperature diode, the current supply of the two ovens (each line distributed among 3 pins to reduce the current load on each pin, i.e. $4 \cdot 3 = 12$ pins in total), and two thermocouple pairs to measure the ovens' temperatures, see fig. 4.13. In total, these use up 44 pins, leaving 7 spare.

On the vacuum side, as discussed in sect. 3.6.3, the DC and temperature

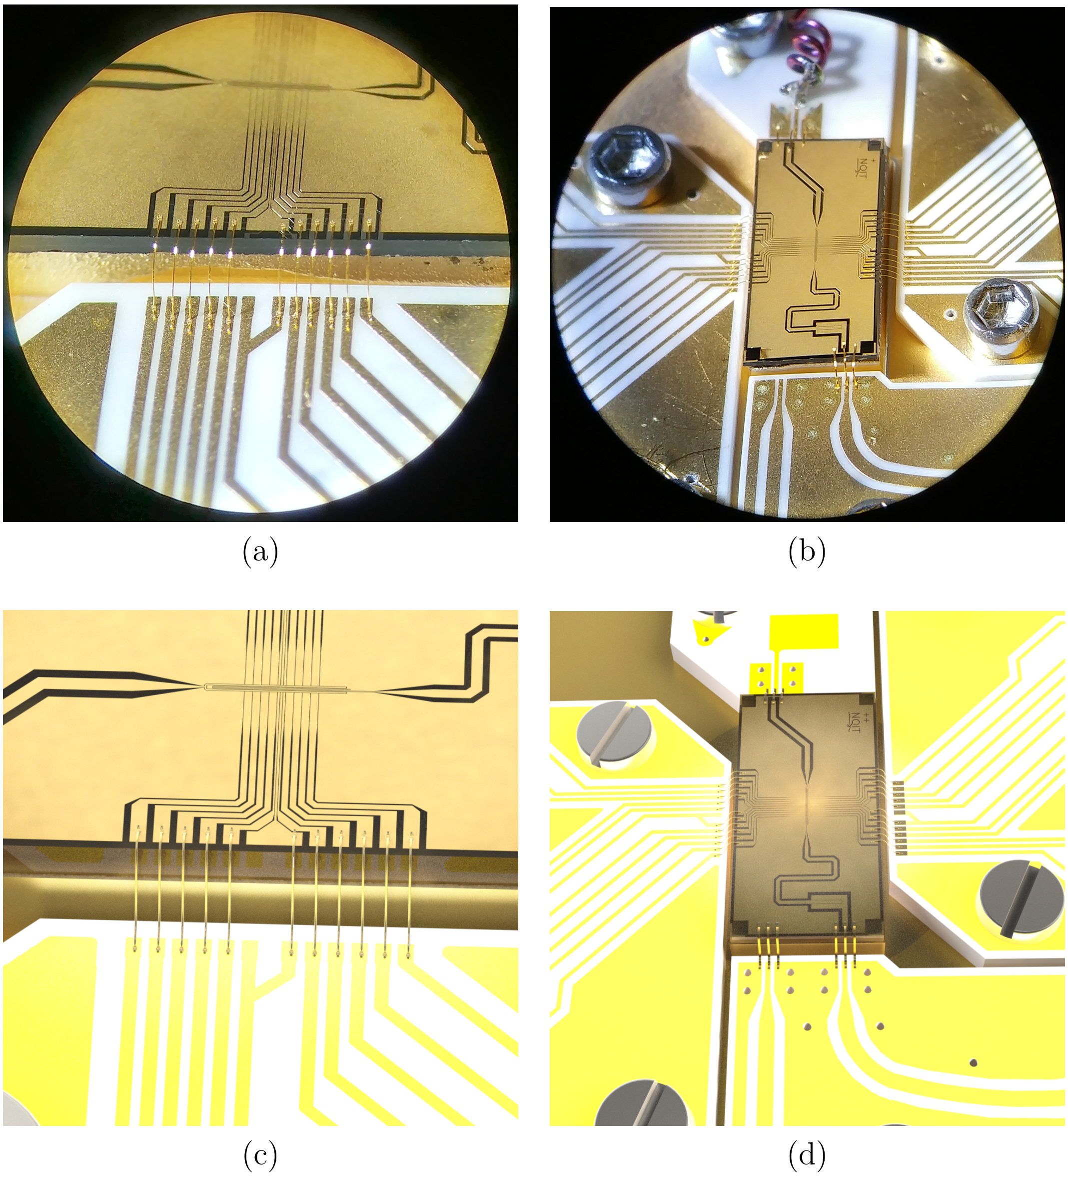

Figure 4.12: Successful gold wire (a) and ribbon (b) bonds. For a sense of scale, the screws (b) are M1.6, the ion trap (b) is 5 mm by 10 mm, and the pitch of the DC bonds (a) is 250 μm. They are taken with an LG G5 camera through the bonder's microscope. These can be compared to the CAD model (c) and (d) of the same views, rendered with Solidworks Professional 2016. The only unplanned change between model (d) and implementation (b) is an inductor; the added coil at the top PCB on (b).

wires going from the μD connector, which is a custom PEEK connector from Allectra Ltd, to the pillbox are made of $\varnothing$0.40 mm, 10 cm long BeCu from MWS Wire Industries Inc. to bring down the thermal conductivity. They are insulated using Kapton tubes that are slid over them, part 312-KAP-TUBE-05-300 from Allectra Ltd. The ovens' wires are made of Kapton coated $\varnothing$0.4 mm copper, part 311-KAP-040-10M from Allectra Ltd, again three for each polarity. The ovens' thermocouples are Kapton coated, type K, $\varnothing$0.16 mm wire, part 1512070 from MDC Vacuum Ltd. They are attached to the ovens using an existing spot welding process, consisting of two spot welds per thermocouple [p. 2f, 6], for which a stored energy power supply is used, part 125ADP from Amada Miyachi Inc. Two spot welds are needed to avoid the oven's current from influencing the thermocouple voltage. The DC wires are crimped onto sub-D HD female crimp pins, part 212-PINFHD-26, using the crimp tool 214-CTOOL-HQ, both from Allectra Ltd. On the DC PCBs, their male counterparts, part 212-PINMHD-26, are soldered. As they are meant for wire to wire connections, the male crimp pins feature a crimp section that goes unused; these are removed to give better optical access into the pillbox.

The air side μD cable shown in fig. 4.13, uses a male 51 pin μD connector, which is spliced onto three cables: Firstly, a shielded 25 pin D-sub connector (for the DC lines), part 1656291 from Phoenix Contact, secondly a shielded 15 pin D-sub cable, part 1223041 from RS Components Ltd, soldered onto a male 15 pin D-sub connector (for the ovens), and thirdly onto a Cat 6 Ethernet cable (for the temperature PCB). The latter is a readily available choice, used for its shielding and twisted pairs[14].

[14]There are two commonly used connector standards for 15 pin D-sub, one featuring 2 rows and one featuring 3 rows, the latter being used for VGA. Most readily available D-sub cables are not rated to the > 3 $^{A}/_{\text{pin}}$ of current required by the ovens, and hence a 3 row cable specified to 5 $^{A}/_{\text{pin}}$ is soldered onto the 2 row connector standard used in this thesis.

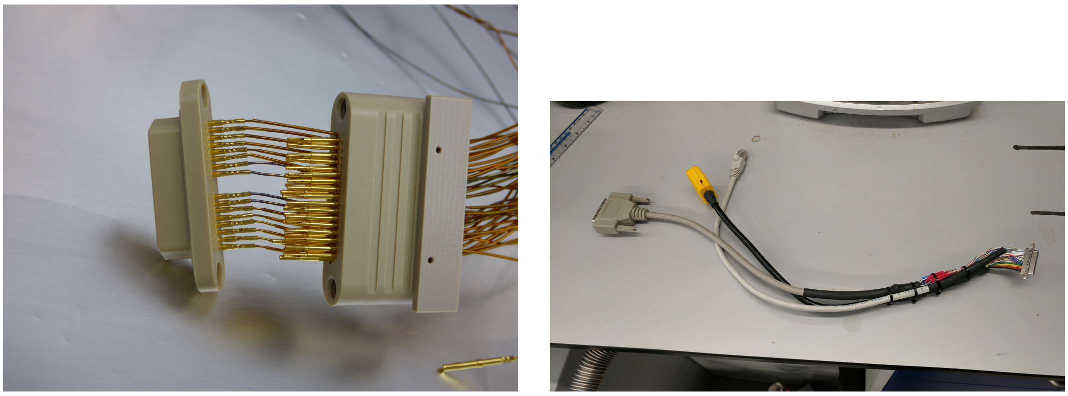

(a) (b)

Figure 4.13: μD cable assemblies. (a) Shown here is the partially assembled in-vacuum μD connector. It is a two part design made of PEEK, between which the pins are confined once the two parts are screwed together. (b) This is the custom air-side cable, distributing the base flange's 51 μD connection (right) onto a 25 pin D-sub (leftmost connector, for the DC signals), a 15 pin D-sub (yellow connector shell, for the ovens), and an Ethernet cable (for the temperature PCB). Taken with (a) a Panasonic DMC-FZ72 and (b) an LG G5 camera.

4.3.3 Microwave cable

In order to transfer the microwaves from the base flange to the pillbox, a semi-rigid SMK cable is used. It needs to have a low thermal but high electrical conductivity, which are partially conflicting goals. The low thermal conductivity is achieved by using SS for both the outer conductor, and the inner conductor. Due to the skin effect, the region in which electrical signals in the inner conductor of a coaxial cable travel is only a few micrometers thick[15]. This means the conductivity of the inner conductor can be substantially improved by using thin layers of good conductors like copper or silver surrounding the SS. In this experiment, the SS inner conductor is surrounded by a few hundred microns of copper, followed by a few microns of silver. Overall, the thickness of the copper and silver are too small to matter from a thermal conductivity point of view, but greatly increase the electrical conductivity at microwave frequencies.

The design used for the pillbox and base flange calls for two bends in the microwave cable; one 90° bend as close as possible to the pillbox end, and one

[15]depending on the frequency of the signals and conductivity of the conductor

180° bend roughly in the middle of the cable. The cables are purchased from Allectra Ltd, part 390-SMK-MM-100-MCOA, and come with a 90° bend as close as the manufacturer could reliably produce[16]; a bend close to the SMK connector is necessary due to geometric constrains inside the vacuum system. The 180° bend is produced in-house with a cable bender, part T570-086 from Seven Associates Inc.

According to the manufacturer, the cable has a dampening of ≈ 0.6 $^{dB}/_{\mathrm{m}}$ at 3.2 GHz, thus about 0.06 dB across the cable, however this does not take into account insertion losses.

4.3.4 RF Drive Chain

As the ion traps used are Paul traps, they need to be supplied with an electric signal which falls in the RF band. Using the trap's geometry, the required voltage at the trap is $80 - 120$ $\mathrm{V}_{\mathrm{Peak}}$, at a frequency of $60 - 80$ MHz. Assuming a simple 50 Ω system, this would imply ≈ 10 W of power going into the system to drive this voltage; this would cause thermal problems and cause instability in the trapping field. In order to use less power, a resonant transformer is used, with the trap side being resonant to the chosen frequency. While this fixes the choice of frequency, it increases the *voltage* on the resonant side by about an order of magnitude, bringing down the *power* required at the input by two orders of magnitude.

A synthesizer, part N5181B MXG from Keysight, is used to create the RF used for trapping. It is connected to an amplifier, part ZHL-03-5WF+, a subsequent 3 dB attenuator, part BW-S3W5+, and a directional coupler, ZFDC-20-3-S+, all from Micro circuits. The latter is used for calibrating the power. From the directional coupler, a 50 cm long SMA cable connects it to the RF resonator PCB, explained below, attached to a UHV feedthrough ending in a bare ⌀1 mm pin on the vacuum side. A barrel connector, part FHP-BECU-1.5IL-CON from LewVac

[16] ≈ 3 mm from the SMK connector

Components Ltd, connects the pin to a ⌀0.80 mm BeCu wire from MWS Wire Industries Inc., insulated using ⌀1.4 mm Kapton tubing, part 312-KAP-TUBE-14-300 from Allectra Ltd. The BeCu wire is then led through a hole in the side of the pillbox, mechanically stabilised using a PEEK spacer, and soldered onto the in-vacuum RF PCB, see sect. 3.8, which is wire bonded onto the trap.

4.4 Vacuum Creation

4.4.1 Ultra High Vacuum Theory

An overview of methods used in this work to minimize the outgassing sources is shown in table 4.2. The precise amount of gas coming from all of these sources however, is not calculated.

Source	**Description**
Diffusion	Gas stored in a material diffuses out of it
Permeation	Gas from ambient pressure permeates through a chamber wall
Vaporization	Material evaporates below its vapour pressure
Desorption	Gas chemically or physically stored on a surface evaporates
Real leak	Gas enters through a small hole in the chamber wall
Virtual leak	A trapped pocket of gas inside the chamber slowly enters the chamber
Back-streaming	Certain types of pumps can be a small source of gas

Table 4.1: Description of outgassing sources

Source	**Ways of minimizing the source used in this work**
Diffusion	Material choice and heating to speed up diffusion to exhaust the source
Permeation	Material choice for the walls
Vaporization	Bulk or surface material choice (e.g. gold plating)
Desorption	Solvents and acid to clean the exposed surfaces
Real leak	Careful handling, inspection for damage
Virtual leak	Design to avoid trapped gas, e.g. by venting holes/slots
Back-streaming	Valve off that type of pump after use

Table 4.2: Outgassing sources and ways of minimizing them

The items with the highest likelihood of significantly contributing to the amount of gas are inside the main vacuum chamber, see fig. 4.14. Using the molecular flow model, the vacuum conductivity for (a) is $C_a = 140\,\mathrm{l/s}$, assuming a simple right angle at the corner as the best approximation of the real design, which is a 5-way connector. The actual conductivity of the 5-way piece will lower the overall conductivity somewhat.

For case (b), the conductance between the pillbox and the main chamber is $C_{aux} \approx 10\,\mathrm{l/s}$, making the conductance $C_b = \left(\frac{1}{C_a} + \frac{1}{C_{aux}}\right)^{-1} = 9\,\mathrm{l/s} \approx C_{aux}$. Therefore, the conductance out of the pillbox limits the achievable pressure in case (b).

Additional pumping concepts beside just one relatively strong pump being relatively far from the likely sources of gas are also considered. One option would be to incorporate another smaller version of a NEXTorr pump directly into the main chamber, instead of connected via pipes. Unfortunately, they would not decrease the pressure by more than a factor of 2 as the conductivity between the main chamber and the main pump is quite good.

A second option is to include getters inside the pillbox. The specific ones that are being considered are St-172 sintered porous getters by SAES getters. The St-172 alloy has a magnetic permeability of $\mu = 1.0001$, which should have a negligible effect on the magnetic field uniformity [p. 7, 1].

Actively heating them to $(450 - 950)\,°\mathrm{C}$, as they are usually operated, would potentially contaminate the pillbox from the inside, while just relying on a passive heating from the UHV bakeout procedure would only activate them to $\approx 50\,\%$. Another problem is that sintered porous getters can release particles, and by using a specially coated "No Particles" version, the pumping speed is further reduced by $35\,\%$. The effect of using them not at room temperature, but rather at e.g. 77 K reduces it by another $20\,\%$. Being deployed inside a confined space, e.g. against a wall on two sides, actually does not decrease the pumping speed by too much, as it is more the bulk of the porous getter as opposed to the surface area, that is responsible for the pumping. Combining these effects means that the biggest size getter, giving $C_{getter} \approx 5\,\mathrm{l/s}$, that fits into the pillbox at most twice, gives $2 \cdot C_{getter} \approx C_b$. Therefore, it would only change the pressure by a minor degree, and because of potential complications due to particles and experimental difficulty, they are not used in the current experiment.

Finally, the pillbox is cooled down to cryogenic temperatures, i.e. there will be some amount of cryogenic pumping. This type of pumping is highly temperature and gas dependent. For the case of LN_2, most of the more complex gas molecules are being pumped, but the more simple ones like H_2 up to CH_4 are not. It is therefore unclear if the pillbox conductance should be maximized or minimized. As a starting point, in the first iteration of the experiment, the pillbox conductance is minimized.

$$\begin{array}{ccccccccc}
\boxed{\text{Main Chamber}} & \Longrightarrow & \boxed{\text{CF40}} & \Longrightarrow & \boxed{90^\circ\ \text{CF63}} & \Longrightarrow & \boxed{\text{Pump}} & \Longrightarrow & \boxed{P=0} \\
\Uparrow & & & & & & & & \\
\boxed{\text{Leak}} & & & & & & & &
\end{array}$$

(a)

$$\begin{array}{ccccccccccc}
\boxed{\text{Pillbox}} & \Longrightarrow & \boxed{\text{Main Chamber}} & \Longrightarrow & \boxed{\text{CF40}} & \Longrightarrow & \boxed{90^\circ\ \text{CF63}} & \Longrightarrow & \boxed{\text{Pump}} & \Longrightarrow & \boxed{P=0} \\
\Uparrow & & & & & & & & & & \\
\boxed{\text{Leak}} & & & & & & & & & &
\end{array}$$

(b)

Figure 4.14: Conductance model used to calculate pressure ratios. Model (a) assumes that the main source of gas is coming from inside the main chamber, but outside the pillbox. Model (b) assumes it is coming from inside the pillbox.

4.4.2 Vacuum Conductivities

We assume that a gas molecule will hit the walls of a chamber many times before hitting another gas molecule. This holds under the condition

$$P\,d < 6.6\,\mathrm{Pa}\cdot\mathrm{mm} \tag{4.4}$$

Using a pressure $P = 10^{-8}\,\mathrm{Pa}$, and a CF 40 inner pipe diameter of $d = 40\,\mathrm{mm}$, this becomes $P \cdot d = 4 \cdot 10^{-7}\,\mathrm{Pa}\cdot\mathrm{mm} \ll 6.6\,\mathrm{Pa}\cdot\mathrm{mm}$. Even for the smaller holes in the pillbox, equ. 4.4 holds by many orders of magnitude. Therefore the space in the entire vacuum chamber is well within the molecular flow region.

The amount of gas Q flowing from one separated region to the next is given by

$$Q = C \cdot (P_1 - P_2) \tag{4.5}$$

where C is the conductivity, and $P_1 - P_2$ the pressure difference. In this model, a molecule colliding with a wall will be emitted along a direction α from the surface normal with a probability of $\cos\alpha$. Therefore, the molecule will most likely be ejected at a normal direction to the surface. Notably, this is independent of the incoming direction. The conductivity C can be calculated analytically for simple cases. For other geometries a Monte Carlo technique is used. This involves calculating the path of many individual gas molecules to get a good approximation of the actual conductivity. Table 4.3 gives a list of conductivities of two geometries.

Geometry	a
Orifice	$\frac{v}{4}A$
Long round tube	$\frac{\pi}{12}v\frac{d^3}{l}$

Table 4.3: Vacuum conductivities of various geometries. Taken from [p. 32-49, 21].

The conductivity of other geometries has been calculated with transmission probability values a from [p. 37-43, 21]. Here, a is the probability that a randomly selected gas molecule which enters one aperture will come out at the other end of the geometry. Notably, this is independent of how long the gas molecule actually takes to get from one end to the other, as long as it eventually reaches it. The conductivity C can be calculated using the transmission probability a via

$$C = a\frac{v}{4}A_1 \tag{4.6}$$

where A_1 is the inlet area and $v = \sqrt{\frac{8k_BT}{\pi m}}$ is the average velocity of the gas molecules [p. 10, 21], using the Boltzmann constant k_B, the temperature T, and the molecular mass m. Therefore, the conductivity is both temperature, and gas molecule dependent. Combining transmission probabilities can be done with the following equation:

$$\frac{1}{a_{1\to 3}} = \frac{1}{a_1} + \frac{A_1}{A_2} \cdot \frac{1}{a_2} + \frac{A_1}{A_3} \cdot \frac{1}{a_3} - 2\frac{A_1}{A_2} \tag{4.7}$$

where the transmission probabilities and inlet apertures of elements 1 to 3 are a_1, a_2, a_3 and A_1, A_2, A_3, respectively, and $a_{1\to 3}$ the total transmission probability. In case all apertures A_1 to A_3 are equal, equ. 4.7 simplifies to

$$\frac{1}{a_{1\to 3}} = \frac{1}{a_1} + \frac{1}{a_2} + \frac{1}{a_3} - 2 \tag{4.8}$$

For just 2 elements this then becomes $\frac{1}{a_{1\to 2}} = \frac{1}{a_1} + \frac{1}{a_2} - 1$. More complicated cases can be calculated with equ. 4.10. Similarly, series conductances are calculated via [p. 39, 21]:

$$\frac{1}{C_T} = \frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} + \dots \tag{4.9}$$

The total transmission probability $a_{1\to n}$ for molecular flow with n elements, inlet areas A_i, and transmission probabilities a_i is given by

$$\frac{1}{A_1}\left(\frac{1 - a_{1\to n}}{a_{1\to n}}\right) = \sum_{1}^{n} \frac{1}{A_i}\left(\frac{1 - a_i}{a_i}\right) + \sum_{1}^{n-1}\left(\frac{1}{A_{i+1}} - \frac{1}{A_i}\right)\delta_{i,i+1} \tag{4.10}$$

with $\delta_{i,i+1} = 1$ for $A_{i+1} < A_i$, and $\delta_{i,i+1} = 0$ for $A_{i+1} \geq A_i$ [p. 47, 21].

This model is used to calculate the vacuum conductivities from the pillbox to the vacuum pumps. The models used are shown in fig. 4.14. The formulae are applied to the vacuum system in sect. 4.4.4.

	Gas →	H_2	C	O	OH	H_2O	N_2	CO	CO_2
	Mass [amu] →	**2**	**12**	**16**	**17**	**18**	**28**	**28**	**44**
T significance	**T [K]** ↓								
Soft bake	500	2300	940	810	790	770	620	620	490
Room	300	1800	730	630	610	590	480	480	380
LN_2	77	900	370	320	310	300	240	240	190
Target T	20	460	190	160	160	150	120	120	100
LHe	4	210	84	73	71	69	55	55	44

Table 4.4: Mean velocities of gasses in m/s important for vacuum systems, at various temperatures, assuming a Maxwell-Boltzmann distribution. Freezing points have been ignored in this calculation, partially as these are pressure dependent, see fig. 4.15.

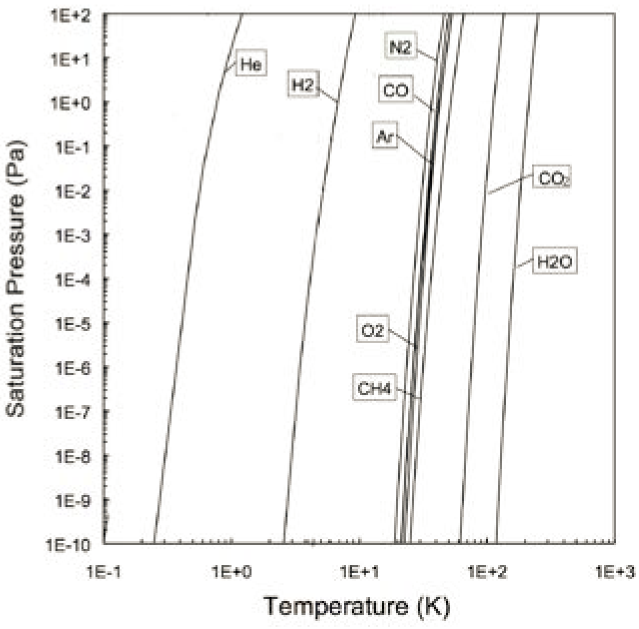

Figure 4.15: Saturation pressures of various gases important for UHV creation. Below about ≈ 16 K, the saturation pressures of most gases go below 10^{-12} mbar, making the partial pressure contribution of these gases on the ion lifetime negligible. This makes this temperature a good economical choice as, given the same heat load, 4 K is about an order of magnitude more expensive to maintain with a liquid helium flow cryostat. [27]

4.4.3 Bakeout Oven

For the UHV baking steps,[17] a custom oven from CDS Airtek was chosen[18], featuring a maximum operating temperature $T_{max} = 500°\text{C}$, 9 kW electrical heating power, a nominally available space of 95 cm width, 75 cm depth, and 80 cm height, forced air circulation creating a temperature uniformity of $\pm 5\,°\text{C}$ at a distance of $\geq 50\,\text{mm}$ from the chamber walls. It features three Ø80 mm access ports, three shelves with adjustable heights, and Ethernet enabled PID controller, and is shown in fig. 4.16.

(a) (b)

Figure 4.16: Bakeout oven setup. The bakeout oven (blue) takes up most of the available space, while the movable shelving (left of oven) houses all auxiliary vacuum equipment. (a) The CAD model used for planning the setup. Notably, to ensure that the oven fits both through the hallway and into its alcove, the relevant dimensions of the floor are measured and a basic floor model is recreated. Rendered with Solidworks Professional 2016. This CAD model can be compared to (b) an image of the oven setup as used in a bake. A titanium sublimation pump used for some of the bakeouts can be powered by the same Agilent 6671A high current supply (below oven) later used for the main field coils. Taken with an LG G5 camera.

Into the ports, a custom SS metal tube of the same diameter as the port is forced. These are intended to stay there permanently and seal in the oven's insulation material from human contact. Into these now smooth ports, custom SS metal cylinders with a smaller diameter than the ports, with lids on both ends, are inserted. They are hollow and feature an additional wall half way between the

[17] As this is a major piece of equipment, 16 separate companies were tracked down and quote requests sent. 8 of them sent a price, and 6 sent dimensions.

[18] Considerable effort is put into making sure the oven fits through the hallway and into the existing alcove.

warm and cold sides to decrease convective thermal conductivity. They are sealed against the port with over-braided glass rope “O”-rings, part PGL.40/HT from Pyroglass Ltd. For the feedthrough of cables, e.g. for thermocouples or vacuum pump powers, one plug features a slot; the cables’ heads are first fed through the port without the plug, then the cables are inserted to the slot, and the plug is pushed into the port. Any gaps are then sealed with crumbled aluminium foil and more glass rope. To attach the experiment’s vacuum system to the outside vacuum infrastructure, a short CF40 nipple is used, again sealed with glass rope surrounding it. Inside the oven, this connects to the experiment via a flexible CF40 hose. Outside the oven, a 6-way nipple is used to connect to the bakeout oven’s turbomolecular pump, ion pump, ion gauge, and residual gas analyser, which will be explained in this and the following sections.

For convenience, the height of the top shelf in the oven can be chosen to be 90 cm above the floor, in line with the 89 cm height of the clean room tables, and 91 cm of the optics tables. As there are no steps between these three locations, a trolley could potentially be used to safely transport the vacuum system. Instead, an aluminium stand, made of the 20 mm by 20 mm Bosch Rexroth (BR) series from KJN Automation Ltd, is used for both assembly and transport between cleanroom, oven, and laboratory. However, the aluminium alloys have melting temperatures not much higher than those found in the bakes, and the structural integrity could be compromised; therefore the stand is not used during the bakes. The vacuum system is placed directly on the shelving.

The temperature is monitored by an 8-channel Type N thermocouple hub, with thermocouples, part 3630301 from RS Components, attached to the oven’s ion pump[19], and to key points of anything inside the oven, e.g. any temperature sensitive parts.

The oven needs to be able to support the weight of the experiment and its

[19]Initially, to monitor the ion pump during baking outside the oven using heater tape, and later to monitor the room temperature.

connected parts. The experiment's mass is 28 kg, and the used Ti Sub nipple has a mass of 4 kg. Combined with various smaller additional masses, the shelves in the oven need to support $\approx 40\,\mathrm{kg}$. As the shelves are only rated to support a mass of 25 kg, two are used on top of each other; the top one is flexible enough to bow down to be supported by the bottom one. While this leaves little margin for additional weight, the oven could be upgraded in various ways to support more weight, e.g. by creating a sturdy custom shelve that hooks into the wall supports in multiple places. However, ≈ 10 bakes with experiments of similar masses have been completed without incidents.

As the experiment's vacuum system needs to be continuously evacuated during a bake, a dedicated system is designed. For initial evacuation from atmospheric pressure, a HiCube 80 Eco turbomolecular pump system from Pfeiffer Vacuum Ltd is used, part PM 015 887-T, combined with a venting valve, part PM Z01 290 B, and a PKR 361 active pirani/cold cathode gauge, part PT T03 350 010. For pressures below what the turbomolecular pump can achieve, a StarCell Vacion Plus 40 ion pump, part 9191240 from Agilent Technologies LDA UK Ltd is used; it has twice the pumping speed as the Vacion Plus 20 from the same supplier used in the experiment's vacuum system. The Vacion Plus 40 is connected with a 200°C bake-able cable to a MiniVac Controller, parts 9290705 and 9290290, sourced from the same supplier. Due to the limited vacuum conductivity between the experiment and the bakeout oven's ion pump, an additional pump is used inside the oven; a 3 filament titanium sublimation pump, part 360819 from Gamma Vacuum. Besides pumps, an ion gauge and a residual gas analyser (RGA), part RGA 100 from Stanford Research Systems, are included for diagnostic purposes, a Varian Multi Gauge controller from Agilent Technologies LDA UK Ltd reads the pressure from the ion gauge. The turbomolecular and ion pumps, and the RGA, can all be individually valved off by ZCR40R all metal valves by Vacgen, the same kind used in the experiment's vacuum system. Additionally, another such valve is located right outside the access port, separating the oven's vacuum system into a hot and cold side, for reasons elaborated on in sect. 4.4.9.

As the available space for the combination of bakeout oven and all of the just mentioned auxiliary equipment is very limited, the oven's size is chosen to both allow for this experiment's longer than usual size[20] and to still give enough space for the auxiliary equipment. As a result, the latter is stacked on multiple shelves, which are mounted on wheels, so they can be removed for better access. The combination of the bakeout oven and its auxiliary equipment can be seen in fig. 4.16.

The oven's Ethernet capability is used by an interface, written by another group member, to remotely initiate temperature ramps, and to monitor the temperature as measured by the oven itself. Additionally, the ion gauge's pressure, ion pump's current, and the thermocouple readings are logged.

4.4.4 Pumps And Gauge

Connected to the vacuum system is an RGA 100 residual gas analyser from SRS, capable of measuring masses from 1 to 100 amu. An ion gauge measures the pressure during the bake-out. An existing turbo pump is used for the rough pumping of the vacuum system. A second, stronger, ion pump is used to support the internal ion pump in the baking process. At the end of the last baking step, an ZCR40R all metal valve is used to shut separate the vacuum between the inside of the oven and the outside. The same kind of valve is installed in front of the turbo pump, the ion pump, the RGA, and the ion gauge. They are used to shut off flow for maintenance, and - in the case of the turbo pump - also during normal operation.

In order to choose an appropriately sized ion pump, it is useful to know the conductance between the pillbox and the ion pump. The ZCR40R valves are spe-

[20] As compared to the group's usual experiments; this is mainly due to the length of the existing cryostat

cified to a conductance of $C_V = 34\,\mathrm{l/s}$, at ambient temperature for N_2. Between the pillbox and the bakeout oven's ion pump, two valves are attached directly to the vacuum system, one between the in- and outside of the oven, and one directly in front of the ion pump, i.e. four in total. The contribution of the piping to the vacuum conductivity is shown in tab. 4.5. Using the transmission probability of the piping a_{tot}, and CF40 piping on both sides, the piping has a conductivity of $C_{\mathrm{p}} = a_{tot}\frac{\nu}{4}\pi r_{CF40}^2 = 13\,\mathrm{l/s}$, for which the gas velocity of N_2 is used from tab. 4.4. The contribution of the four valves is $C_{\mathrm{V,tot}} = \left(\frac{1}{C_V} + \frac{1}{C_V} + \frac{1}{C_V} + \frac{1}{C_V}\right)^{-1} = 9\,\mathrm{l/s}$. Combined, their conductivity is $C_{\mathrm{tot}} = \left(\frac{1}{C_p} + \frac{1}{C_{v,tot}}\right)^{-1} \approx 5\,\mathrm{l/s}$. As both the ion cleaning and flooding the vacuum chamber during opening between bakes uses argon, the ion pump needs to be able to handle noble gasses, too. Therefore, a Agilent Starcell is chosen, with a higher capability of pumping noble gasses than the standard ion pumps from Agilent, at the cost of a slightly smaller speed of pumping other gasses. A Starcell 40 has a nominal pumping speed of 40 l/s, and can pump N_2 at 34 l/s. Once saturated with Argon, it can still pump at C_{sat} =10 l/s, which combined with the conductance to the pump C_{tot}, gives a pumping speed of $C_{\mathrm{pump}} = \left(\frac{1}{C_{sat}} + \frac{1}{C_{tot}}\right)^{-1} = 3\,\mathrm{l/s}$.

4.4.5 Cleaning

In general, creating UHV conditions in experiments requires careful cleaning of all components that go into the vacuum chamber. The chemicals used in cleaning are:

1. Neutracon, a laboratory detergent, part 10355650 from Fisher Scientific UK Ltd

2. Deoxidine 624, aka Bonderite 624, a mixture of phosphoric acid with proprietary additions, part MP0003 from Las Aerospace Ltd

3. Acetone, $\geq$ 99.9% purity, HPLC grade, part 34850-2.5L-M from Sigma-

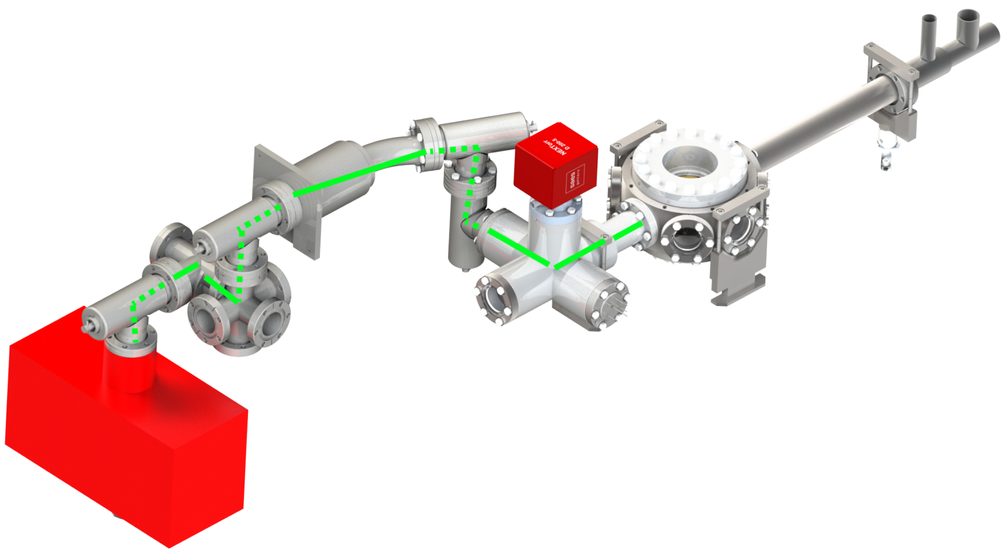

Element	**r**[mm]	**l**[mm]	**a**
5-way piece	20	87	0.35
	31.5	83	0.46
	31.5	83	0.46
Pipe	20	≈200	0.2
Pipe	20	≈100	0.32
Pipe	31.5	≈50	0.59
		a_{tot}:	0.083

Table 4.5: Vacuum conductivities model between the pillbox and the second ion pump, neglecting valves. From the pillbox' vacuum chamber, the first connection is a 5-way piece – modelled as a CF63 tube, then two CF40 tubes. Afterwards, there are two right angle valves in series, with their conductivities as specified in the text. Following the two valves, a $\sim 200\,\mathrm{mm}$ long CF40 hose was planned to then connect to a right angle valve (which separates the inside of the oven from the outside), which then uses two short CF40 tubes to connect to a final right angle valve. This last valve is used to keep the ion pump under vacuum while venting the rest of the system. In the table, N_2 at ambient temperature is assumed and the valves are omitted. The smaller radius, $r = 20\,\mathrm{mm}$, is that of CF40 piping, while the larger is that of CF63.

Aldrich Co Ltd

4. Isopropyl alcohol (IPA), $\geq 99.8\%$ purity, HPLC grade, part 34863-2.5L-M from Sigma-Aldrich Co Ltd

5. Deionised water (DI), 18.2 MΩ· cm resistivity, created with a laboratory water supply, part PURELAB Option-Q7 from ELGA LabWater

Generally, the parts to be cleaned are placed in a vessel just big enough to fit the part; this is usually a 1 litre glass beaker, an existing ⌀25 cm SS cooking pot[21], or a SS container with fitting sieve[22], parts K932 and GM319 from Vogue. For a lot of small parts, e.g. 50 M1.6 screws, existing SS tea strainers in a glass beaker are used to enable batch cleaning and drying. For the longest part, the cryostat, a custom SS vessel is designed.

The chosen vessel is then filled with the first chemical from the above list. It is then placed in an ultrasonic bath, part F0031802 from Ultrawave Ltd, filled with tap water at a temperature of $T = 30°\text{C}$ for $t = 5$ min. An ultrasonic power of 400 W is used, at a frequency of 30-40 kHz. Depending on which chemical was just used, they are then rinsed in DI and/or blown dry with compressed air. The overall process is shown in table 4.6, a picture can be seen in fig. 4.19a.

The compressed air used in table 4.6 is taken from the pre-filtered departmental supply. Although the filtering at the source would be sufficient, the problem is the piping to the clean room being standard household copper plumbing; this adds small amounts of contaminants, that are unacceptable to dry e.g. the ion trap's surface with. Therefore, a multi-stage filter, parts UD9+ and QD10+ from Atlas Copco, which are specified to remove oil to below 3 ppb, and particles to below 0.9 $\mu g/\text{m}^3$. Pressure is adjusted using a regulator[23], and measured using an analogue gauge, parts R07-200-RNKG and 18-013-989 from Norgren. To

[21] Importantly, with a smooth surface on the internal area

[22] These are usually used in gastronomy; the standard "gastronorm" size $1/2$ is used here, and accessories like sieves and lids are readily available

[23] As the compressed air is filtered afterwards, a standard, readily available regulator can be used.

avoid any further contamination after this filtering, clean plumbing from SMC Pneumatics UK Ltd is used, namely their KPH series of fittings in combination with plastic tubing of parts TPS1065BU and TCU0805B-1, and blow guns, parts VMG11BU-F02-02.

Most parts are cleaned using this process, although some parts require a modification, e.g. as some parts are incompatible with phosphoric acid. For example, using the above 5 step process, freshly cleaned OFHC Cu and BeCu oxidises in air within minutes to hours, even when thoroughly dried with compressed air. However, when the last step, DI water, is omitted, the surface stays unoxidised.

This is particularly important to keep in mind when working with BeCu. The oxide layer can be thick enough that the crimp connections used, see sect. 4.3.2, do not reliably break through. This can result in an intermittent electrical connection, while yielding a strong mechanical connection. Coupled with the only subtle difference in colour upon oxidation, this increases troubleshooting difficulty.

For PCBs that are soldered with flux, they are first cleaned in a solution of 5 g NaOH, 99 % purity, part GPC8035-Y from APC Pure, per 100 ml DI water. While submerged, the PCB to be cleaned is brushed[24], the solution is then replaced with fresh solution. This process is repeated several times. When no more bubbles are created, the PCB is then cleaned with the above chemicals. Due to concerns of potential damage to the PCB, this is done with manual instead of ultrasonic agitation.

4.4.6 Trap Cleaning Procedure

The ion trap itself is cleaned using a different process to the generic one just described. No ultrasonic agitation is used here, due to concerns whether the fine structures of the central electrode can withstand extended ultrasonic agitation.

[24]As this step comes before the above 5 step process, the used brush is an ordinary one that comes with a flux remover, part FLU400DB from Electrolube

Solvent	T[°C]	t[min]	Rinse	Drying
3% Neutracon in DI	30	5	DI	No drying
10% Deoxidine in DI	30	5	DI	Compressed air
Acetone	30	5	-	Compressed air
IPA	30	5	-	Compressed air
DI	30	5	DI	Compressed air

Table 4.6: UHV cleaning steps, based on the group's existing recipe. These are the generic steps followed for most parts, although some parts require extra consideration.

To achieve at least a small level of agitation, a pair of cleaned tweezers is used to move the ion trap in the solvents. 1-Methyl-2-pyrrolidone (NMP), 99 % purity, part 10773891 from Fisher Scientific UK Ltd, is used as the first step, followed by the solvents of table 4.6, except the first one. Each step is held for 1 minute. This process is repeated another time using fresh solvents. Afterwards, the traps are visually inspected for damage, first by eye, then using a microscope, part BX61 from Olympus, see fig. 4.17. Typical damage includes scratch marks due to slipping of tweezers; the microscope images are also used to confirm the feature size of the electrodes, as seen in fig. 4.3.

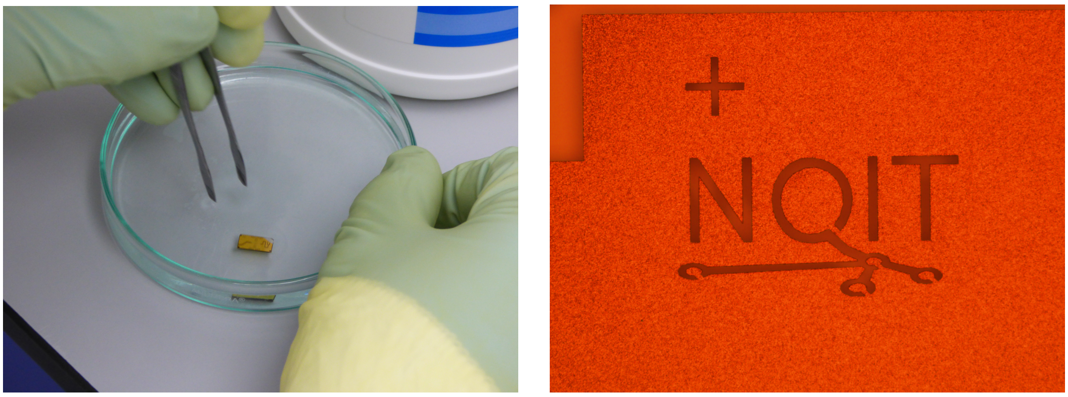

(a) (b)

Figure 4.17: Ion trap cleaning process. (a) The ion traps are first cleaned in a series of solvents, which are explained in the text. This particular ion trap comes from the edge of a wafer, as seen by the missing electroplating in the top right corner (just visible in this picture); it is just used to practice the cleaning process. (b) Afterwards, they are inspected for damage using a microscope. Shown here is the NQIT logo located in the top left corner of each trap. Taken with (a) a Panasonic DMC-FZ72 camera and (b) an Olympus BX61 microscope.

4.4.7 System Assembly

The entire assembly is carried out in the group's ISO class 7 cleanroom. Every tool used is cleaned with the method outlined in table 4.6. While working on the vacuum system, hooded coveralls (part TCX1), face masks (part FMNS), and latex gloves (part GLTNS1) are used, all from Connect 2 Cleanrooms Ltd. After cleaning, parts are wrapped in clean aluminium foil, part UHV-001500 from All-foils Inc, which is also used to create clean working spaces. An example of this assembly environment can be seen in fig. 4.19. CF flange bolts are tightened using a torque wrench, part 7000 A from Wera, using the torque recommended by Lewvac, $9\,\mathrm{N \cdot m}$ for CF16, $16\,\mathrm{N \cdot m}$ for CF40, and $20\,\mathrm{N \cdot m}$ for CF63-CF250.

The procedure for filling the oven tubes is as follows; as both calcium and strontium quickly oxidise in air, their time spent exposed to oxygen is minimised. Their vacuum sealed packaging is broken in a open-lidded box under a constant stream of argon to displace atmospheric oxygen, see fig. 4.18. Compressed argon comes as 99.998 % purity, part "pureshield" argon from BOC Ltd, but is contaminated with similar copper plumbing as the compressed air. It is therefore handled using an identical set of filters and subsequent plumbing, see sect. 4.4.5. The vacuum system itself is also flooded with argon. Using a scalpel, flakes of calcium are cut into small pieces, which are transferred to the vacuum system using tweezers. On the way, they are exposed to atmosphere for a few seconds, which is not enough to cause critical levels of oxidation. They are then inserted into the oven tubes, and packed more tightly using a stainless steel rod of slightly smaller diameter than the oven tubes. Using a pair of pliers, the top ends of the oven tubes are then crimped shut. The argon flooding is stopped, the vacuum system sealed, and an already connected turbomolecular pump is used to create a high vacuum to stop any further oxidation of the calcium. Afterwards, the outer valve of the vacuum system is closed, and the whole system is transferred to the bakeout oven.

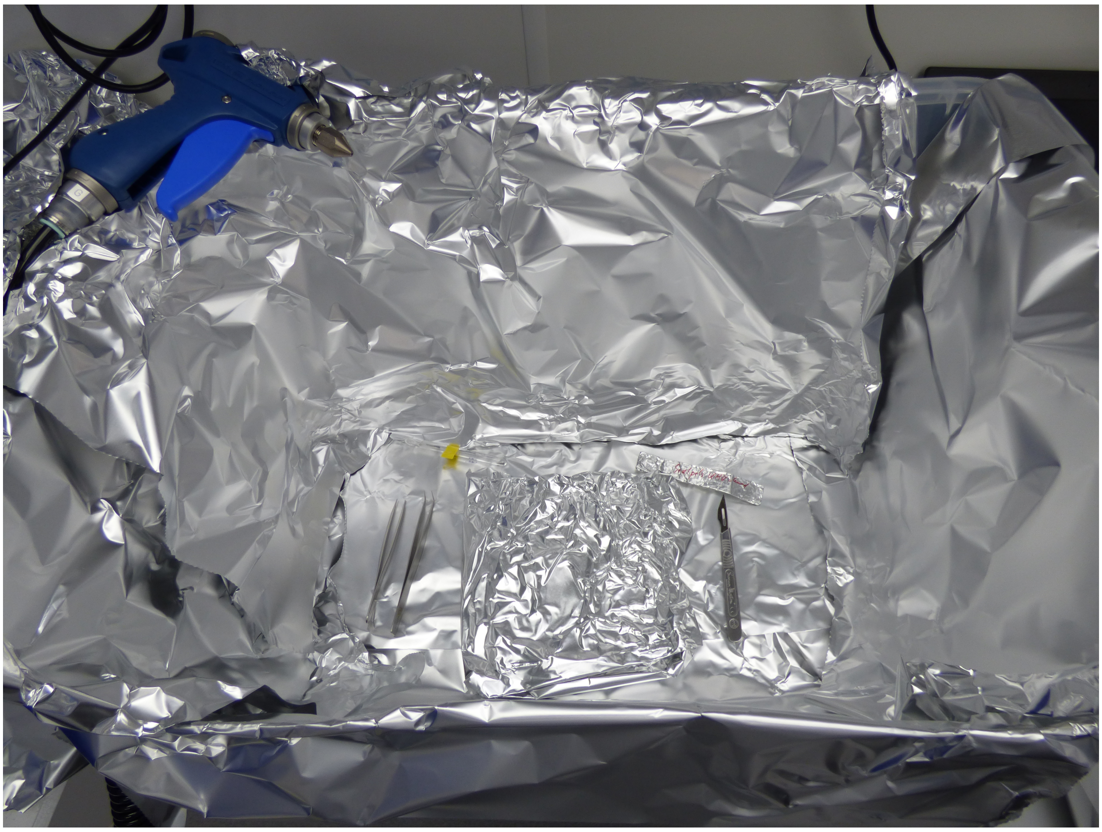

Figure 4.18: Argon environment enclosure for working with calcium and strontium in preparation of oven loading. Argon is continuously leaked into the box via the compressed argon gun (top left). The vacuum sealed glass vial (yellow tag) is then broken, and pieces of calcium cut apart small enough to fit into the ⌀1 mm oven tube. Both the plastic box and the aluminium cutting block are covered in clean aluminium foil. Taken with a Panasonic DMC-FZ72 camera.

The cleanliness achieved when using the cleaning process outlined in sect. 4.4.5, creates its own problems; for example when freshly cleaned parts made of similar metal compositions come into contact, cold welding can take place. This was observed when screwing cleaned SS parts together, e.g. on the screws fastening the μD connectors to the base flange (both on the air and vacuum sides), in which case they are removed using a left-handed drill[25]. Subsequent cold welding is avoided by the use of a UHV-compatible dry lubricant, MoS_2, part 330-MOS2-10G from Allectra Ltd.

4.4.8 Helium Leak Test

After the assembly of the system, and before baking, it is checked for leaks. At this stage, fixing a leak might be as simple as tightening the screws on a flange, or it might involve breaking vacuum. As opening the vacuum system to air always needs to be followed by a bake, it is preferable to check for leaks at this stage as opposed to after a bake. The method used to check for leaks is to spray helium on the air side of potential leaks, and to check whether the helium content in the internal high vacuum changes.

First, the experiment's vacuum system is connected to the bakeout oven's vacuum system, and the turbomolecular pump is used to create a high vacuum, as measured by the experiment's ion gauge, as explained in sect. 4.4.3. The RGA is then used to measure the partial pressures of various gases (identified by their masses) inside the vacuum[26]. As the amount of atmospheric helium is only ≈ 5 ppm (by volume), and because there should be no significant source of helium inside the system, its partial pressure should quickly drop below $\approx 10^{-13}$ mbar[27], the measurement sensitivity of the RGA. The RGA is then set for maximum sensitivity, while only measuring helium, as the partial pressure of other gases

[25]This is common method used to remove stuck screws

[26]This RGA only detects gases of ≤ 100 amu; however this covers all of the gases usually important for UHV, see table 4.4.

[27]If it does not, this in itself is a sign of a leak.

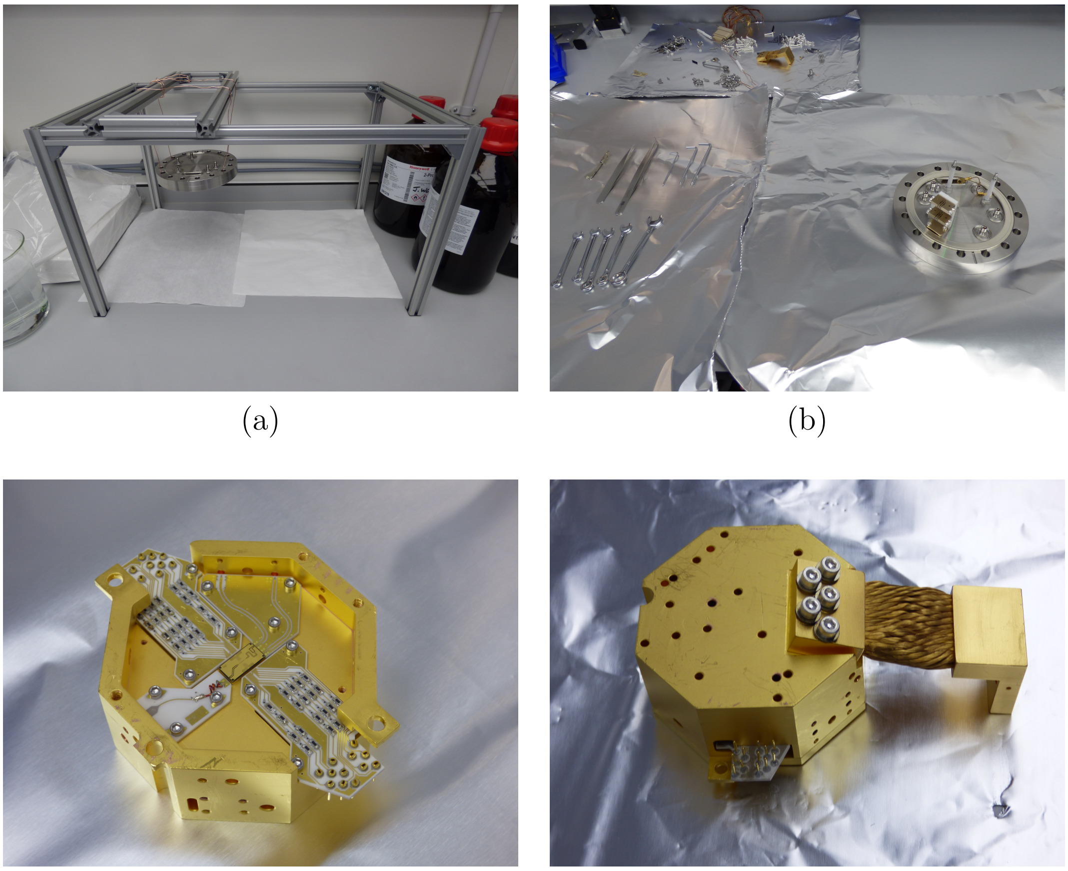

(a) (b)

(c) (d)

Figure 4.19: Clean room assembly environment. (a) During cleaning, parts are suspended using bare copper wires, shown here is the base flange. (b) For assembly, the tools undergo the same cleaning procedure as the parts going into the chamber. A partially assembled base flange can be seen on the right, tools for assembly are on the left, and additional parts to fit onto the base flange are in the back. (c) The pillbox with DC, μW, and RF PCBs attached, ready for wire bonding. (d) The thermal braid attaches to the bottom of the pillbox, secured with screws. As these are inaccessible after mounting the pillbox, they are carefully tightened to full torque; this also ensures good thermal conductivity. All taken with a Panasonic DMC-FZ72 camera.

would damage the RGA in this mode. Now, $\approx 99\,\%$ purity helium is sprayed onto the air side of any potential sites for leaks, i.e. the windows, electrical feedthroughs, the cryostat, ion pumps & gauge, valves, and onto all CF joints. The latter actually have a small recess for this purpose to allow helium to be sprayed directly onto the gasket. Any rise in the helium's partial pressure inside the vacuum can then be attributed to a leak sucking through helium from the currently sprayed air side; which could then be fixed. No leaks were found in any of the bakes.

4.4.9 Bake procedure

Unfortunately, thorough cleaning of the vacuum system on its own is not enough to create the target pressure of $< 1 \cdot 10^{-11}$ mbar; a subsequent bake is required. Table 4.1 lists seven possible sources of outgassing. Three of these sources can be treated using elevated temperatures while simultaneously maintaining the vacuum; namely diffusion, vaporization, and desorption are all processes that happen much quicker at the typical baking temperatures of 190°C and 350°C. The goal of speeding up this process is to deplete these outgassing sources during the bake, so that they are greatly suppressed upon returning to lower temperatures.

The higher the temperature used, the faster the outgassing happens; however not all components can withstand the higher temperatures. As a compromise, three stages of baking are used, with decreasing temperature.

The first stage of baking used hereafter is called an air bake; all stainless steel components capable of withstanding the oven's maximum temperature of $T_{max} =$ 500°C, are baked *without* assembly. This creates an oxide layer on all exposed surfaces, which is known to reduce their amount of outgassing. To preserve their cleanliness during baking, they are wrapped in the same kind of aluminium foil used in their cleaning.

The second stage consists of assembling all components capable of withstanding 350°C, notably excluding the base flange with ion trap and all windows, and is referred to as a hard bake. Instead, blank flanges are used to cover the CF40 and CF100 joints, parts FL-40CF and FL-100CF from LewVac Components Ltd. They are then connected to the bakeout oven's vacuum system, and continuously evacuated with the help of its turbomolecular and ion pumps. The temperature is then ramped up to 350°C, at a rate of 1.5 $^{\circ C}/_{\min}$. The pressure increases by several orders of magnitude while ramping up. While holding this temperature, the pressure decreases exponentially. The pressure P can be fitted to

$$P = a \cdot e^{-b \cdot t} + c \tag{4.11}$$

The pressure this fit is applied to is shown in fig. 4.20; however the fit is only applied to the pressures with the turbomolecular pump disconnected. From a starting pressure of $a = 1.1 \cdot 10^{-5}$ mbar, the pressure decreases with the constant $b = 0.51\,^1/_{\text{day}}$, i.e. by a factor of $^1/_e = 0.37$ in the first $^1/_b = 2$ days, asymptotically reaching the final pressure of $c = 7.8 \cdot 10^{-7}$ mbar. From this model, a rule of thumb can be derived; if pressure drops by $< 10\,^{\%}/_{\text{day}}$, the bake can be stopped, as there is little extra pressure reduction to be gained. For this system, the amount of pressure drop that can still be expected is $\frac{P-c}{c} = 20\,\%$, derived in sect. 7.4, for which it is not worth prolonging the bake. While ramping down the temperature, at some point the un-baked part of the vacuum system has higher pressure than the baked part. At this cross-over point, the valve between the two should be closed to avoid back-streaming of gas from the un-baked into the baked vacuum system. To find this point, the valve is periodically shut during the ramp-down; if the pressure inside the baked part decreases upon shutting the valve, the valve should remain closed. Once the vacuum system has thermalised to room temperature, a new fit of the pressure is created using equ. 4.11. The addition of more components in the next stage, in particular ones made of higher outgassing materials and harder-to-clean bits, means that the final pressure c

is likely to increase compared to that reached in this stage. Only if the final predicted pressure c is acceptable, i.e. $\approx 1 \cdot 10^{-11}$ mbar, the system is prepared for stage three.

The third and final stage of baking is referred to as the soft bake. After the hard bake, the vacuum system, once it is cooled down, is vented to air. All components inside the vacuum system are inserted, in particular the base flange with all electrical components, as well as all windows. As the last step before closing the lid on the vacuum system, the oven tubes are filled. Again, the oven's temperature is ramped up, but at a slower rate, e.g. 0.5 $^{\circ C}/_{\mathrm{min}}$, well below the CF40 window's specified maximum rate of 1 $^{\circ C}/_{\mathrm{min}}$. The same windows also limit the maximum temperature to 200°C; to give an extra margin of safety, the oven's temperature is set to 190°C.

An overview of the baking steps is shown in tab. 4.8. Both for the hard and the soft bakes, any components that can be actively degassed are treated thus. Generally, the method used for active outgassing is to slowly ramp up the current used while observing the system's pressure. When it spikes, hold the current until the pressure returns to ≈ 2 times the initial pressure. Then, continue the current ramp until the maximum desired current. An overview of this can be seen in table 4.7.

4.4.10 Bake protocols

After an initial air bake, the vacuum system is assembled and hard baked. Previously in the group, hard bakes have been performed at 350°C, however in this case the components are all rated to 450°C, and as previously explained higher temperatures greatly accelerate outgassing, so the baking temperature is increased to match. The bake itself created a satisfactory pressure, see fig. 4.20. However, the copper gaskets sealing the CF joints seized so strongly that a custom tool proved necessary to increase the leverage to ≈30 cm for CF40 joints and ≈50 cm

	Before	During	After
TSP	48 A, all 3 filaments	43 A on 1 filament, as needed	Valved off
Ion pumps	1 min on	1 min tests, then continuously on	Continuously on
Ca oven	2 A	-	2 A
Ion gauges	Controller's degass cycle	Controller's degass cycle	Controller's degass cycle
Getter	Activation cycle (60 W)	Conditioning cycle (20 W)	Activation cycle (60 W)

Table 4.7: Active outgassing steps of various components before, during, and after a bake. All three of the TSP's filaments are initially outgassed, although only one is used during the bake; it is valved off after the bake as it is located outside the experiments vacuum system. The ion pumps can only be run continuously at a pressure of $< 10^{-5}$ mbar; above that pressure they are outgassed by a one minute run every 24 hours. However, the NEXTorr's ion pump is not run during the bake as its magnets are not bake-able. The calcium oven is only outgassed before and after, as the used air side μD cable is not bake-able. The ion gauges are degassed every 12 hours using their controller's built-in degass cycle. The getter is activated by the application of 60 W of power for ≈ 3 hours before and after the bake, and conditioned with 20 W in between.

Type	T [°C]	T limited by	t
Air	≈ 500	Oven T_{max}	$1 \ldots 10$ hrs
Hard	≈ 350	Gasket seizing	$\frac{1}{2} \ldots 7$ days
Soft	≈ 200	Windows	$7 \ldots 14$ days

Table 4.8: Baking steps. The usual method is to air bake the bulk of the CF parts, as well as the Ca/Sr oven tubes, to create an outgassing reducing layer of oxides. This is followed by a vacuum hard bake of the assembled system, without its windows. Finally, the latter are added and the system reaches its final pressure.

for CF100 joints to remove them, and only combined with strong hammer taps, using a 0.3 kg hammer, part 200C.26 from Facom, was this successful. The reasons for this are likely a combination of factors. Firstly, the torque used for this bake was not well controlled, and likely too high. Secondly, the gaskets used are annealed and silver plated ones, parts CG40-39.2AGA and CG100AGA for CF40 and CF100 respectively, from LewVac Components Ltd. While the silver plating should decrease occurrence of seizing, annealing of the gaskets softens the copper. This increases the contact area between the gasket and the CF flange. Combined with a likely too high torque, and higher than usual temperature, these soft gaskets then seized. All gaskets are replaced and future hard bakes are planned to be performed at 350°C.

Next, the windows are added ready for a soft bake. As the components for the base flange were still under development, this soft bake was not the final one; rather it was used to check that the target pressure of $\approx 10^{-11}$ mbar can be reached, and for initial testing of the cryogenic system. An overview of the bake can be seen in fig. 4.21. Fig. 4.22 shows the experiment's vacuum system placed into a bakeout oven, ready for the first soft bake. The cryostat is only capable of being baked without damage if both the UHV side and the liquid cryogen sides are being evacuated. This necessitates a separate turbomolecular pump being connected to the liquid cryogen side; the UHV side is inside the vacuum system, i.e. being pumped using the main system pumps. This extra turbomolecular pump is connected to the exhaust port of the cryostat, the cryostat's other port is plugged with a custom closed nut, that compresses a PTFE ring. The exhaust port has a KF25 connection; unfortunately the standard o-ring material for KF, Viton, is only marginally compatible with extended soft bake temperatures. Therefore, ones made of the high temperature alternative material Kalrez are chosen, rated for continuous use up to 260°C. However, their quoted lead time was several weeks, so for the first soft bake, a PTFE version was designed

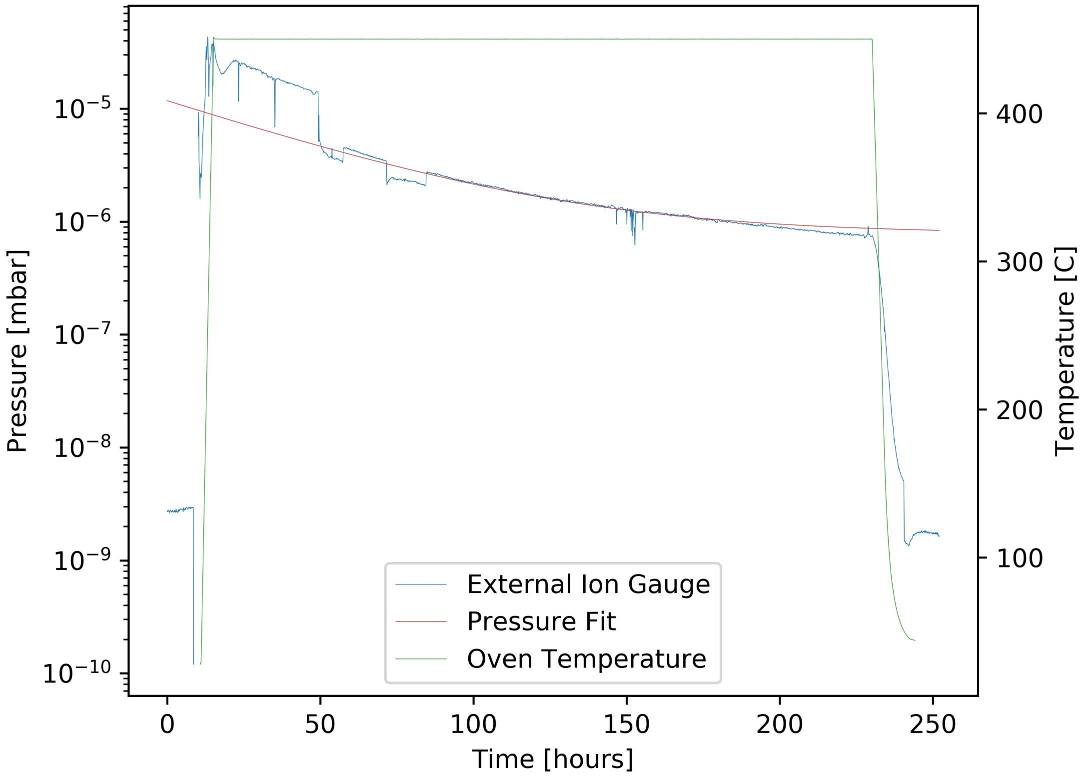

Figure 4.20: Pressure and temperature of the vacuum system during the hard bake. As the temperature of this bake is too high for standard cables, only external readouts are used. Before the bake, the pressure is measured at $3 \cdot 10^{-9}$ mbar. The temperature is then ramped up by $1.75\,{}^{C}/_{\mathrm{min}}$ to 450°C, causing the pressure to spike to $4 \cdot 10^{-5}$ mbar. For the first 100 hours of the bake, the pressure spikes/drops are due to the pumps being manipulated; the ion pump is turned on to check if the pressure is low enough for it to be left on. The turbo pump is valved off to check if it is still contributing to a lowering of the pressure, or if its base pressure is reached and increases the system's pressure. Upon ramping down by a nominal $1.2\,{}^{C}/_{\mathrm{min}}$, the pressure plummets to $5 \cdot 10^{-9}$ mbar, at which point the valve between the parts inside and outside of the oven is closed. This is to avoid the parts of the vacuum system that were not baked to increase the pressure of the parts that were baked. The final pressure achieved inside the baked parts of the vacuum system is assumed to be $< 10^{-11}$ mbar, as subsequent lower temperature bakes with additional in-vacuum components achieved this pressure.

and fabricated in-house.

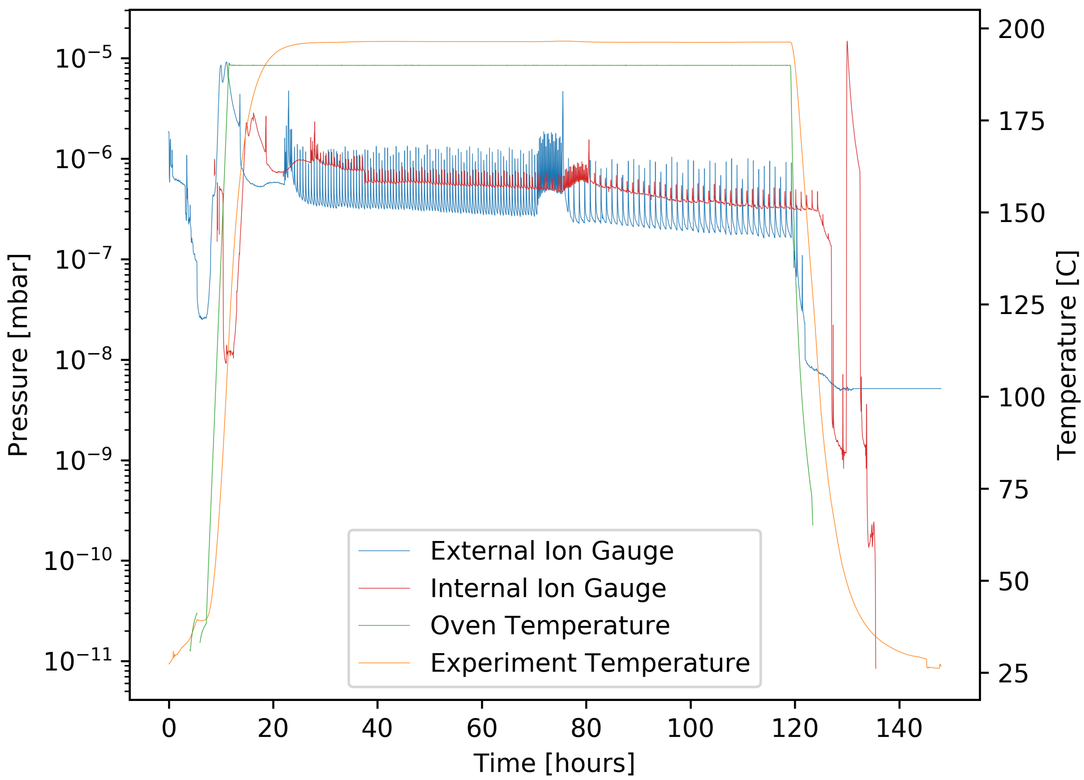

Figure 4.21: Pressures and temperatures of the vacuum system during first soft bake. In contrast to the hard bake, the pressure before the bake is much higher, $1{\cdot}10^{-8}$ mbar. After the temperature ramp of $0.5\,{}^{C}\!/_{\mathrm{min}}$ to 190°C, the pressure spikes to $\approx 10^{-6}$ mbar. In this bake, a titanium sublimation pump is included inside the oven. It is fired in regular intervals, causing the many pressure spikes seen between 20 and 120 hours into the bake. The temperature is then ramped down by a nominal $1.5\,{}^{C}\!/_{\mathrm{min}}$, although the oven's temperature significantly lags behind this. The experiment is insulated in aluminium foil, see fig. 4.22, to smooth out temperature changes. The effect is that the experiment's temperature, as measured by a thermocouple attached to a CF40 window, lags behind the oven's temperature by ≈ 2 hours. The final pressure is too low for the ion gauge to read, i.e. $< 10^{-11}$ mbar.

Once the components for the base flange are assembled, the vacuum system is vented to air, and the base flange inserted. Initially the plan was to only use the NEXTorr D200-5 for pumping. However, due to concerns of increased gas load coming from the harder-to-clean components on the base flange, an additional Vacion Plus 20 is added, which is attached at this stage. An overview of this

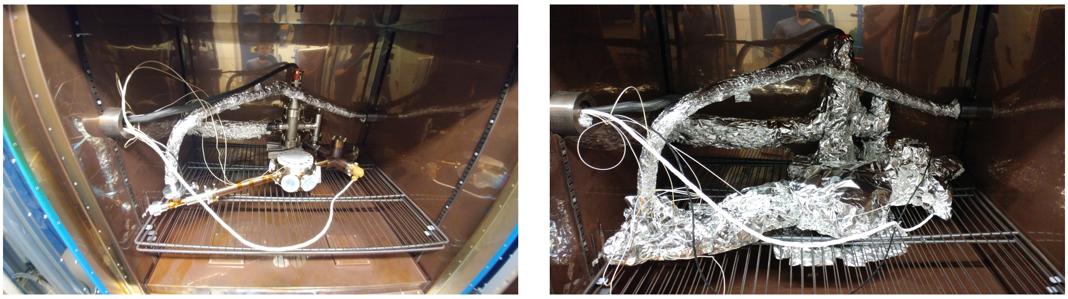

(a) (b)

Figure 4.22: The vacuum system is placed in the bakeout oven. In (a) the various electrical connections are visible; thermocouples (smaller white cables) are attached to various temperature sensitive points on the vacuum chamber to monitor compliance with maximum temperature specifications; mainly the CF40 windows are limited to 200°C. The ion gauge (bigger white cable), getter (smaller black cable), and titanium sublimation pump (bigger black cable) are all connected. In (b), the system is covered in several layers of aluminium foil to reduce the thermal conductivity between the oven's air and the vacuum system; this is done to smooth out any temperature fluctuations in the oven's temperature. Also, the plastic caps are removed before putting on the foil, as they are not rated to temperature. To avoid scratching the windows, small "tents" are built out of aluminium foil to make sure none of the foil touches the windows. Before these are put on, all windows are carefully de-dusted using filtered compressed air. Taken with an LG G5 camera.

second soft bake can be found in fig. 4.23. The pressure a few weeks after baking is below the measurement sensitivity of the ion gauge, $< 1 \cdot 10^{-11}$ mbar.

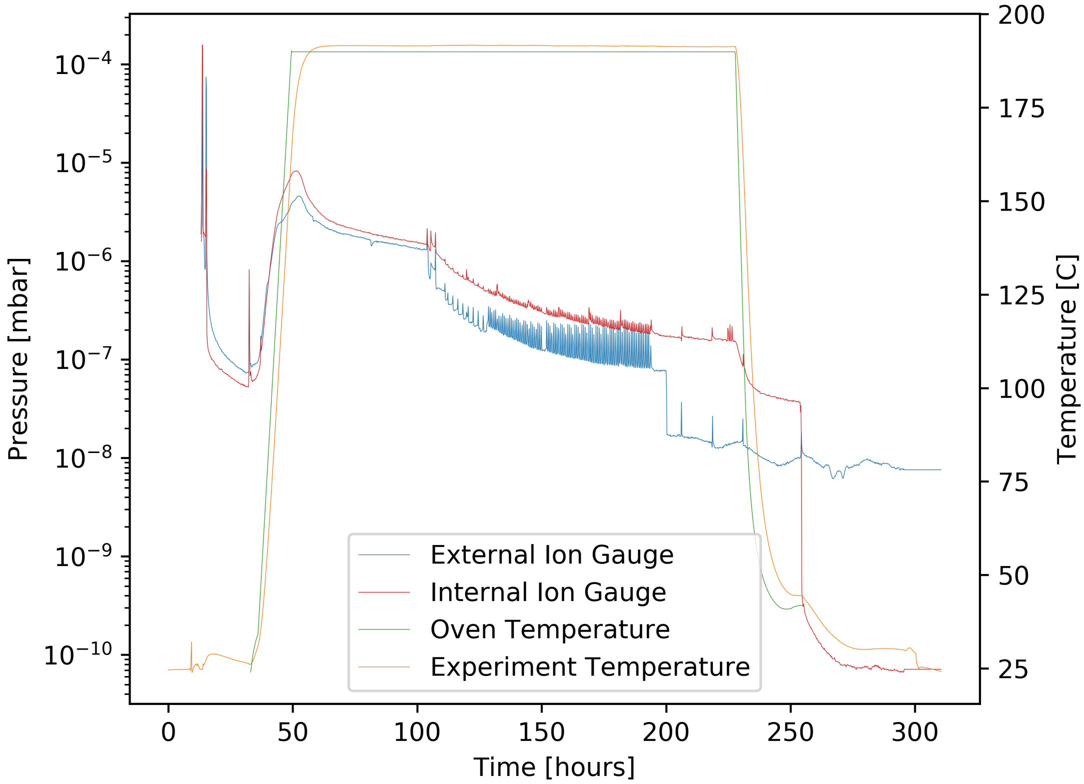

Figure 4.23: Pressure and temperature of the vacuum system during second soft bake. The pressure before the bake is measured at $5 \cdot 10^{-8}$ mbar. After the temperature ramp of $0.5\,{}^{C}\!/_{\mathrm{min}}$ to 190°C, the pressures during the bake are affected in the same way as the previous bakes, fig. 4.20 and 4.21. The final pressure reached directly following the bake is $7 \cdot 10^{-11}$ mbar, however in subsequent months the pressure became too low for the ion gauge to read, i.e. $< 1 \cdot 10^{-11}$ mbar.

Chapter 5

Experimental Results

In the previous chapters, it was described how the experiment was designed and fabricated. This chapter presents the results achieved so far, starting with a characterisation of the ion trap chip, sect. 5.1. The successful attempts at trapping ions and the initial results are described in sect. 5.2. The effect of various parameters on the loading rate is investigated in sect. 5.3. The stability of a few of the experiment's parameters is examined in sect. 5.4. In sect. 5.5 the trap's microwave magnetic field gradient simulations are compared to measurements. This is a crucial parameter for two-qubit gate operations.

5.1 Ex-Situ Ion trap characterisation

The ion trap design relies on a microwave resonator to amplify the magnetic fields that the ions are exposed to, see sect. 3.2. The resonator is built into the ion trap chip itself. While it is possible to simulate the resonance frequency of this resonator, fabrication tolerances and imperfections in the materials can lead to deviations of the resonance frequency from the model. For example, the resistance of the gold at cryogenic temperatures strongly depends on the purity of the electroplated gold, which is not a parameter measured in the fabrication process. To overcome this problem, a series of 8 different resonator lengths are fabricated, spaced 100 μm apart, labelled "-4" to "+4". The resonator labelled n is $n \cdot 100\,\mu m - \mathrm{sgn}\,(\mathrm{n}) \cdot 50\,\mu\mathrm{m}$ longer than the nominal design; the case $n = 0$ is omitted.

After fabrication, to test which resonator length is closest to the being resonant to the qubit frequency, $f_{1\to 1} = 3\,123.181\,794\,\mathrm{MHz}$, a PCB is designed to emulate the electronic structure surrounding the ion trap chip in-situ. The PCB, manufactured by Seeed Technology Co., is made out of 0.8 mm thick FR4, with a 35 μm Cu conductive layer, onto which Ni and then Au are plated, using the electroless nickel immersion gold (ENIG) technique. The PCBs are then bolted onto a OFHC holder. The PCBs' microwave lines are soldered onto an (unmodified) SMK connector, see sect. 3.7.1, to which a semi-rigid SMK cable, see sect. 4.3.3, is connected. In turn, the SMK cable is connected to a female-female SMA connector, and then to a high stability flexible SMA cable attached to a Vector Network Analyser (VNA).

Initially, the gold plating was intended to facilitate gold wire bonding from the PCB onto the ion trapping chip. However, the wire bonding was unsuccessful using gold wires. This was partially due to the ongoing wire bonding problems, see sect. 4.3.1. A thicker layer of Au would be needed for more reliable bonding. To avoid these problems of Au wires on ENIG, Al wires are instead used for this ion trap characterisation.

To roughly emulate the temperature, a styrofoam bucket filled with liquid nitrogen is used. The PCBs are first submerged in the liquid until bubbles are no longer rapidly forming, i.e. the PCB has reached thermal equilibrium. Afterwards, they are elevated to just above the liquid level, with the OFHC holder still partially submerged. This way, the PCB is still held closely at the boiling point of liquid nitrogen, 77 K, while the ion trap's resonator is surrounded by gaseous nitrogen, which, compared to liquid nitrogen, more closely resembles the vacuum this test is intended to emulate. The test setup can be seen in fig. 5.1.

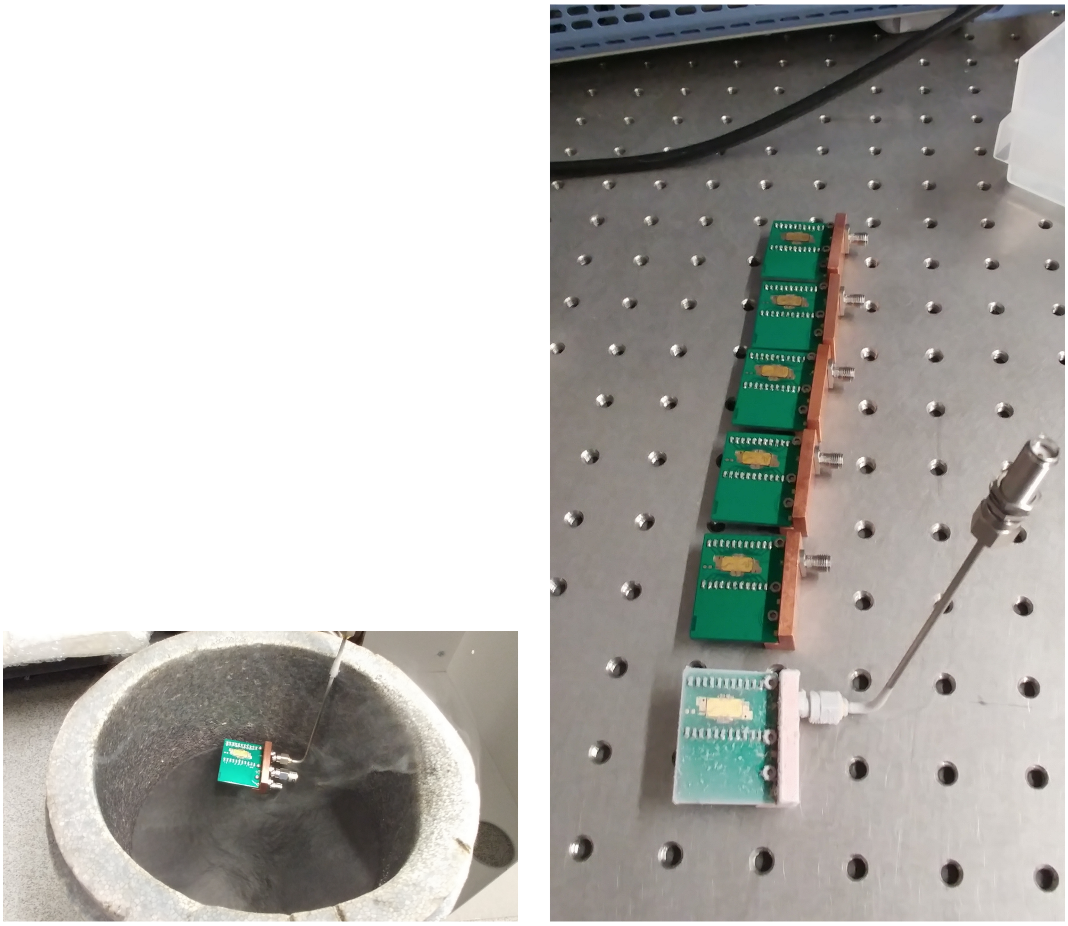

(a) (b)

Figure 5.1: Liquid nitrogen testing of ion traps. (a) The traps are soldered and wire bonded onto "mock-up" PCBs, described in sect. 3.8, that closely resemble the electronic structure at microwave frequencies surrounding the ion trap. They are then connected to the same type of SMA cable used in the experiment, and submerged in liquid nitrogen. After thermalisation, indicated by bubbles stopping to form, the traps are held just above the liquid level. The traps are not submerged, but the PCB holder, made of copper, is partially submerged and acts as a heat sink. While in this state, a measurement of the reflected power, S_{11} is taken, see fig. 5.2. (b) As there are 8 (two are not shown) different resonator lengths, this procedure is repeated for each. The bottom trap is still cold from a recent measurement. Both taken with an LG G5 camera.

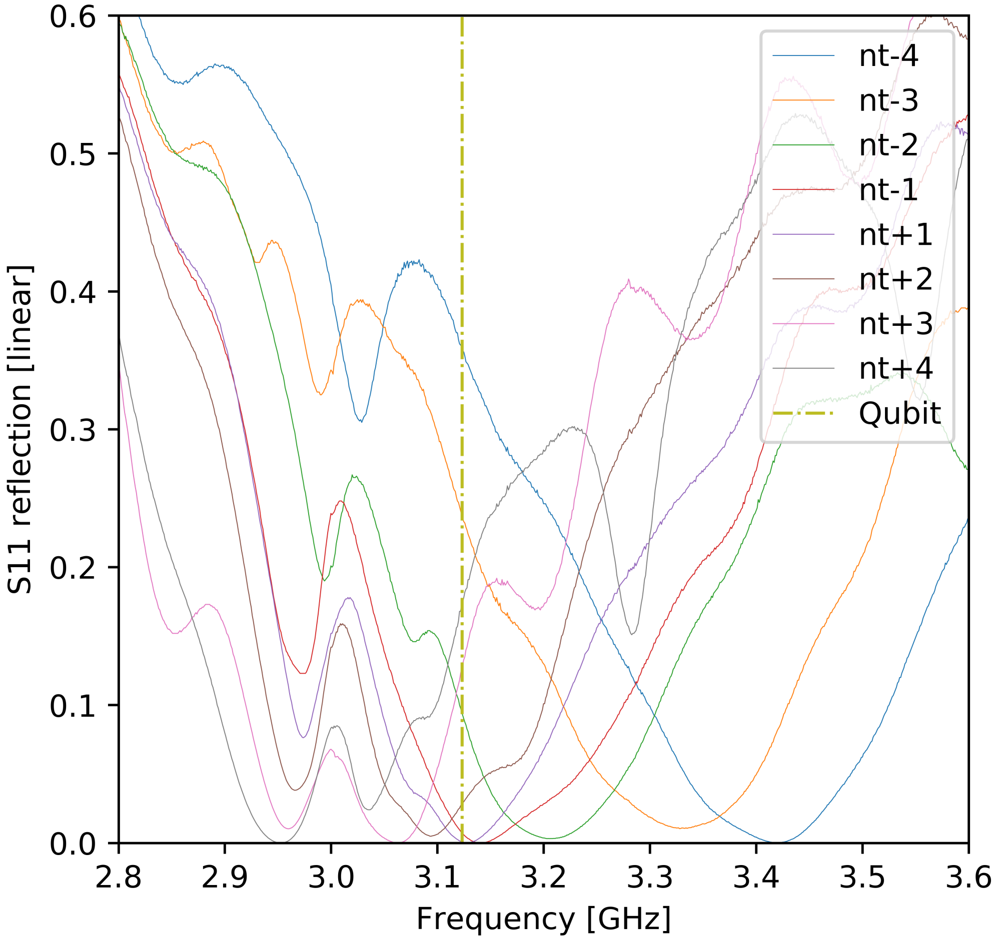

(a)

Design name	-4	-3	-2	-1	+1	+2	+3	+4
Offset from nominal [μm]	-350	-250	-150	-50	50	150	250	350
Offset from qubit frequency [MHz]	312	210	84	18	5	-28	-60	-170

(b)

Figure 5.2: S_{11} measurement of all 8 different resonator lengths of the ion traps, taken at a temperature of 77 K. The designs "+1" and "+2" come closest to the qubit frequency $f_{1->1} = 3.123182\,\text{GHz}$. Taken with a Rohde & Schwarz ZVL13 Vector Network Analyzer, and calibrated using a Rohde & Schwarz ZV-Z135 calkit to reduce the influence of the wiring on the measurement.

As a rule of thumb, the minimum acceptable trap design would reflect at most half of the incoming power; this could then be countered by increasing the microwave power applied to the resonator input. This would be acceptable as most of the power that is reflected from the trap is transmitted back outside the vacuum chamber, and would therefore not increase the heat load on the cryogenic system. However, a smaller amount of reflected power would reduce the required microwave amplifier output power.

VNA measurements are taken for all 8 resonator lengths and shown in fig. 5.2. However, due to the low Q of the resonators, all resonator lengths reflect less than half of the incoming power. This means that, at least for liquid nitrogen temperatures and the imperfect facsimile of the "mock-up" PCBs, any of the designs would be viable, although trap designs with low-reflecting are still preferable. The observed shift of the frequency steps due to the 100 μm resonator length step size is $\approx$ 60 MHz on average. In addition, there are further reflections in the measurement setup, and a resulting 154 MHz repeating pattern of local peaks in reflections, see 5.2b. The overall FWHM of each resonance is 400 ± 50 MHz, resulting in a Q factor of 8 ± 1, therefore the resonator length only needs to be correct within $\pm 200\,\mu$m. The resonator length "+1" is chosen as the final trap to be used.

5.2 Initial trapping

After choosing the ion trap's resonator length, it is bonded onto the pillbox, see sect. 4.2.2, and then wire bonded to the surrounding PCBs 4.3.1. Afterwards, the vacuum system is assembled, see sect. 4.4, moved to an optics table, and laser beams are directed towards it. After testing all electrical connections for continuity and resolving various problems related to starting a new experiment, preparations are made to trap an ion for the first time in this trap. As there are quite a few potential causes for not seeing any trapped ions appear, the first step is to look for fluorescence of neutral Ca.

5.2.1 Neutral Fluorescence

As mentioned in sect. 2.4, this work uses a two step ionisation process. The first step uses a 423 nm laser to excite the $4s^2\ ^1S_0$ to $4s4p\ ^1P_1$ transition in neutral Ca. The frequency detunings for the 423 nm laser are referenced to the existing wavemeter setting of 709.07812 THz, which is different to the 709.07824 THz measured by NIST; presumably at one time this was measured to be the accurate transition frequency for ^{40}Ca for this wavemeter. As this is a shared resource across multiple experiments, and each experiment having its chosen detuning referenced to this parameter, the wavemeter reference is not corrected. Furthermore, since this wavemeter has not been calibrated for multiple years, all 423 nm detunings reported here should be taken with some uncertainty.

To create the neutral Ca, the atomic oven, see subsec. 3.3.4.2, is heated to nominally 330 C. The 423 nm laser's frequency is then scanned, and the photons arriving at the PMT counted, see sect. 3.13. The resulting photon counts can be seen in fig. 5.3.

The observation of neutral fluorescence in this trap checks, among other things, that:

- The atomic oven is indeed creating neutral Ca
- The 423 nm laser's frequency can be focused onto roughly the right area[1]
- The imaging system, and the experiment's control system, are functioning and able to count photons

[1]However, as the imaging system's shutters initially are completely open, the fluorescence could be coming from an area outside the trapping region.

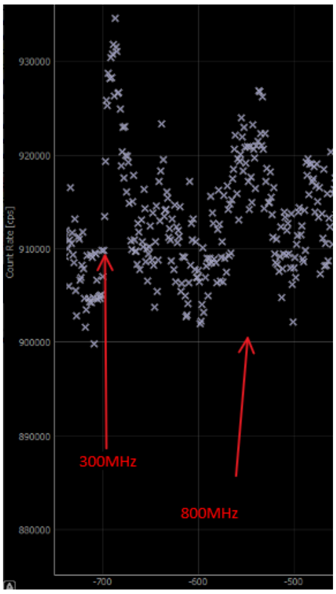

Figure 5.3: Frequency scan of the 423 nm laser, measuring the fluorescence of the neutral Ca atoms, at a magnetic field of $\approx 0.3\,\mathrm{mT}$. The fluorescence in counts per second is plotted against time; here the 423 nm laser is set to slowly change frequency in a triangular pattern, and is measured asynchronously. The frequency is changed from 0 ± 200 MHz to 1000 ± 200 MHz, and this is synchronised by eye to find the resonances. More accurate measurements are described in sect. 5.3. For 300 ± 200 MHz, a strong resonance is found, and for 800 ± 200 MHz, a weaker one. These correspond to ^{40}Ca and ^{43}Ca, respectively.

5.2.2 ^{40}Ca

After having observed neutral Ca fluorescence, the rest of the lasers are also directed towards the trapping region, and the RF Paul trap and corresponding DC electrodes, are turned on. As $^{40}\mathrm{Ca}^+$ is both a simpler system compared to $^{43}\mathrm{Ca}^+$, and more abundant in the atomic oven, it is first attempted to trap $^{40}\mathrm{Ca}^+$.

The minimal set of lasers, apart from those for photo-ionisation, used for trapping in this thesis is a 397 nm and an 866 nm. However, to avoid any problems with dark states, see sect. 2.6, an 850 nm laser and an 854 nm laser are used instead of the 866 nm. At this point in the experiment, beam paths for the 854 nm and 866 nm lasers already exist; as all lasers going to the experiment are fiberised, the fibres for the 866 nm laser and 850 nm laser are simply switched. As the wavelengths of the three red lasers are so close to each other, dispersion in the optics is minimal, and neglected.

Notably, only one beam path is necessary to trap and cool $^{40}\mathrm{Ca}^+$ and $^{43}\mathrm{Ca}^+$.

To avoid any problems of too low a laser intensity at the ion, owing to imperfect beam pointing and/or detuning, the laser powers used for initial trapping are fairly high.

After systematically trying various free parameters, e.g. the RF input power (which affects the trap's stability parameter q), laser frequencies, and precise location of the laser foci, trapping attempts proved successful. Subsequently, the 866 nm laser is used instead of the 850 nm laser, in line with the cooling scheme, see sect. 2.6.

To distinguish a trapped ion from background scatter, the red lasers are turned off for a short time[2]. This results in the ion being rapidly ($\ll 1\,\mu$s) shelved in $3\mathrm{D}_{3/2}$. To get the signal that is due to the ion, the signal rate received while red lasers are off is subtracted from that when they are on. This background-

[2]This is done either by manually blocking the red beam path, e.g. by hand, or by switching the AOMs off. Uncooled, the ion is seen to remain trapped for up to ≈ 1 min, but this is not systematically measured.

subtraction method is used throughout the work described in this chapter.

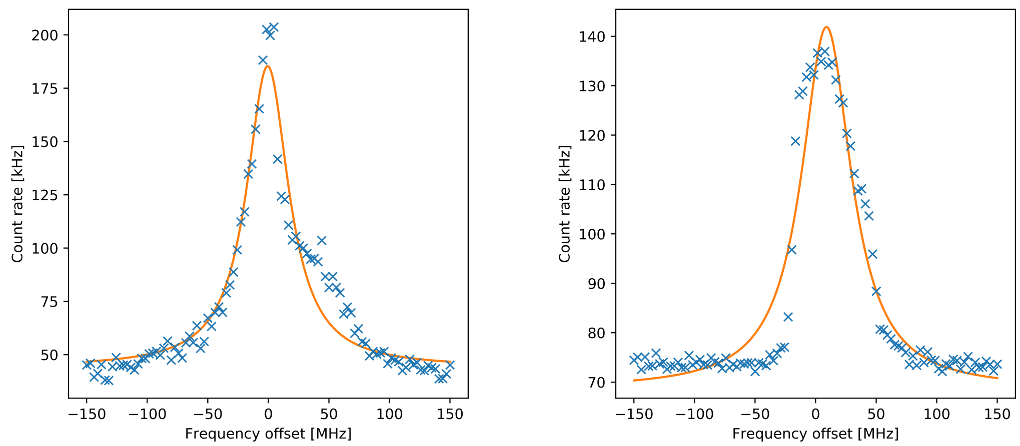

(a) (b)

Figure 5.4: Frequency scans in $^{40}Ca^{+}$, plotted against count rate measured by the imaging system. Frequency is changed by adjusting the frequency of the double-passed AOM, the x-axis gives the resulting change in laser frequency of the (a) 397 nm laser and (b) 866 nm laser. Both scanned by ±150 MHz, from the previously established resonance, fitted here to a Lorentzian lineshape. The deviation of the scans from a Lorentzian is due to $^{40}Ca^{+}$ not being a two level system; if the detuning of the 397 nm and the 866 nm lasers are equal, the lambda system consisting of $4S_{1/2}$, $4P_{1/2}$ and $3D_{3/2}$ creates a dark resonance. In the resonant case, this decreases the fluorescence, which can be seen here at (a) +20 MHz, and (b) -30 MHz. Note that in (a), the fit is slightly misleading; what appears to be an increase in fluorescence at +50 MHz is actually a dip in fluorescence at +20 MHz over the case without a dark resonance. Dark resonances could be avoided by switching to a cooling scheme using the 850 nm and 854 nm lasers instead of the 866 nm laser.

The first parameters measured using the ion are the laser wavelengths involved in the cooling scheme, see fig. 5.4. In order to scan a laser's frequency, and measure the resulting fluorescence of the ion, a simple method would be to slowly adjust the laser from one frequency to the next. However, this causes two problems; such a method makes it harder to distinguish other effects, e.g. a slowly drifting laser intensity, from the effect to be measured. Secondly, in the case of the 397 nm laser frequency, scanning blue of resonance turns the cooling scheme from sect. 2.6 into a heating scheme, quickly resulting in an increase in the ion's thermal energy above the trap's depth, and ultimately a lost ion.

In order to avoid these problems, the following method was adopted. Let f_0 be a laser frequency suitable to cool the ion by laser cooling. To probe the fluorescence at any other given frequency f, the frequency of the light incident on the ion was quickly switched to f using an AOM, and then switched back to f_0 again after accumulating fluorescence for a short time of order 10 ms.

This way, the ion only receives a small increase in temperature, not enough to be heated out of the trap, and is immediately cooled afterwards. This also removes any effect of hysteresis, as all frequency measurements all start out with a freshly cooled ion. The order of frequency measurements is also randomised to distinguish outside influences against the one to be measured. For scanning laser frequencies scans via AOMs are preferable to implement this randomisation. AOMs have a certain frequency, e.g. 260 MHz, with a peak diffraction efficiency, which decreases the further the applied frequency differs from it. There is a feedback system which stabilises the laser power arriving at the ion, see fig. 3.21; this system also compensates for of decreased diffraction efficiencies by applying additional power to the AOMs. For scans beyond a range of $\approx \pm 100$ MHz, the diffraction efficiency of the AOMs drops so much that the power cannot be further increased without risking damage to the AOMs.

These wide-range scans are then done via scanning the piezos of the lasers' cavities. However, these scans are then taken taken in non-random, ascending order due to the inherent hysteresis of piezo drives. The maximum range of these scans depends on the maximum piezo voltage, which translates to typically $\approx \pm 1$ GHz, although this is in practise typically limited to $\approx \pm 0.5$ GHz for the blue lasers and $\approx \pm 1$ GHz for the red lasers.

5.2.3 ^{43}Ca

In order to trap $^{43}Ca^+$, first $^{40}Ca^+$ is trapped at a magnetic field of ≈ 0.3 mT, referred to as "low field", in contrast to ≈ 28.8 mT, "high field". The cavities are then adjusted by the isotope and hyperfine shifts from tab. 2.2. This necessitates an EOM frequency beyond the range the high field EOM can supply, so a different one is used for low field. Due to the higher number of states, and the wider spreading of states in $^{43}Ca^+$ compared to $^{40}Ca^+$, the laser intensities are increased; $^{43}Ca^+$ was then successfully trapped. After some initial low field measurements are taken, see sect. 5.2.7 and 5.5, the EOM is changed to high field. As changing the EOMs involved realigning various optics, no additional low field work is subsequently done with $^{43}Ca^+$.

For high field, typical laser frequency scans can be seen in fig. 5.5.

5.2.4 Signal to background ratios

A large number of parameters influence the fluorescence signal to background ratios (SBR). While the laser detunings and intensities are comparatively easy to reproduce, the optics involved typically need manual adjustment. The imaging system for example has 5 degrees of freedom (x, y, z axes, and tip/tilt) on the primary lens, and 3 degrees of freedom on the secondary lens. For the alignment of the cooling lasers, controlling the last mirror's tip and tilt angles via piezo drives (which to first order mainly changes the position of the beam at the ion) is certainly helpful, the angle of these beams still needs to be optimised to avoid scattering too much laser light off the surface of the chip into the imaging system. Also the position of the focusing lens needs to be optimised.

All of these degrees of freedom need to be controlled and optimised; and in particular those that are optimised by hand usually could be optimised further. In this work, the achieved signal to background ratios for $^{40}Ca^+$ are about 20:1, and 5:1 for $^{43}Ca^+$, both at a magnetic field of ~ 0.3 mT. For $^{43}Ca^+$ at a higher

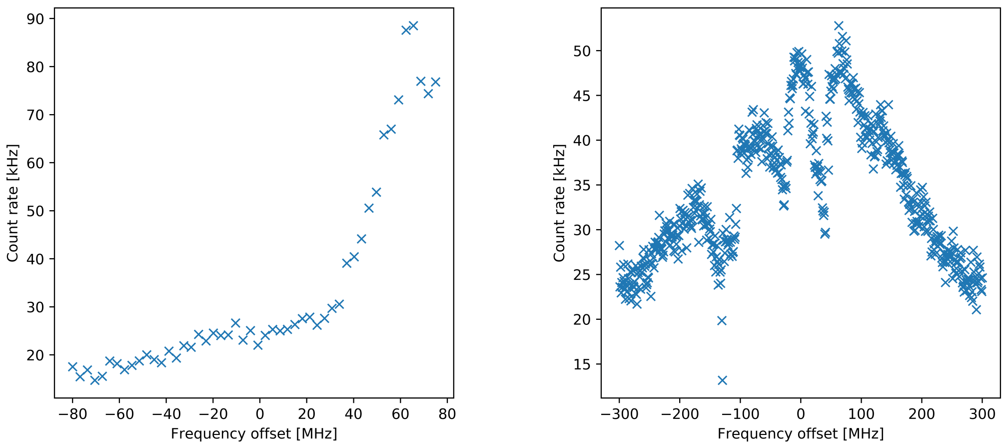

(a) (b)

Figure 5.5: Frequency scans in $^{43}Ca^{+}$ at a magnetic field of $\approx 28.8\,\mathrm{mT}$, plotted against count rate measured by the imaging system. Frequency is changed by adjusting the voltage of the piezo inside the reference cavities the lasers are locked to. The x-axis gives the resulting change in laser frequency of the (a) 397 nm laser and (b) 866 nm laser. As the method of adjusting the frequency via the cavity piezo is comparatively slow, the 397 nm laser cannot be scanned blue of resonance without heating the ion out of the trap; hence the 397 nm laser scan is only shown red of resonance. Compared to fig. 5.4, the 866 nm laser scan shows a much more complicated structure with many such dark resonances, owing to the hyperfine structure of $^{43}Ca^{+}$. This is also why fits have not yet been implemented for these scans. The parameters used to obtain these scans have been optimised for high fluorescence and well-resolved dark resonances, and will be presented in sect. 5.2.4.

magnetic field of ~ 28.8 mT, the parameters influencing the fluorescence rate are further optimised to achieve 10:1. In this thesis, laser intensities are given as multiples of a saturation intensity I_0, defines as the intensity where the rates of spontaneous and stimulated emission are equal.

Specifically, these parameters are a 397 nm laser power of $P_{397} = 400\, I_0$ at a nominal detuning of $\Delta f_{397} = -20\,\mathrm{MHz}$, an 866 nm laser power of $P_{866} = 1500\, I_0$ at a nominal detuning of $\Delta f_{866} = +120\,\mathrm{MHz}$, and a nominal EOM power of $P_{\mathrm{EOM}} = 0\,\mathrm{dBm}$ at a frequency of $f_{\mathrm{EOM}} = 2603\,\mathrm{MHz}$. Examples of typical SBRs and the resulting thresholding parameters are shown in fig. 5.6. These fluorescence rates and signal to noise ratios are sufficient for all experiments attempted in this thesis.

5.2.5 Beam pointing

Laser beam pointing onto the ion is a crucial part of the experiment. Typical beam waist diameters are $\approx 20\,\mu$m, while the ion's height above the surface is $40\,\mu$m, so any deviation from the optimal focus position and beam pointing leads to higher rates of scatter of the trap surface and/or lower intensity at the ion. Therefore, using a pair of mirrors to "walk" the beam into a path parallel to the surface, and subsequently getting the lens position and thus the focus within a Rayleigh length of ≈ 2 mm is an important first step in the alignment of the lasers.

As laser beam pointings tend to drift over time, e.g. due to thermal fluctuations, they periodically need realignment. As there are 5 beams paths[3] planned in this experiment, this task is automated. As detailed in sect. 3.10, on every beam path there is a combination of two mirrors on kinematic mounts before the laser is focused onto the ion; one of those mirrors is also controllable via piezo drives. For the cooling beams, the voltage on one axis is scanned[4], and the res-

[3] 1 for cooling, 2 for state prep and readout, and 2 for raman spectroscopy

[4] To minimize the effects of hysteresis, the scan is simply a stepped linear increase in voltage,

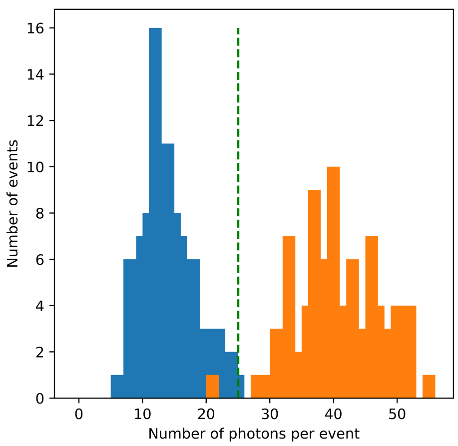

(a)

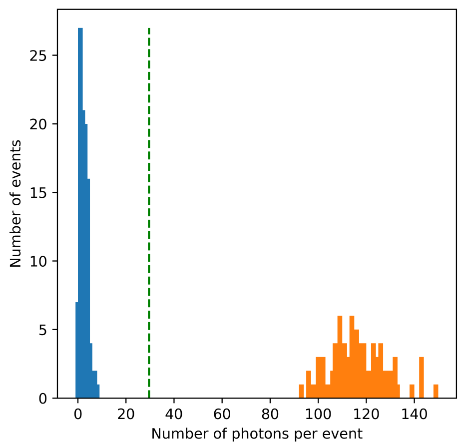

(b)

Figure 5.6: Thresholding examples. To distinguish between a bright (orange) and dark ion (blue), a simple thresholding method (green) is used. To get a bright ion, it is simply made to fluoresce using the laser parameters used for readout. For a dark ion, only the red lasers are turned off, which do not affect the PMT, but result in both $^{40}Ca^{+}$ and $^{43}Ca^{+}$ being shelved onto the $3D_{5/2}$ states, effectively eliminating its fluorescence. Here, for a fixed time of $t = 1\,\text{ms}$, photons are counted with a dark ion, and then for another $t = 1\,\text{ms}$ with a bright ion. The process is then repeated 1000 times to gather shot noise statistics and then plotted here. The x-axis shows bins of a particular number of received photons. The y-axis shows how many times this particular number of photons has been received. These results are then fitted to Poisson distributions. The optimal threshold to set is where the fits of the two distributions intersect. Shown here are examples of comparatively (a) bad and (b) good discrimination on $^{40}Ca^{+}$.

ulting fluorescence and background scatter rates measured for each voltage. A Lorentzian function is then fitted onto the fluorescence to determine the peak, the mirror's piezo in this axis is then set accordingly. The piezo drive has a hysteresis of $\approx 10\%$ of the range travelled, this effect is minimised by only scanning a small distance, about a beam waist radius in either direction of the last known optimal; therefore the resulting error is small and neglected.

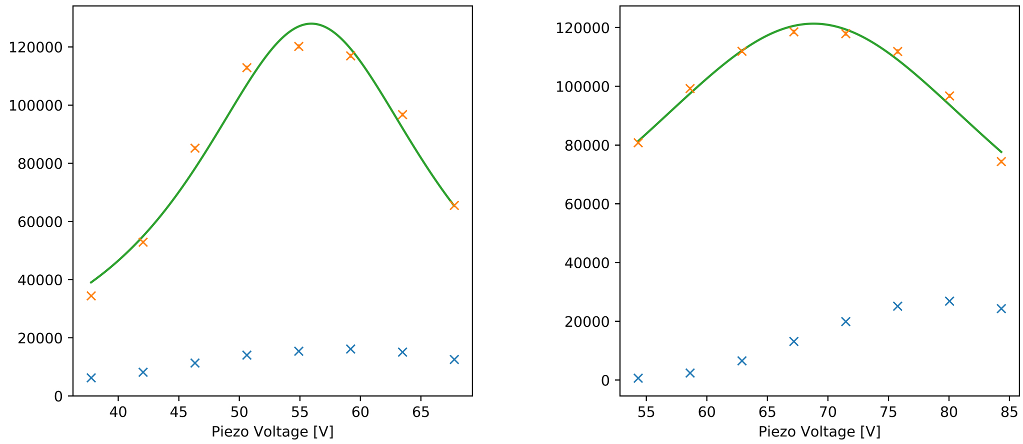

(a) (b)

Figure 5.7: Beam pointing scans in $^{43}Ca^{+}$, plotted against count rate measured by the imaging system. The x-axis shows the voltage applied to the piezo drive inside the mirror mount to the (a) x-axis (horizontal) and (b) y-axis (vertical). The detected fluorescence rate (orange) is fitted to a Lorentzian lineshape (green) and contrasted to the background detection rate (blue). As a rough guide, expected to be accurate to a factor of 2, the beam position at the ion changes $\approx 1\,{}^{\mu m}/_{V}$, depending on the exact path the lasers take through their focusing lens. For the vertical direction, it is expected that a lower beam height results in a much higher background count. For the horizontal direction, the background count mainly depends on the imaging system's shutters; if they are wide open, then the background should not change, if they are nearly closed, the background should change similarly to the ion's fluorescence.

An example of a typical beam pointing scan can be seen in fig. 5.7.

This process is then repeated for the second axis of that mirror. For the other beam paths, different methods for determining the laser's intensity are used; e.g. for the 393/397-σ_+ beam path, the 393 nm laser's shelving speed is used. Here,

as opposed to using a randomised order.

the shelving time is deliberately chosen to be insufficient for reliable shelving, e.g. by using $1/10$ to $1/2$ the usual time. Therefore, when trying to shelve, there is only an e.g. 50% chance the ion will be shelved, even when the laser is perfectly aligned, see sect. 2.7. The shelving probability is then very sensitive to the laser's intensity at the ion's position. The sequence of state-prep, shelving, and readout is then repeated e.g. 100 times to receive a shelving probability. Another Lorentzian[5] is then fitted in a similar way to that of the cooling beam path, and the mirror's piezo voltage set.

Similar methods are used to servo the other beam paths. For a long term (48 hours) study of the stability of the beam pointing, see sect. 5.4.

5.2.6 Motional frequencies

To measure the motional frequencies of the ion, an RF signal is applied on a DC pad of the ion trap. To decouple the DC level from the RF, a transformer is used, part WBC1-1LB from Coilcraft. The choice of DC pad affects how strongly the RF signal will couple to which of the ion's three axes of motion. An overview of the trap's electrodes was shown in fig. 3.4c, but for easier reference another overview is shown in fig. 5.8.

Choosing electrodes DC_left_6 or DC_right_6 results in the strongest coupling to the radial motional modes, but has no coupling to the axial motional mode. For experimental simplicity, a single electrode would ideally be used to couple to all motional modes. As a compromise, one of the four adjacent electrodes DC_left_6 and DC_right_6 can be chosen, as they can give significant coupling to all three motional modes. Nominally DC_left_5 was chosen for the tickle; however a symmetry error in the custom-made cabling means that this was actually DC_left_7, however this does not affect the measured motional frequencies. The RF signal is greatly attenuated by the low pass filter on the DC PCB;

[5]While a Lorentzian is not the correct function to use in this case, it is sufficient for most cases.

however a deliberate RF signal is still able to excite the ion's motion.

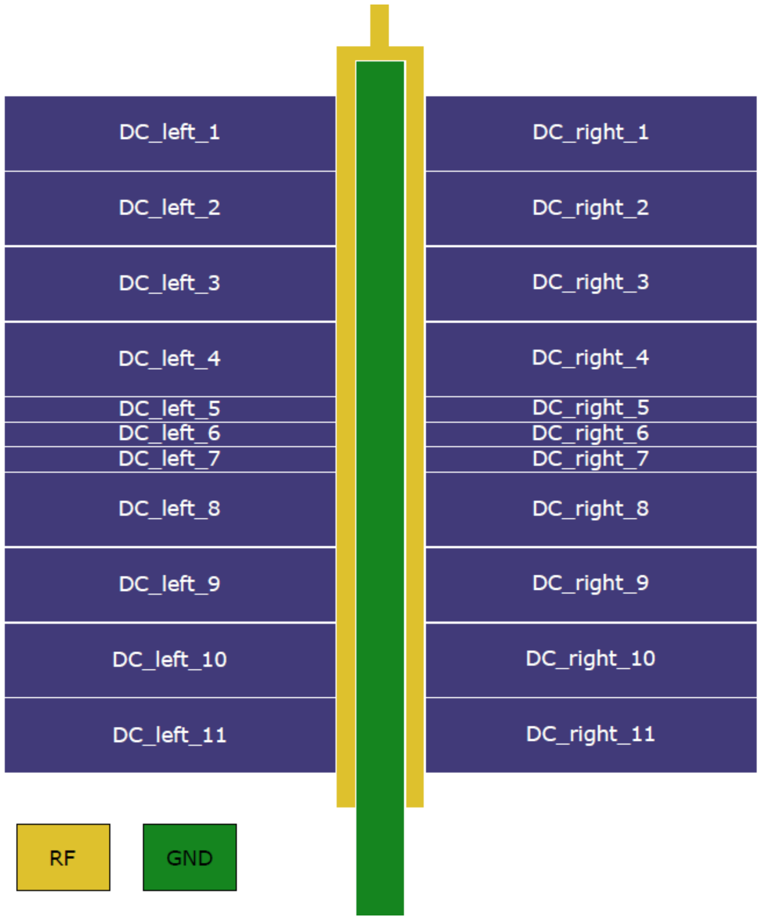

Figure 5.8: Overview of the electrodes with the used naming scheme for their pads.

To detect the three motional frequencies, RF is applied at one frequency, and the resulting fluorescence measured. Notably, these initial measurements are taken without correlating the ion's fluorescence with the phase of the RF used for trapping, as opposed to the correlation method employed in the micromotion compensation, see sect. 5.2.8.

Depending mainly on the frequency and amplitude of the applied RF signal, the fluorescence changes, and can become both higher and lower; if a fluorescence dip instead of a peak is observed, the RF amplitude is lowered to observe peaks. Typical peaks can be seen in fig. 5.9.

The motional frequencies are first measured with $^{40}Ca^{+}$. Using a classical harmonic oscillator as a model, the frequency of motion is $f = \frac{1}{2\pi} \cdot \sqrt{\frac{k}{m}}$, with the ion's mass m and a spring constant k. Therefore, the ratio of isotope frequencies is $\frac{f_{43}}{f_{40}} = \sqrt{\frac{m_{40}}{m_{43}}} = 0.964$. Thus, the motional frequencies of $^{43}\mathrm{Ca}^+$ can be found 3.6% lower than those of $^{40}\mathrm{Ca}^+$.

5.2.7 Rabi flopping

The goal of this experiment is to store information in the $4\mathrm{S}_{1/2}|4,+1\rangle$ and $4\mathrm{S}_{1/2}|3,+1\rangle$ states. To get to either, the population has to be transferred from $4\mathrm{S}_{1/2}|4,+4\rangle$ using several microwave pulses. To calibrate the frequency and duration of these pulses, rabi flops are used. Initially, the frequency is only known approximately – as the magnetic field and a.c. Zeeman shifts are not yet calibrated – and the duration is unknown. The pulse sequence for this experiment is shown in fig. 5.10.

To get the frequency, the duration is set to be be much higher than the expected t_π time[6]. Then the frequency is scanned, i.e. the sequence of state-prep, microwave pulse, shelving, and readout is repeated e.g. 100 times to receive a shelving probability at each frequency point. Once the frequency is known, the duration can be optimised. Each optimisation of one parameter increases the precision with which the other can be determined. Therefore, the frequency and duration are iteratively optimised on each transition.

The first transitions to be optimised are the σ-polarised $4\mathrm{S}_{1/2}|4,+4\rangle \rightarrow 4\mathrm{S}_{1/2}|3,+3\rangle$, and the π-polarised $4\mathrm{S}_{1/2}|4,+3\rangle \rightarrow 4\mathrm{S}_{1/2}|3,+3\rangle$, the latter using a t_π-pulse of the former as a starting point, see fig. 5.11. Due to the trap geometry detailed in sect. 3.2, the microwave field is a lot stronger in the σ-polarisation than in the π-polarisation, resulting in t_π times of $0.74\,\mu$s and $10.5\,\mu$s, respectively, with nominally the same microwave input power of $p = 21.7\,\mathrm{dBm}$.

As these two transitions diagnose the two microwave polarisations, they are

[6]If the chosen duration happens to be $2n \cdot t_{pi}$ for any integer n, then no signal would be observed. In this unlikely case, the subsequent frequency scan could be repeated with a different duration.

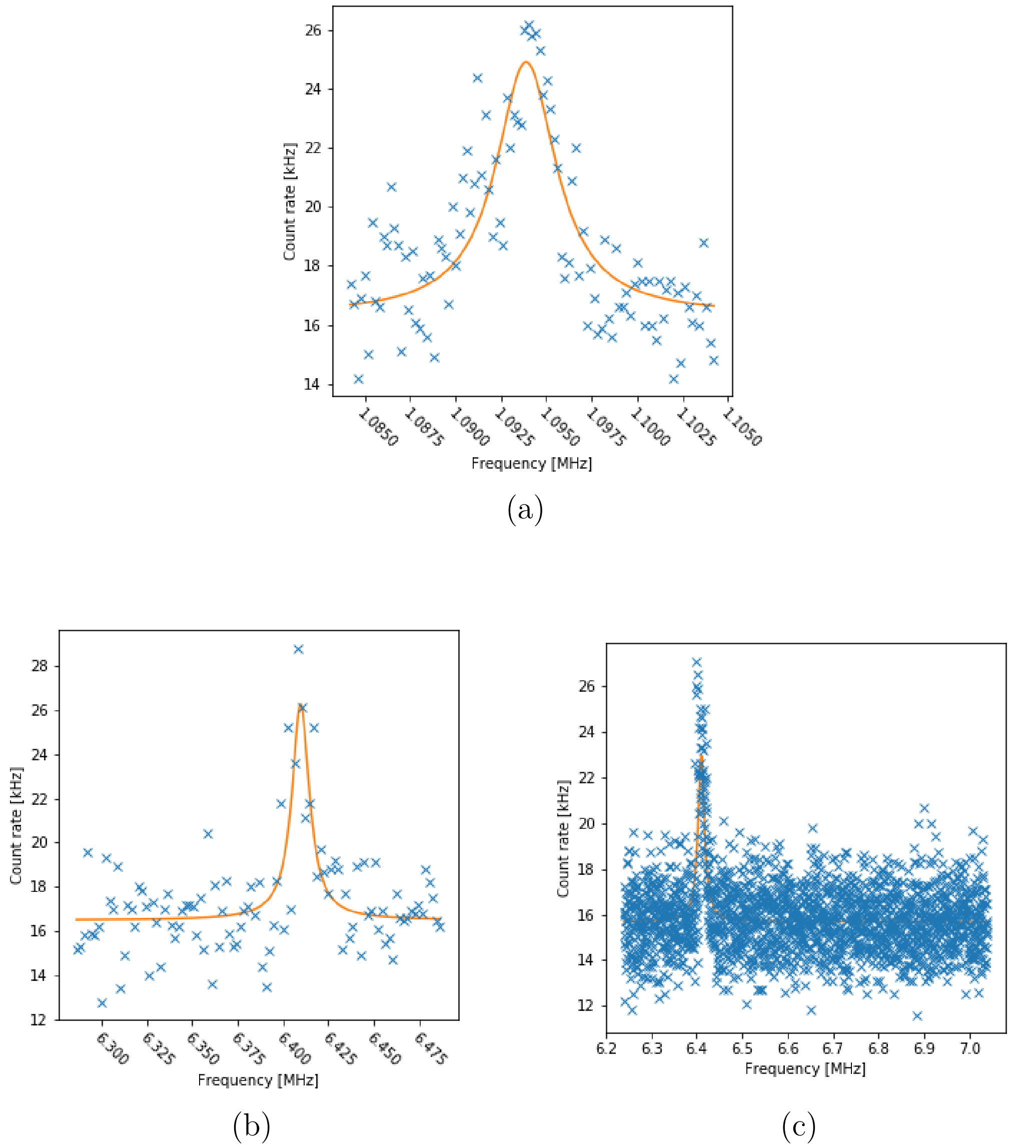

Figure 5.9: Motional frequency scans of $^{43}Ca^{+}$. The fitted frequencies are (a) 1.09 MH, axial, (b) 6.15 MHz, lower radial, and (c) 6.41 MHz, upper radial. Using the same applied signal, they exhibit a different amplitude with the axial frequency being the easiest to detect and the radial being the hardest.

Figure 5.10: The pulse sequence used to probe microwave-based transitions. Depending on the probed transition, there could be only a single microwave pulse, or multiple, as explained in the text. The ideally employed 850 nm laser was not yet operational for the results obtained in this thesis, indicated by the use of bracket, see sect. 2.7.

used for measuring the rabi flopping time as a function of the position of the ion and then compared to the simulations based on the trap geometry, see sect. 5.5.

5.2.8 Micromotion compensation

To measure the micromotion parallel to the static magnetic field, the method used in this work is to correlate the ion's fluorescence rate with the phase of the RF Paul trap. The RF frequency is $f_{RF} = 63.98$ MHz, which is much larger than a typical fluorescence rate of $f_{fluo} = 0.1$ MHz. During one RF cycle, the mean photon count n is thus

$$n = f_{fluo} \cdot f_{RF}^{-1} = 1.6 \cdot 10^{-3} \tag{5.1}$$

Therefore, to detect small changes in the fluorescence rate as a function of the RF frequency phase, many RF cycles have to be averaged; typically an observation time of $\approx 1 \dots 2$ s is used, i.e. $\approx 10^8$ RF cycles. The observation of correlated photons in the micromotion compensation verifies that the phase coherence between the RF source and the readout electronics is stable. This is

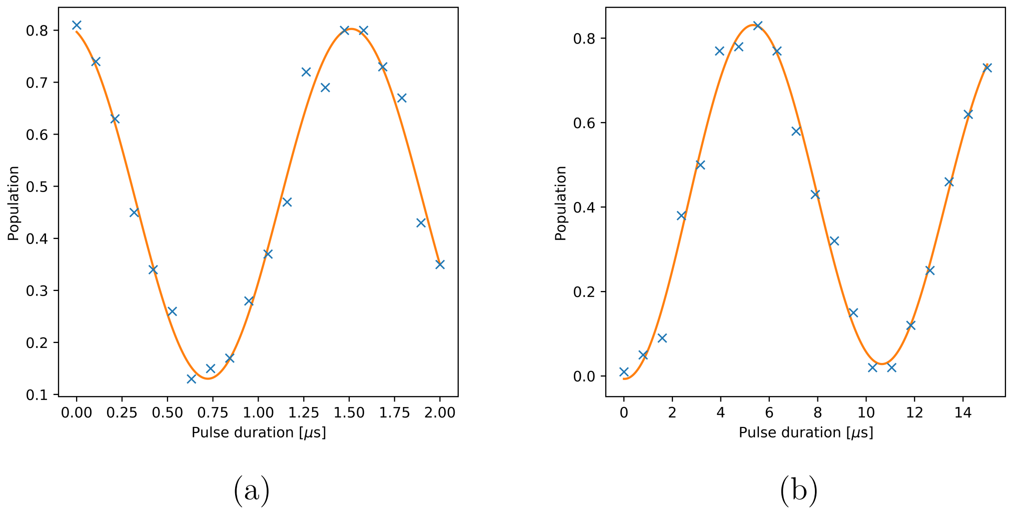

Figure 5.11: Rabi flops in $^{43}\mathrm{Ca}^+$, the shelving probability as a function of duration of the applied microwave field. The probed transitions are (a) $4\mathrm{S}_{1/2}|4,+4\rangle \rightarrow 4\mathrm{S}_{1/2}|3,+3\rangle$, which is σ-polarised, and (b) $4\mathrm{S}_{1/2}|4,+3\rangle \rightarrow 4\mathrm{S}_{1/2}|3,+3\rangle$, which is π-polarised. As explained in sect. 2.7, the maximum visible shelved population in this work is limited to 90%; the minimum population is unaffected by this.

then fitted onto the amplitude and phase of a sinusoid with the fixed frequency f_{RF}. This experiment is repeated for a number, typically $10\ldots30$, of micromotion compensation fields E_x, and the sinusoid's amplitude fitted to a function $A * \mathrm{abs}\,(E_x - E_{x,0})$, with a scaling factor A, to determine the optimal $E_{x,0}$.

In principle, the same method can be used to detect the out-of-plane micromotion. However, none of the in-plane lasers have an out-of-plane directional component, and can therefore not detect this. To compensate this out-of-plane micromotion, the ion is viewed through the imaging system's CCD camera. Images are taken at different out-of-plane micromotion compensation fields, and the apparent size of the ion compared. For higher micromotion, the ion should have a larger apparent size on the camera; however no size difference could be discerned. This could have been due to a non-optimal alignment of the lenses of the imaging system, so this could be retried after the alignment is optimised.

5.2.9 Calibration of imaging system and DC electric field

For measurements taken in sect. 5.5, it is necessary to know the conversion rate of applied compensation fields to the ion's change in position. This can be done in two ways; calculated by using the RF pseudo-potential trap strength, or by directly measuring it via the displacement of images on the CCD.

The camera, Andor iXon Ultra model DU-897U-S50-#EX, has a pixel size of 16 $\mu\mathrm{m}/\mathrm{pixel}$. The magnification factor depends on the relative position of the two lenses in the imaging system. To directly measure the effective pixel size, the imaging system is focused onto the trap, and the trap features measured in pixels and compared to the microscope measurements taken in fig. 4.3. The effective pixel size is measured as 0.16 $\mu\mathrm{m}/\mathrm{pixel}$, giving a magnification factor of 100. Next, the camera is focused on the ion and its displacement is measured as a function of applied compensation field, see fig. 5.15. The resulting scaling is measured as 0.78 $\mu m/(\mathrm{kV/m})$.

The method used for initial alignment – without an ion – of the laser beams is shown in fig. 5.12. Once an ion is trapped, its fluorescence is used to fine tune the laser's alignment.

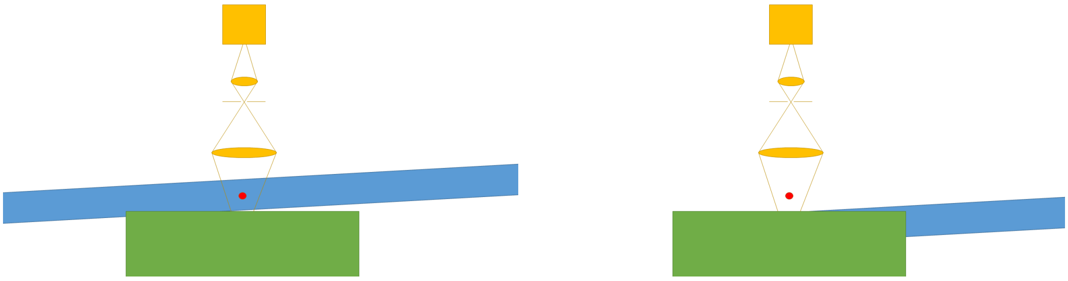

(a) (b)

Figure 5.12: Initial alignment method for the laser beams. (a) The laser beam (blue) is first intentionally angled towards the trap's surface (green). For the horizontal alignment, the beams are moved around until they are in the place where the ion is intended to be trapped (red), as seen by the imaging system (yellow). For the vertical alignment, the laser beam is then moved upwards, away from the trap, and then slowly moved downwards until it is first visible via the imaging system. The laser beam is then stopped from moving, and a flip mirror is used to divert the laser beam to a beam profiler located at the same optical distance from flip mirror as the trap. The laser beam position is on the beam profiler is noted down. (b) The flip mirror is moved out of the path again, and the laser beam is moved further downwards, until it is just no longer visible via the imaging system. The flip mirror is again used to divert the laser beam to the beam profiler. The previous location and current location on the beam profiler are averaged, creating an estimate for the alignment to hit the position below the ion. The ion's trapping height – 40 μm – is then added to this location, and the beam profiler is used to move the laser beam to this new position.

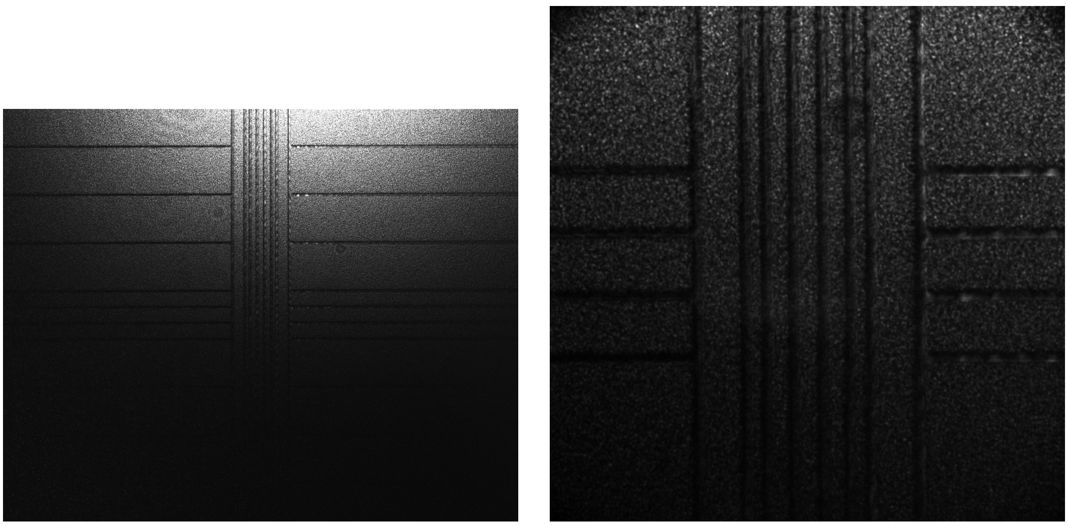

(a) (b)

Figure 5.13: Trap as imaged through the CCD camera, also see fig. 4.3. (a) Initially, a lower magnification is used to find the trap in the first place. (b) This is the normal magnification used throughout this work. Taken with an Andor iXon Ultra camera.

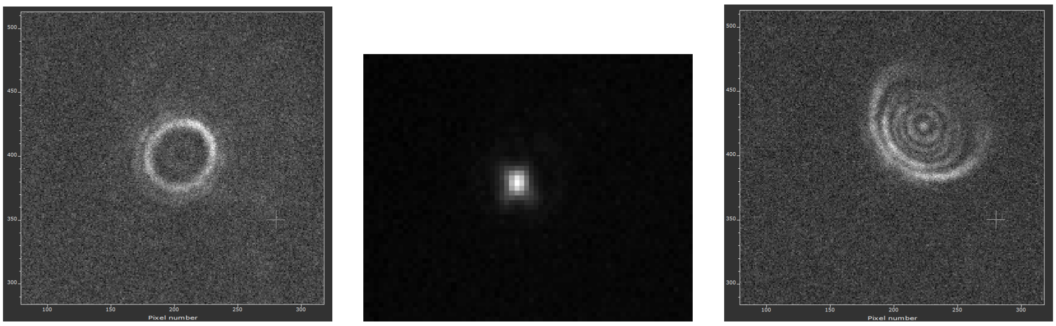

(a) (b) (c)

Figure 5.14: ^{40}Ca^{+}ion as imaged through the CCD camera, (a) below, (b) at, and (c) above focus. For the purposes of imaging, the ion is a point source of light. The point spread function above and below focus shows that the imaging system is not optimally aligned. Taken with an Andor iXon Ultra camera.

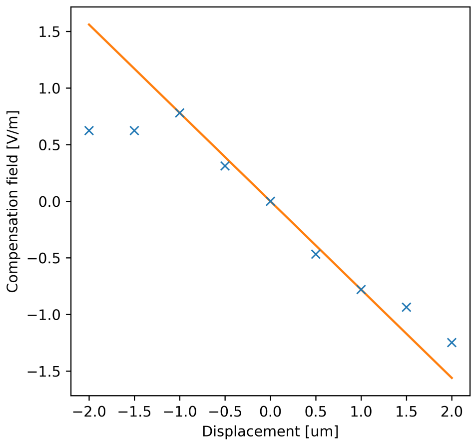

Figure 5.15: ^{40}Ca^{+} displacement as a function of applied compensation field E_x, which is added to nominal trapping field, and converted into voltages using the House model, see sect. 3.2. For each E_x, an image is taken, and its x-position fitted to a Gaussian intensity profile. The resulting centre position on the camera is converted into a displacement using the effective pixel size derived in the text. For displacements of more than $1\,\mu$m, the ion's micromotion blurs out the image too much for a clear centre to be determined, and these points are not used in the fit. The slope is measured to be 0.78 ${}^{\mu m}/_{(\mathrm{V/m})}$. Taken with an Andor iXon Ultra camera.

5.3 Loading Rate Study

5.3.1 Loading Model

The model used to derive the expected behaviour of the loading will be explained in the following sections. This model assumes ^{40}Ca, but is applicable to ^{43}Ca with only minor inaccuracies; adjusting it to a different element would only require some minor modifications, e.g. to the vapour pressure curve. The model will be used to fit the results from the loading rate study, so the dependence of the loading rate on a few key parameters – T_{oven}, V_{RF}, f_{423}, and P_{378} – is highlighted.

The loading rate is a function of the number density in the trap region n_{trap}, the interaction volume V_{inter}, the ionisation probability $P_{ionising}$, and the probability of finding a stable orbit P_{stable}. These are defined and derived in the following sections.

5.3.1.1 Number Density in Trap Region

The Ca in the oven tube is heated up to a temperature T_{oven}. This creates a vapour pressure p inside the oven tube, which is given by

$$p = p_0 \cdot 10^{A + B/T_{oven} + C * \log T_{oven}}$$

using the oven temperature T_{oven} in Kelvin, the atmospheric pressure $p_0 = 1010\,\text{mbar}$ and the unitless fitting constants $A = 10.127$, $B = -9517$, and $C = -1.403$ [2]. For a typical oven temperature of $T_{oven} = 330°\text{C}$, the pressure is $p = 2.8 \cdot 10^{-7}\,\text{mbar}$. Assuming the Ca behaves like an ideal gas, the number density n_{oven} can then be calculated as

$$n_{oven} = \frac{p}{k_B T_{oven}} = \beta P \tag{5.2}$$

using the Boltzmann constant k_B and defining $\beta = \frac{1}{k_B T_{oven}}$. At a typical oven temperature of $T_{oven} = 330°\mathrm{C}$, the number density is $n_{oven} = 3.4 \cdot 10^{15}\,\mathrm{m}^{-3}$. The oven tube has a small hole in it with an area of $A_{oven} \approx 0.2\,\mathrm{mm}^2$. This hole is assumed to be small enough not to affect the pressure inside the oven tube. The effusive flow rate Q_e of atoms out of the hole is

$$Q_\mathrm{e} = \frac{pA}{\sqrt{2\pi m k_B T_{oven}}}$$

For a typical oven temperature of $T_{oven} = 330°\mathrm{C}$, this is $Q_e = 9.5 \cdot 10^{10}\,\mathrm{s}^{-1}$. The number density of the Ca atoms inside the trapping region can then be calculated as [p. 6, 37]

$$\boxed{n_{trap}\left(T_{oven}\right) = \frac{n_{oven}\left(T_{oven}\right) A_{oven}}{4\pi l_{oven}^2} = \frac{\beta\left(T_{oven}\right) P\left(T_{oven}\right) A_{oven}}{4\pi l_{oven}^2}} \tag{5.3}$$

using the distance $l_{oven} = 48\,\mathrm{mm}$ of the oven to the nominal ion location. To a good approximation, the number density n_{trap} is constant across the trapping region. Using the typical oven temperature of $T_{oven} = 330°\mathrm{C}$, the number density in the trapping region is $n_{trap} = 2.3 \cdot 10^{10}\,\mathrm{m}^{-3}$.

5.3.1.2 Interaction Volume

The interaction volume is defined as the overlap of the trapping volume, atom beam, and photoionisation lasers.

The trapping volume – where trapping can occur if a cold ion is placed – is approximated as an ellipsoid, with diameters l_x, l_y, and l_z, with directions defined according to the trap axes from sect. 2.9. The volume is then $V_{trap} = \frac{4\pi}{3} l_x l_y l_z$. As a first approximation, the radial trap axes are assumed to be classical harmonic oscillators. The spring constant k would then be

$$k = 4\pi^2 f^2 m \tag{5.4}$$

with the radial frequency $f = f_{radial,1}$ and $f_{radial,2}$. As the radial motional frequencies are nearly the same, they are approximated as $f_{radial} = f_{radial,1} \approx f_{radial,2}$. Assuming the trap depths are the same across the two radial axes, the resulting diameter $l_{radial} = l_y = l_z$ of the trapping region is

$$l_{radial} = \frac{1}{\pi f_{radial}} \sqrt{\frac{2\psi_{E,RF}}{m}} \tag{5.5}$$

using the trap depth $\psi_{E,RF}$ from 3.2.1. Using equations 3.4 and 3.5, and assuming that the radial frequencies are the same, and the effect of the axial frequency is negligible, the radial frequency f_{radial} is linearly proportional to the RF Voltage V_{RF}. This can be expressed as

$$f_{radial}\left(V_{RF}\right) = F \cdot V_{RF} \tag{5.6}$$

using an empirically determined proportionality factor $F = 77 \pm 5\,\mathrm{kHz\,V^{-1}}$. As a simplification to the trap depth $\psi_{E,RF}$ equation 3.1, it is reduced to

$$\psi_{E,RF} = D \cdot V_{RF}^2 \tag{5.7}$$

using a proportionality factor $D = \frac{Z^2 e^2}{\pi^2 m \Omega^2} \left[\frac{b}{(a+b)^2 + (a+b)\sqrt{2ab+a^2}}\right]^2$, whose variable are explained in sect. 3.2.1, and can be used to calculate the proportionality factor as $D = 4.49\,\mu\mathrm{eV\,V^{-2}}$.

Using these expressions for the radial frequency f_{radial} and the trap depth $\psi_{E,RF}$, the radial diameter l_{radial} becomes

$$l_{radial} = \frac{\sqrt{2D}}{\pi F \sqrt{m}} = 19\,\mu\mathrm{m} \tag{5.8}$$

which is smaller than the beam diameter $2 \cdot \omega_{PI} = 32\,\mu\mathrm{m}$ used for the photoionisation process. Therefore, in the radial direction, the trap volume is limited by the extent of the RF trapping region. Notably, this diameter is not dependent on the RF voltage V_{RF}.

However, for the axial direction, the DC potentials can extend much further than this. In this thesis, six electrodes are used to generate this field, three per side; e.g. DC_left_1-3 and DC_right_1-3 from fig. 5.8. As a first approximation to the extent, the electrodes spacing is used, i.e. $l_x \sim 90\,\mu$m. This would make the trapping volume $V_{trap} \sim \frac{4\pi}{3}\frac{l_x l_y l_z}{8} = 17 \cdot 10^3\,\mu\text{m}^3$. However, l_x is larger than the typical beam diameters; therefore, the extent of the DC potential in the axial direction is not limiting the interaction volume.

The interaction volume is then approximately a cylinder with the diameter l_{radial} defining the base of the cylinder and the beam radius w_{PI} of the photoionisation lasers influencing the height of the cylinder. The photoionisation laser has a 45° angle to the trap x-axis; through geometric considerations, this gives rise to a factor of $\sqrt{2}$; the interaction region is then

$$\boxed{V_{inter} = \pi\frac{l_{radial}^2}{4} \cdot 2\sqrt{2}w_{PI}} \tag{5.9}$$

For the nominal beam diameter of $2 \cdot \omega_{PI} = 32\,\mu$m, this is $V_{inter} = 13 \cdot 10^3\,\mu\text{m}^3$. Using the typical number density $n_{trap} = 2.3 \cdot 10^{10}\,\text{m}^{-3}$ derived in the previous section, the resulting average number of atoms in the interaction volume V_{inter} is $N_{trap} = n_{trap} \cdot V_{inter} = 3 \cdot 10^{-5} \ll 1$.

5.3.1.3 Kinetic Energy Distribution

Not all of the Ca ions created in the interaction volume V_{inter} can be trapped however. To a first approximation, only atoms with a kinetic energy less than the pseudopotential trap depth can be trapped. The distribution of speeds of particles in an effusive beam which pass a given plane in some given interval of time is

$$\rho(v) = (\beta m)^2 \frac{v^3}{2}\exp\left(-\frac{1}{2}\beta m v^2\right) \tag{5.10}$$

The probability P_{slow} of an atom having a speed up to v_{max} is

$$\boxed{P_{slow} = \int_0^{v_{max}} \rho(v)\, dv = 1 - \left(\frac{1}{2} m\beta v_{max}^2 + 1\right) \exp\left(-\frac{1}{2}\beta m v_{max}^2\right) \qquad (5.11)}$$

The probability of having an atom with a kinetic energy that is smaller or equal to half the typical pseudopotential energy of $\frac{1}{2}\psi_{E,RF} = \frac{1}{2} \cdot 34\,\mathrm{meV}$, i.e. $v_{max} = \sqrt{\frac{\psi_{E,RF}}{m}} = 200\,\mathrm{ms}^{-1}$, is $P_{slow} = 11\%$. This somewhat low probability hints that a higher trap depth could improve the loading rate considerably; especially macroscopic traps with trap depths on the order of $\sim 1\,\mathrm{eV}$ would increase the loading rate.

Besides the above kinetic energy stemming from the oven temperature T_{oven}, there is another source of kinetic energy of the atoms; the recoil energy from the electron from the ionisation process. As mentioned in sect. 2.4, the minimum energy a photon can have to ionise an atom in the 4s4p $^1\mathrm{P}_1$state is $E_{min} = \frac{hc}{\lambda_{min}} = 3.1806\,\mathrm{eV}$ with the ionisation wavelength $\lambda_{min} = 389.81\,\mathrm{nm}$, the Planck constant h, and the speed of light c. However, the laser being used to ionise the atoms has a wavelength of $\lambda_{378} = 378\,\mathrm{nm}$, giving rise to a recoil energy of $E_{recoil} = hc\left(\frac{1}{\lambda_{min}} - \frac{1}{\lambda_{378}}\right) = 99.4\,\mathrm{meV}$, distributed among the electron and Ca ion. This energy E_{recoil} is bigger than the trap depth, and would make trapping impossible if the ion would receive all of it; however most of it is given to the electron. Based on conservation of energy and momentum, the energy that the Ca ion receives is

$$E_{Ca,recoil} = E_{recoil}\frac{m_{e^-}}{m_{e^-} + m} = 1.4\,\mu\mathrm{eV} \qquad (5.12)$$

using the electron mass m_{e^-} and the Ca mass m. As this is much smaller than the typical thermal energies, the recoil energy of the Ca $E_{Ca,recoil}$ can be neglected.

5.3.1.4 Ionisation Probability

The typical speed of a Ca atom arriving at the interaction region at a temperature of $T_{oven} = 330°\mathrm{C}$ is

$$v_{rms} = \sqrt{\frac{3k_BT}{m}} \approx 600\,\mathrm{ms}^{-1} \tag{5.13}$$

The typical time the Ca spends inside the interaction region is therefore $t_{rms} = \frac{l_{radial}\sqrt{2}}{v_{rms}} \sim 50\,\mathrm{ns}$. This is a lot larger than the lifetime of the 4s4p $^1\mathrm{P}_1$level $\tau = 4.49\,(4)$ ns, so, to a good approximation, this transition can be described in a steady-state regime. The average upper state population N_2 for this two-level system is [p. 94, 17]

$$\boxed{N_2 = \frac{1}{2}\frac{S}{\kappa^2 + 1 + S}} \tag{5.14}$$

using a saturation parameter S and a detuning κ in units of natural linewidth γ. More explicitly, these are $S = D_{423}\frac{2I}{\mu_0 c\hbar^2 4\gamma^2}$ and $\kappa = \frac{\omega_{12} - 2\pi f_{423}}{2\gamma}$, using the atomic dipole matrix element D_{423}, the laser intensity $I = \frac{P_{423}}{\pi w_{423}^2}$, the natural linewidth $\gamma = 35.4\,(3)$ MHz [p. 2, 37], the frequency of the transition ω_{ab}, the frequency of the laser ω_{423}, the vacuum permeability μ_0, the speed of light c, and the reduced Planck constant $\hbar$.

Typical beam intensities are several saturation intensities, leading to $S \sim 10$, therefore, to a good approximation with detuning $\kappa = 0$, half of the population is transferred from the ground state $4\mathrm{s}^2\,{}^1\mathrm{S}_0$ to the excited state 4s4p $^1\mathrm{P}_1$. At the same saturation $S \sim 10$, with a fairly large detuning of $\kappa = \frac{100\,\mathrm{MHz}}{\gamma}$, the upper state population becomes $N_2 \approx \frac{1}{4}$, which decreases the loading rate only by a factor of 2 compared to $\lim_{s\to\infty} N_2 = \frac{1}{2}$.

From there, photons from a 378 nm laser excite them out into the continuum. Typically, only a small fraction of the atoms get ionised by the 378 nm laser, i.e. this is a rate limiting step. Therefore, the beam waist of the photoionisation lasers is given by the 378 nm laser, as the 423 nm laser is typically saturating its transition, i.e. $\omega_{PI} = \omega_{378}$. The cross section for the photoionizing transition is $Q = 170\,(60)$ Mb [p. 8, 37].

The probability $P_{ionising}$ of an atom with a speed v in the interaction volume

V_{inter} to be ionised is

$$P_{ionising}\left(P_{378}\right) = a_{40} \cdot \frac{Q}{\pi w_{378}^2} \cdot \frac{\lambda_{378} P_{378}}{hc} \cdot N_2 \frac{l_{radial}\sqrt{2}}{v} \tag{5.15}$$

using the 378 nm laser power P_{378}, the 378 nm laser photon energy $\frac{hc}{\lambda_{378}}$, the isotopic abundance of ^{40}Ca in the enriched Ca $a_{40} = 75.6\%$, and the interaction time t. For a typical interaction speed of $v_{rms} = 600\,\frac{\mathrm{m}}{\mathrm{s}}$, laser power $P_{378} = 75\,\mu$W and 378 nm laser beam radius $w_{378} = 16\,\mu$m, the resulting ionisation probability is $P_{ionising} = 0.005\%$. The physical meaning of the terms influencing this probability $P_{ionising}$ are

1. The probability a_{40} of finding a ^{40}Ca isotope

2. The probability $\frac{Q}{\pi w_{378}^2}$ of one photon from the 378 nm laser hitting the cross section Q

3. The rate $R_{photons} = \frac{\lambda_{378} P_{378}}{hc}$ at which photons arrive at a cross section of the interaction volume V_{inter} perpendicular to the propagation of the laser beam

4. The time $N_2 \frac{l_{radial}\sqrt{2}}{v}$ the atom spends in both the 4s4p ^{1}P$_1$level and the interaction region

The probability P_{photon} of a single photon ionising any atom within the interaction volume V_{inter} is

$$\boxed{P_{photon} = a_{40} \cdot \frac{Q \cdot n_{trap} \cdot V_{inter}}{\pi w_{378}^2} \cdot N_2} \tag{5.16}$$

For these values, and assuming $N_2 \approx \frac{1}{2}$, the probability is $P_{photon} = 2.38 \cdot 10^{-15}$. If the photoionisation probability per atom were substantial, a simple sum over the trapping volume would not be appropriate, but as $P_{photon} \ll 1$, this is not the case. The rate at which ions of any speed are generated in the interaction volume V_{inter} is

$$R_{ionised} = P_{photon} \cdot R_{photons} \tag{5.17}$$

which is $R_{ionised} = 0.3\,\mathrm{s}^{-1}$ for typical values. The is also places an upper bound on the rate at which ions can be trapped.

Let v_{max} be the maximum speed at which a resulting ion is slow enough to still have a chance to be trapped. Then

$$R_{slow} = \int_0^{v_{max}} P_{photon} \cdot R_{photons} dv = a_{40} \cdot \frac{Q \cdot n_{trap} \cdot V_{inter}}{\pi w_{378}^2} \cdot N_2 \cdot \frac{\lambda_{378} P_{378}}{hc} \cdot \int_0^{v_{max}} \rho\left(v\right) dv \tag{5.18}$$

For $v_{max} \to \infty$, this reduces back to the previous eq. 5.17, as the probability distribution is normalised, $\int_0^\infty \rho\left(v\right) dv = 1$. Using eq. 5.11, the resulting rate R_{slow} is then

$$R_{slow} = R_{photons} \cdot P_{photon} \cdot P_{slow} = R_{photons} \cdot a_{40} \cdot \frac{Q \cdot n_{trap} \cdot V_{inter}}{\pi w_{378}^2} \cdot N_2 \cdot P_{slow} \tag{5.19}$$

5.3.1.5 Stability Factor

In the pseudopotential approximation, the maximum kinetic energy that an ion can have and still stay trapped is given by the trap depth $\psi_{E,RF}$. However, this is only true for atoms ionised in the centre of the trap. The further away from this centre, the less kinetic energy is necessary to escape the trap. This can be described by the following inequality

$$E_{kin} + E_{pot} \leq \psi_{E,RF} \tag{5.20}$$

with $E_{pot} = \frac{1}{2} k l^2$ being the pseudopotential energy in the radial plane. The timescales of laser cooling are $\sim 1\,\mathrm{ms}$, while the timescales of an ion escaping the trap are given by the RF period, $\sim \frac{1}{64\,\mathrm{MHz}} = 16\,\mathrm{ns}$. Therefore, laser cooling is required for an ion to stay trapped, but it does not help to trap an ion to begin with [p. 23, 30]. As a first approximation, to account for the effect of a random starting location and its associated potential energy E_{pot}, the maximum kinetic energy $E_{kin,max} = \frac{1}{2}mv_{max}^2$ is limited to half the trap depth $\psi_{E,RF}$, i.e.

$$\boxed{v_{max} = \sqrt{\frac{\psi_{E,RF}}{m}} = V_{RF}\sqrt{\frac{D}{m}}} \tag{5.21}$$

So far, the trap has been assumed to be stable for any particle with low enough kinetic energy. To account for the effect of a potentially unstable orbit, a trap stability probability P_{stable} is introduced. It gives the probability of a new ion to reach a stable orbit, and stay confined within the trap volume. This is a function of the phase of the RF, ion location, speed, and direction. It is typically estimated using numerical simulations, e.g. [p. 23ff, 30]. Judging from the cited numerical simulations, as a first, very rough, approximation, it can be thought of as depending only on the trap stability parameter q

$$P_{stable} = P_{stable,0}\left(1 - \frac{q}{q_{unstable}}\right) \tag{5.22}$$

using the the trap stability parameter where trapping is no longer possible, $q_{unstable} = 0.91$. The factor of P_0 comes from the observation that, even in a stable trap, the maximum fraction of trapped ions is limited [p. 25, 30], and as a very rough initial estimate is $P_{stable,0} = \frac{1}{2}$. This uses the trap stability parameter q as defined in sect. 3.2.1. As a simplification, it is reduced to

$$q = GV_{RF} \tag{5.23}$$

using a proportionality constant $G = -\frac{16Ze}{\pi m\Omega^2}\frac{b}{(a+b)^2\sqrt{2ab+a^2}} = 3.20 \cdot 10^{-3}\,\mathrm{V}^{-1}$. Therefore, the previous equation becomes

$$\boxed{P_{stable}\left(V_{RF}\right) = P_{stable,0}\left(1 - \frac{GV_{RF}}{q_{unstable}}\right)} \tag{5.24}$$

which for a typical values becomes $P_{stable} = 0.34$.

5.3.1.6 Loading Rate

Using the rate of ion creation R_{slow} with a velocity of up to v_{max} from eq. 5.19, and inserting the maximum velocity $v_{max} = V_{RF}\sqrt{\frac{D}{m}}$, eq. 5.21, and multiplying by the probability of finding a stable orbit P_{stable}, eq. 5.24, the resulting rate of trapped ions is $R_{ionised} = P_{stable} \cdot R_{slow}$. Expanding this yields

$$\boxed{R_{ionised} = P_{stable}\left(V_{RF}\right) \cdot R_{photons}\left(P_{378}\right) \cdot P_{photon}\left(T_{oven}, f_{423}\right) \cdot P_{slow}\left(V_{RF}, T_{oven}\right)} \tag{5.25}$$

For typical values listed below, this is $R_{trapped} \sim 0.013\,\mathrm{s}^{-1}$, i.e. the mean time to load is $T_{mean} = \frac{1}{R_{trapped}} = 80\,\mathrm{s}$. This can be compared to the experimental result shown in fig. 5.17a, where at nominally the same parameters the observed mean time to load was $t_{mean} = 117\,\mathrm{s}$. Considering the number of simplifying assumptions which had to be introduced in the model, it should only be expected to give the correct order of magnitude for the overall absolute rate, so the observed agreement to within a factor 1.5, for this parameter set, is encouraging. Possible sources of errors include misalignment of the lasers, uncertainty in the oven temperature T_{oven}, and the approximations made for to get the trap stability probability P_{stable} and the maximum speed v_{max}.

Without doing numerical simulations, an improvement to this model would be to add a position dependence to the maximum velocity v_{max}, so that eq. 5.20 is true at every point. Additionally, numerical simulations could give a more accurate result for the trap stability probability P_{stable}, as well as replacing the pseudopotential approximation. In this model, this could be implemented by nominally setting the maximum velocity $v_{max} \to \infty$, and making the trap stability probability P_{stable} a function of initial position and ion velocity v, as well as on the phase of the RF.

For ease of reference, the variables and terms are listed in tables 5.1 and 5.2.

For the loading rate study, the following proportionalities are used for fitting:

Variable	Value
Q	170 (60) Mb
$q_{unstable}$	0.91
A	10.127
B	−9517
C	−1.403
p_0	1010 mbar
Z	1
m	40 amu
γ	35.4 (3) MHz
D_{423}	
f_0	709.07812 THz

(a)

Variable	Value
a_{40}	75.6%
A_{oven}	$0.2\,\mathrm{mm}^2$
l_{oven}	48 mm
$P_{423,0}$	493 nW
$w_{378} = w_{423}$	16 μm
a	60 μm
b	23.6 μm
Ω	$2\pi \cdot 63.98$ MHz

(b)

Variable	Value
T_{oven}	330°C
V_{RF}	90 V
P_{378}	75 μW
P_{423}	$10 \cdot P_{423,0}$
f_{423}	$f_0 + 250$ MHz

(c)

Table 5.1: Overview of the variables, and their nominal values, used in the loading rate model. (a) This list shows the physical constants which cannot be changed when trapping $^{40}Ca^+$. (b) These variable could be changed, but would require some work to do so. (c) These variables are straightforward to adjust in the laboratory, and most of them are scanned in the following loading rate study

Variable	Expression	Typical value
N_2	$\frac{1}{2}\frac{S}{\kappa^2+1+S}$	$\frac{1}{2}$
S	$S = D_{423}\frac{2I}{\mu_0 c\hbar^2 4\gamma^2}$	10
κ	$\frac{\omega_{12}-2\pi f_{423}}{2\gamma}$	5
I	$\frac{P_{423}}{\pi w_{423}^2}$	$90\,\mathrm{kW\,m^{-2}}$
β	$\frac{1}{k_B T_{oven}}$	$1.2 \cdot 10^{20}\,\mathrm{J^{-1}}$
p	$p_0 \cdot 10^{A+B/T_{oven}+C*\log T_{oven}}$	$2.8 \cdot 10^{-7}$ mbar
n_{trap}	$\frac{\beta(T_{oven})P(T_{oven})A_{oven}}{4\pi l_{oven}^2}$	$2.3 \cdot 10^{10}\,\mathrm{m^{-3}}$
G	$-\frac{16Ze}{\pi m\Omega^2}\frac{b}{(a+b)^2\sqrt{2ab+a^2}}$	$3.20 \cdot 10^{-3}\,\mathrm{V^{-1}}$
D	$\frac{Z^2e^2}{\pi^2 m\Omega^2}\left[\frac{b}{(a+b)^2+(a+b)\sqrt{2ab+a^2}}\right]^2$	$4.49\,\mu\mathrm{eV\,V^{-2}}$
l_{radial}	$\frac{\sqrt{2D}}{\pi F\sqrt{m}}$	19 μm
V_{inter}	$\pi\frac{l_{radial}^2}{4} \cdot 2\sqrt{2}w_{PI}$	$13 \cdot 10^3\,\mu\mathrm{m}^3$
P_{stable}	$P_{stable,0}\left(1-\frac{GV_{RF}}{q_{unstable}}\right)$	0.34
$R_{photons}$	$\frac{\lambda_{378}P_{378}}{hc}$	$1.4 \cdot 10^{14}\,\mathrm{s^{-1}}$
P_{photon}	$a_{40} \cdot \frac{Q \cdot n_{trap}(T_{oven}) \cdot V_{inter}}{\pi w_{378}^2} \cdot N_2\left(f_{423}\right)$	$2.38 \cdot 10^{-15}$
P_{slow}	$1-\left(\frac{1}{2}m\beta v_{max}^2+1\right)\exp\left(-\frac{1}{2}\beta m v_{max}^2\right)$	11%

Table 5.2: Overview of the derived variables used in the loading rate model, along with their defining equations, and values when using the typical variable values from tab. 5.1.

$$R_{trap}\left(P_{378}\right) \propto P_{378} \tag{5.26}$$

$$R_{trap}\left(f_{423}\right) \propto \frac{S}{\left(\frac{\omega_{12}-2\pi f_{423}}{\gamma}\right)+1+S} \tag{5.27}$$

$$R_{trap}\left(V_{RF}\right) \propto \left(1-\frac{GV_{RF}}{q_{unstable}}\right)\cdot\left[1-\left(\frac{1}{2}\beta DV_{RF}+1\right)\exp\left(-\frac{1}{2}\beta DV_{RF}\right)\right] \tag{5.28}$$

$$R_{trap}\left(T_{oven}\right) \propto \frac{\beta\left(T_{oven}\right)p\left(T_{oven}\right)A_{oven}}{4\pi l_{oven}^{2}}\cdot\left[1-\left(\frac{1}{2}\beta DV_{RF}+1\right)\exp\left(-\frac{1}{2}\beta DV_{RF}\right)\right] \tag{5.29}$$

For the dependence on the oven temperature T_{oven}, a good approximation is $R_{trap}\left(T_{oven}\right) \propto p\left(T_{oven}\right)$, as the vapour pressure dependence dominates over the kinetic energy dependence for typical values.

5.3.2 Experimental Results

On several occasions during the initial studies, no ions could be loaded, and at other times the loading rate seemed to fluctuate a lot. It is not trivial to load an ion trap situated so close to the electrodes. To investigate this, a systematic study of the effect of various parameters on the loading rate is undertaken.

The composition of the Ca inside the oven is 24.3% ^{43}Ca, and 75.6% ^{40}Ca, see sect. 3.3.4.2. Therefore, the mean time to load $\mathrm{t_{mean}}$ for ^{40}Ca should be $\frac{75.6\%}{24.3\%} = 3.1$ times that of ^{43}Ca, all else being equal. To get better statistics, this loading rate study is undertaken with ^{40}Ca; this should closely match that of ^{43}Ca up to this factor.

Starting at room temperature, the atomic oven heats up to 330 C in $t_{warmup} \approx$ 10 s, which directly affects the time it takes from starting a loading experiment to having loaded an ion.

This loading rate study does not intend to find absolute values for the loading rate as a function of the parameters, only the relative effect of changing a particular parameter. The model used for describing the Paul trap is given in sect. 3.2.1. This model only describes the effect of the RF component of the trap; the DC component is neglected. This is due to the additional complexity of having an angle θ between the ion's radial trapping as defined by the RF and DC components. However, as a first approximation, the effect of the DC component on the radial trapping can be neglected if the trapping strength due to the RF component is much larger than that due to the DC component. This is the case if the axial secular frequency is much smaller than the radial frequencies, i.e. $f_{axial} \ll f_{radial,1}, f_{radial,2}$. In this experiment, this is typically the case, as $f_{axial} \approx 1\,\text{MHz}$ and $f_{radial} \approx 6\,\text{MHz}$, as measured in sect. 5.2.6.

The investigated parameters are the axial secular frequency f_{axial}, RF trapping voltage V_{RF}, angle θ of the ion's principal trapping[7] axis y' to the ion trap's axis y, 423 nm laser frequency offset f_{423}, 378 nm laser power P_{378}, and temperature of the atomic oven T_{oven}, see fig. 5.16. Also, the long term stability of the loading rate over several days is tested, see fig. 5.17.

f_{axial} The axial secular frequency f_{axial}, together with the RF trapping voltage V_{RF}, changes the trap stability parameters, and thus affects whether an ion will find a stable orbit. The loading rate model introduced in the previous section assumes the axial frequency is low enough not to affect the radial directions, but high enough to still confine the ion. The expected behaviour is plateau-like, with no change in loading rate for a fairly wide range, and a sharp drop in loading rate outside of this plateau. As the mean time to load was initially quite large, $t_{mean} > 5$ min, a wide scan of axial frequencies $f_{axial} = 0.5 \ldots 3$ MHz is done, with com-

[7] An angle of $\theta = 0°$ implies y' to coincide with y

paratively small statistics to quickly decrease the mean time to load to a more to a value that does not impede experimental progress. This initial wide sweep can be seen in fig. 5.16a, plotted in green[8]. While this wide sweep only features small statistics for its extreme values, it appears a factor of ≈ 10 can be gained in the loading rate by using an axial frequency of 1 MHz over the initially used 3 MHz. Two months later[9], a more narrow scan around $f_{axial} = 1.0\,\mathrm{MHz}$ was undertaken; it is plotted in blue, with $n = 6$ loading events per point. A sharp minimum appeared to be at an axial frequency of $f_{axial} = 1.1\,\mathrm{MHz}$, but another narrower scan with $n = 5$ loading events per point, plotted in orange, indicates that this was just statistical. Overall, a reasonable conjecture is that the mean time to load t_{mean} is unaffected by the choice of axial frequency f_{axial} as long as it is within roughly the range of $f_{axial} = 0.75\,\mathrm{MHz} \ldots 1.5\,\mathrm{MHz}$. For the rest of the loading rate study, an axial frequency of $f_{axial} = 1.0\,\mathrm{MHz}$ is used. Due to the small statistics involved, it is not conclusive that the axial frequency f_{axial} has a significant impact on the mean time to load t_{mean}. Thus the data is consistent with the expected plateau-like behaviour as a function of the axial secular frequency.

V_{RF} The RF trapping voltage V_{RF}, has two effects in the loading rate model; it increases the trap stability parameter q, leading to a lower probability P_{stable} of reaching a stable orbit, which increases the mean time to load t_{mean}. However, the RF trapping voltage V_{RF} also increases the trap depth $\psi_{E,RF}$, leading to a higher maximum velocity v_{max} that can still be trapped; this decreases the mean time to load

[8]For axial frequencies of $f_{axial} = 0.5\,\mathrm{MHz}, 1.0\,\mathrm{MHz}, 2.0\,\mathrm{MHz}, 3.0\,\mathrm{MHz}$, the resulting mean times to load are $t_{mean} = 200\,\mathrm{s}, 20\,\mathrm{s}, 60\,\mathrm{s}, 400\,\mathrm{s}$, with observed loading events numbering $n = 2, 14, 6, 1$, respectively. For the single loading event, the plotted standard deviation is half its mean time to load.

[9]With slightly different loading parameters, i.e. the datasets cannot be compared directly, but rather be viewed on their own.

t_{mean}. These two effects are competing, so it could be expected that the RF trapping voltage V_{RF} has a global minimum. Due to saturation effects in the chosen RF amplifier, see sect. 4.3.4, the synthesizer output power P_{synth} only roughly translates to a linear increase in RF trapping voltage V_{RF}. Over a wide range of RF trapping voltages[10] $V_{RF} = 80 \ldots 90\,\mathrm{V}$ ($q = 0.24 \ldots 0.27$, $\psi_{E,RF} = 27 \ldots 34\,\mathrm{meV}$, and $P_{synth} = -6.65 \ldots -4.65\,\mathrm{dBm}$), the mean time to load t_{mean} does not change significantly; at the more extreme values of $V_{RF} = 90 \ldots 110\,\mathrm{V}$ ($q = 0.27 \ldots 0.33$, $\psi_{E,RF} = 34 \ldots 51\,\mathrm{meV}$, and $P_{synth} = -4.65 \ldots -2.65\,\mathrm{dBm}$), t_{mean} increases significantly. An RF trapping voltage of $V_{RF} = 90 \pm 10\,\mathrm{V}$ ($q = 0.27 \pm 0.03$, $\psi_{E,RF} = 34 \pm 8\,\mathrm{meV}$) corresponding to a synthesizer output power of $P_{synth} = -4.65\,\mathrm{dBm}$, is chosen for the rest of the loading rate study. The mean time to load t_{mean} exhibits a minimum likely to be in the range $V_{RF} = [70\,\mathrm{V} \ldots 90\,\mathrm{V}]$, however more measurements at value outside the measured range are needed to confirm. The fit to the model does not work well for the RF voltage V_{RF}. The increase in the mean time to load t_{mean} from $V_{RF} = 80\,\mathrm{V}$ to $110\,\mathrm{V}$ suggests a stronger dependence on the trap stability parameter q than assumed in the model, and the drop in the mean time to load t_{mean} from $V_{RF} = 60\,\mathrm{V}$ to $80\,\mathrm{V}$ might indicate that the approximations made for the maximum velocity v_{max} only coarsely predict the data. Both the above mentioned terms affecting the mean time to load t_{mean} are only modelled very roughly; numerical simulations would likely be needed for more accurate results.

θ The angle θ of the ion's principal trapping axis y' to the ion trap's axis y is determined by the applied DC voltages. The loading rate model assumes an angle of $\theta = 0$; including this angle in the model would

[10] RF Trapping voltages are given in voltage amplitudes.

increase its complexity further, and is not attempted in this thesis. It is expected that the angle θ does not affect the loading rate, as long as it is non-zero to allow for cooling of all axes of motion. Fig. 5.16d shows the mean time to load t_{mean} as a function of this angle $\theta = 0 \ldots 40°$; angles $\theta > 40°$ require DC voltages outside the range that the used DAC is able to supply, $[-10\,\mathrm{V}, 10\,\mathrm{V}]$. Angles in the range of $\theta = 0 \ldots 25°$ show no statistically significant variation in the mean time to load t_{mean}, however for the higher angles of $\theta = 32°$ & $\theta = 40°$, no loading events are registered within 30 minutes each. To avoid scaling fig. 5.16d too much, a mean time to load of $t_{mean} = 1000\,\mathrm{s}$ is used to indicate this potential infinity. Negative angles are assumed to affect the mean time to load in the same way as their absolute values. As no dependence on the angle θ within the acceptable range of $\theta = 0 \ldots 25°$ is found, the previous default angle of $\theta = 15°$ is left unchanged for future experiments. It is was not investigated why angles θ of 32° and above did not succeed. For an angle of $\theta = 0°$, it was somewhat surprising that trapping was successful. However, this might be due to the rapid measurement method of immediately releasing an ion once it has been confirmed to have been trapped; any slow build-up of motion in the uncooled axis might only result in a lost ion after some time. Additionally, there might be a small coupling between the two radial motional modes or small residual angle.

f_{423} The 423 nm laser provides the first step of the ionisation process, from $4s^{2\,1}S_0$ to $4s4p\,^1P_1$, see sect. 2.4. For this laser, a high enough power P_{423} is used such that all Ca atoms of the chosen isotope are excited, to a good approximation. However this is assuming that the laser's frequency f_{423} is close to resonant to the isotope-dependent transition frequency. As explained in sect. 5.2.1, using the existing ionisation setup, the frequency needs to be measured. Fig. 5.16e shows the mean time to load t_{mean} as a function of detuning from a frequency of 709.07812 THz. As the adopted laser power P_{423} is high – about ten time the saturation power $P_{423,0}$, the expected size of the feature is > 100 MHz, with a saturated global minimum in the mean time to load t_{mean} and a slow drop-off away from the minimum. Between $f_{423} =$ 250 MHz and 350 MHz, no significant difference is found. The loading rate model is used to fit to the data, which it predicts within error bars. The fitted optimal frequency is 300 MHz. All future trapping attempts for ^{40}Ca are made at a nominal $f_{423,\mathrm{Ca}^{40}} = 250$ MHz, as this is the detuning previous data has been taken at. However, as will be explained in sect. 5.4, the wavemeter is expected to drift by ± 100 MHz. Assuming it drifts from $f_{423} = 250$ MHz to $f_{423} = 150$ MHz, based on the fit, this wavemeter drift could cause a factor of about 2 in expected difference in the mean time to load, and would be a contributing factor to fluctuations in the mean time to load t_{mean}. More measurements with a recently calibrated wavemeter would be needed to confirm this more precisely. The isotope shift between ^{40}Ca and ^{43}Ca is $f_\Delta =$ 612 MHz, this is added to the ^{40}Ca frequency offset to give $f_{423,^{43}\mathrm{Ca}} =$ 862 MHz.

P_{378} The 378 nm laser provides the second step of the two stage ionisation process, from $4\mathrm{s}^{2\,1}\mathrm{S}_0$ to $4\mathrm{s}4\mathrm{p}^1\mathrm{P}_1$ and then to the continuum, see sect. 2.4. Its frequency is therefore unimportant, as long as it is high enough to excite the electron out to the continuum. Assuming there is enough 423 nm laser intensity at the right wavelength, the resulting cloud of Ca is mostly in the exited state $4\mathrm{s}4\mathrm{p}^1\mathrm{P}_1$. Therefore, an increase in the 378 nm intensity I_{378} should make a linear difference in the mean time to load t_{mean}. Only two data points are taken here, the second one at

2.5 times the initial 378 nm intensity, see fig. 5.17a and 5.17b. However, these are taken with good statistics, 538 and 75 loading events, respectively. The loading rate model predicts an inversely linear relationship between the mean time to load t_{mean} and the 378 nm laser power, eq. 5.26. Their fitted load times are $t_{1/2}(I_{378}) = 88 \pm 5\,\mathrm{s}$ and $t_{1/2}(2.5 \cdot I_{378}) = 35 \pm 5\,\mathrm{s}$, which follow the expected linear relationship to within their error bars.

T_{oven} The loading rate model predicts that the temperature of the atomic oven T_{oven} changes the loading rate as per eq. 5.29. This has two main contributions; firstly, the vapour pressure $p(T_{oven})$ increases the number density, decreasing the mean time to load t_{mean}. Secondly, the kinetic energy E_{kin} of the atoms increases, which decreases the fraction of atoms P_{slow} slow enough to be captured. Fig. 5.17a-c show the loading statistics for oven temperatures of $T = 330°\mathrm{C}$, 340° C, and 350° C, with respective fitted times of $t_{1/2}$=35±4 s, 29±3 s, and 11±1 s. The fit to the loading rate model (orange) is reasonably good; larger statistics could lead to the results more closely following the model. Notably, the fit (green) using just the dependence on the vapour pressure $p(T_{oven})$ already predicts the mean times to load t_{mean} almost as well as the full model, and is already a good approximation. The temperature of the atomic oven T increases the vapour pressure $p(T_{oven})$ of Ca by a factor of $\approx 1.05\,\mathrm{K}^{-1}$ around the used temperatures of $T = 320 \ldots 350°\mathrm{C}$. As a rule of thumb, this means that the mean time to load should follow $t_{mean}(T) \approx 1.7 \cdot t_{mean}(T + 10°\mathrm{C})$.

In conclusion, all of the investigated parameters change the loading rate. The loading rate model predicts the dependence on the 423 nm laser frequency f_{423}, the 378 nm laser power P_{378}, and the oven temperature T_{oven} fairly well. However, the dependence on the RF voltage V_{RF} indicates shortcomings in the model, that need improved expressions for the probability of finding a stable orbit P_{orbit} or numerical simulations to more accurately predict the mean time to load t_{mean}. The approximations made for the loading rate model neglect any dependence on the axial frequency f_{axial} and the angle θ. However, values that allow for reliable loading are experimentally determined.

For subsequent experiments, the temperature of the atomic oven T_{oven} has been increased to $T = 340°\mathrm{C}$, limited by concerns over too much accumulated deposition of calcium onto the trap. There are also some concerns over increasing the 378 nm laser intensity I_{378}, as it might charge some parts of the trap, although as investigated in sect. 5.4, these concerns seem to be unfounded and the 378 nm laser intensity I_{378} is increased by the factor of 2.5, limited by concerns over laser induced fibre damage. These improvements mean that typical loading times are now $10 \ldots 20\,\mathrm{s}$ for ^{40}Ca and about three times that for ^{43}Ca. Following these changes, the loading rate is no longer an obstacle to experimental progress.

Over the course of 24 hours, 12 ions are lost, therefore the average observed lifetime of an ion in this trap is 2.0 ± 0.5 hours. This means the vacuum pressure is sufficiently low for single ion experiments, even at room temperature.

The main goal of this loading rate study was to find values for the parameters to allow for both reliable, and reasonably fast, loading. These values have been found, and loading no longer hinders experimental progress in other areas.

Effects that were not looked at in this study include the magnetic field B, Doppler cooling parameters, 423 nm laser intensity, and 423 nm & 378 nm laser beam pointing, polarisations and focus. For example, for the chosen qubit, the magnetic field $B = 28.8\,\mathrm{mT}$ has a significant effect on how effectively 423 nm laser polarisations couple to neutral Ca, see fig. 2.1.

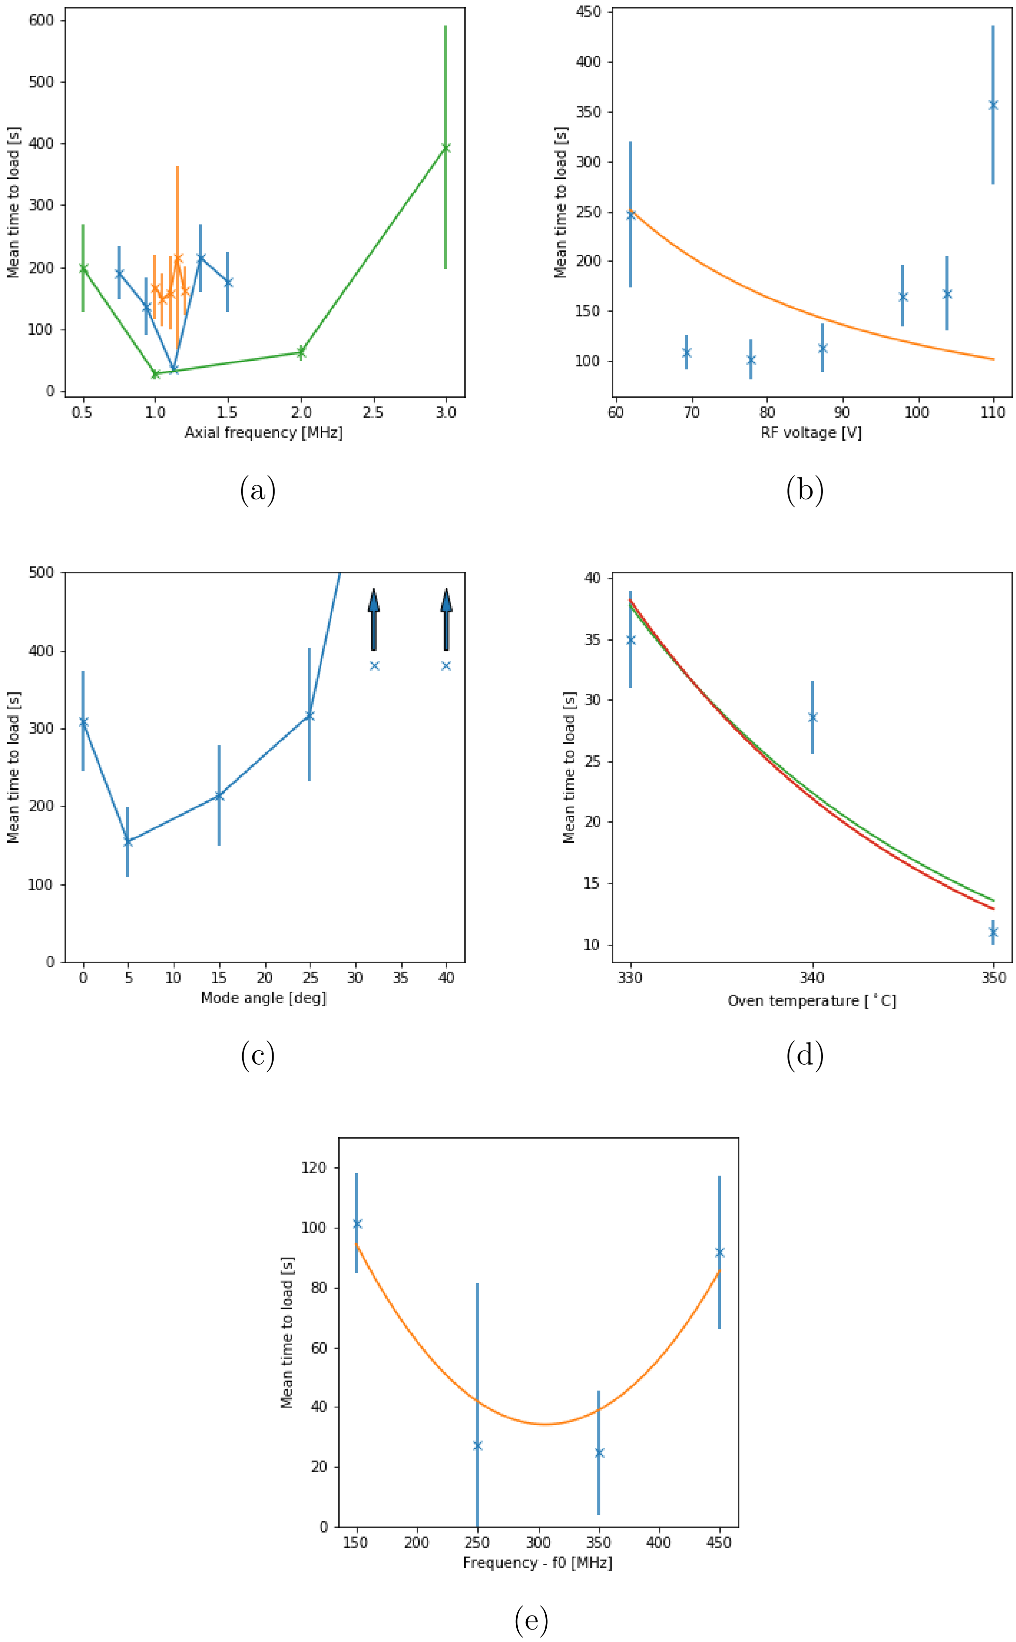

Figure 5.16: Mean time to load T_{mean} as a function of various parameters. Error bars are standard deviations. The details are explained in sect. 5.3. (a) Axial secular frequency, (b) RF trapping voltage, fitted to eq. 5.28, (c) Trap potential angle, (d) Temperature of the atomic oven fitted to eq. 5.29 (orange). The other fit (green) is neglecting all terms within eq. 5.29 except for the vapour pressure $p\left(T_{oven}\right)$ term; as can be seen, this provides a good approximation. (e) The 423 nm laser's detuning from the nominal reference f_0 fitted to eq. 5.27.

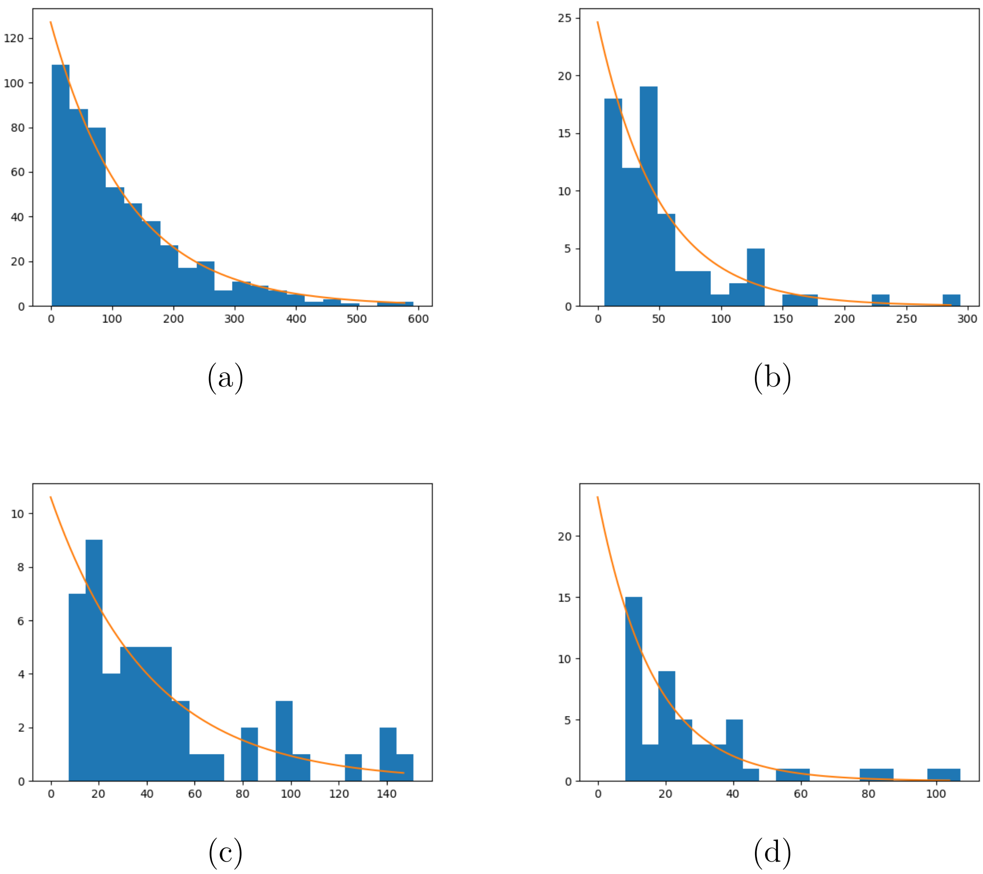

(a) (b) (c) (d)

Figure 5.17: Long term stability of the loading rate. (a) Shown here are 538 loading events over the course of 24 hours at a nominal oven temperature of $330\,\mathrm{C}^\circ$. The mean time to load $\mathrm{T}_{\mathrm{mean}}$ is 117 s. Fitting an exponential decay to this, $N_0 e^{-t/\tau}$, the $1/e$ decay constant τ is measured to be 127 s, i.e. the half time is $\tau_{1/2} = 88$ s. If there were any problems with a parameter drifting over time, especially in a bi-modal way, it would show up in this graph as a deviation from the fit. However, no statistically significant deviation is measured, therefore $\mathrm{T}_{\mathrm{mean}}$ is stable. (b) Shown here are 75 loading events with a half time of $\tau_{1/2} = 35 \pm 6\,\mathrm{s}$, again at a nominal oven temperature of $330\,\mathrm{C}^\circ$, but with 2.5 times the 378 nm laser intensity used in (a). (b) Shown here are 50 loading events with a half time of $\tau_{1/2} = 29 \pm 5\,\mathrm{s}$, at an increased nominal oven temperature of $340\,\mathrm{C}^\circ$, with 2.5 times the 378 nm laser intensity used in (a). (d) Shown here are 50 loading events with a half time of $\tau_{1/2} = 11 \pm 3\,\mathrm{s}$, at an increased nominal oven temperature of $350\,\mathrm{C}^\circ$, with 2.5 times the 378 nm laser intensity used in (a). This shows the loading rate's linear dependence on the 378 nm laser intensity, and the exponential dependence on the oven temperature.

5.4 Drifts in trap parameters over time

Interleaved with the loading rate study of sect. 5.3, other important parameters in the experiment are measured to ensure that any potential change in the loading rate can be traced back to a specific cause. These logged parameters are the temperature of the pillbox T_{PB}, the optimal voltage on the mirror of the Doppler beams in two axes, V_x and V_y, the wavemeter readings of the lasers involved in Doppler cooling, λ_{397}, λ_{866} and λ_{854}, the ion's three motional frequencies, f_{ax}, f_{r1} and f_{r2}, and lastly the ion's optimal micromotion compensation field [11] E_x. These parameters are plotted in figs. 5.18 and 5.19.

T_{PB} As mentioned in sect. 3.8, the PCB designed to measure the pillbox's temperature is not installed. Therefore, the cryostat's temperature diode is used instead, see fig. 5.18a. However, its thermal connection to the pillbox is not well defined, so this should be taken more as a qualitative indicator of temperature change. While all of these measurements are taken at room temperature, there is a significant increase in temperature while the Ca oven is at its temperature of $\approx 330°\mathrm{C}$. As there is a vacuum gap between the oven and the pillbox, and thermal conductivity from the oven to the baseflange and then to the pillbox is negligible, this increase in temperature is due to black body radiation. This increase could be more significant at lower temperatures as the pillbox, made mainly out of OFHC, has a much lower specific heat capacity at lower temperatures, e.g. $c_{16\,\mathrm{K}} = 3.8\,{}^{J}/_{\mathrm{kg\cdot K}}$, compared to room temperature, $c_{300\,\mathrm{K}} = 377\,{}^{J}/_{\mathrm{kg\cdot K}}$. Therefore, the same amount of thermal radiation changes the temperature of the pillbox 100 times more at cryogenic temperatures compared to room temperature, all else being equal. The maximum measured temperature change is $3°\mathrm{C}$,

[11] Due to limitations in the measurement setup, only the in-plane micromotion compensation is measured.

which seems to level out after firing the ovens on a partial duty cycle for $\approx 16\,\text{hrs}$, although the real temperature change could be higher.

$V_{x,y}$ The beam pointing is servoed every few minutes, and the optimal piezo voltage logged. For the first $\approx 1\,\text{day}$ in fig. 5.18b, the pillbox's temperature is constant, and nominally nothing is changing. During this period, the intrinsic drift rates of the Doppler beams are measured to be ${}^{dV_x}\!/_{dt} \approx 0$ for the horizontal axis, and ${}^{dV_y}\!/_{dt} \approx 1.5\,{}^{V}\!/_{\text{day}}$ for the vertical. The noise in the measurement of $\approx 0.5\,\text{v}_{\text{peak}-\text{peak}}$ is assumed to be due to hysteresis in the piezo drivers and shot noise in the ion's fluorescence. Later, the oven is fired and the pillbox temperature is changing, between 27 and 32 hours and in the period after 48 hours, as seen in fig. 5.18a. During these times of temperature change, both axes respond with a change in the measured optimal position. The horizontal axis seems to change by ${}^{dV_x}\!/_{dT} \approx 1.5\,{}^{V}\!/_{\text{K}}$ and the vertical axis by ${}^{dV_y}\!/_{dT} \approx 0.5\,{}^{V}\!/_{\text{K}}$, assuming the intrinsic drift rate remains constant. The piezo drivers have a maximum voltage of 150 V. Assuming the measured temperature change is accurate and the temperature dependencies are constant, the piezos would not be able to follow the change in ion position when the trap is cooled, as this would need a range of ${}_{\Delta}V_x = (T_{hot} - T_{cold}) \cdot {}^{dV_x}\!/_{dT} = (300\,\text{K} - 16\,\text{K}) \cdot 1.5\,{}^{V}\!/_{\text{K}} = 430\,\text{V}$[12]. It is both possible that the temperature change of the pillbox is underestimated by cryostat's diode, and that the temperature dependencies change over temperature. However, in the case these ranges are accurate, the cool-down could be done in e.g. four stages. At each of them, the mirrors would be adjusted by hand and the travel range recuperated. This procedure creates some overhead, as it needs to be repeated

[12]While the vertical axis fairs better, ${}_{\Delta}V_y = 140\,\text{V}$, this is at the limit of the travel range and any slight changes in the dependencies could lead to the travel range exceeding the maximum voltage.

for each of the five axes, potentially four times, and at each cool-down and warm-up cycle.

λ A wavemeter is used to measure the frequencies of the lasers, and the result is subtracted from the nominal $^{40}\mathrm{Ca}^{+}$ resonance , see fig. 5.18c. The 397 and 866 nm lasers are servoed onto the ion's resonance every few minutes, with an estimated servo noise of $< 10\,\mathrm{MHz}$. Therefore, the drifts visible in the measured frequency of these two lasers are due to the wavemeter's intrinsic drift. As can be seen, the drifts are $^{df_{397}}/_{dt} = 100\,^{MHz}/_{48\,\mathrm{hrs}}$ and $^{df_{866}}/_{dt} = 50\,^{MHz}/_{24\,\mathrm{hrs}}$. Assuming the turning points in the drifts are the maximum drifts, that the drift distance in the other direction is of the same magnitude, and that the drifts have a small wavelength dependence, the long term stability of the red lasers' frequency readings should be trusted to $\pm 50\,\mathrm{MHz}$ and the blue lasers to $\pm 100\,\mathrm{MHz}$. These are important to note for trapping ions if reliant on a wavemeter that has been calibrated more than a few hours ago. There was a period of $\approx 8\,\mathrm{hrs}$ where the measured 397 nm laser frequency was suddenly more than 100 MHz from the nominal resonance. This is due to the laser's mode changing and going out of lock. Assuming this laser was already going multi-mode in the preceding $\approx 4\,\mathrm{hrs}$, this would explain why no ions could be trapped during these $\approx 12\,\mathrm{hrs}$. While the 397 and the 866 nm lasers are locked to cavities which are then servoed onto the ion, the 854 nm laser is locked to the wavemeter directly, without a cavity. Therefore, its "measured frequency" is constant. The gaps in the 854 nm laser's frequency indicate that this laser has jumped mode[13], however as this laser is only necessary for de-shelving, its absence does not affect the measurements taken in the loading rate study. The 423 nm laser's frequency is only locked to the wavemeter

[13]Typically, the red lasers jump mode cleanly from one to the next, while the blue lasers have large periods of multi-mode behaviour in-between single modes.

during times of loading, and drifts by $-50\,{}^{MHz}\!/_{\mathrm{hr}}$ otherwise.

f Measurements of the ion's three motional frequencies are shown in fig. 5.19a. There is a repeating pattern with a cycle time of $\approx 30\,\mathrm{min}$, which is presumed to be due to air conditioning, beyond which the frequencies are within measurement noise during the first 24 hrs. After increasing the temperature of the pillbox due to the loading rate study, the effect of a change in temperature is $\sim 0.03\,{}^{\%}\!/_{\mathrm{K}}$ and $\sim 0.1\,{}^{\%}\!/_{\mathrm{K}}$ for axial, and radial motional frequencies, respectively. For two-qubit gates, the secular frequency should have a detuning error δ that is small compared to the gate's rabi frequency Ω, aimed to be $\Omega \approx 2\pi \cdot 3\,\mathrm{kHz}$ [p. 215f, 23]. The presumably air conditioning detuning error δ_{AC} is about $\delta_{AC} = 2\pi \cdot 10\,\mathrm{kHz}$ peak to peak. This would cause significant errors, however these can be avoided by periodic calibration of the shift, as well as by tracking down and mitigating the causes of the periodicity, by e.g. thermal insulation and improved fibre coupling. Measurement noise was about $2\pi \cdot 1\,\mathrm{kHz}$, i.e. the remaining secular frequency drift that is not periodic has an upper bound of $\delta_{res} < 2\pi \cdot 1\,\mathrm{kHz}$. Assuming the worst case $\delta = 2\pi \cdot 1\,\mathrm{kHz}$, the two-qubit gate error would be $\sim 10 \ldots 20\%$ [fig. 6.16, 23], although the secular frequency drift, and therefore two-qubit gate error, could be much better than this. An upgrade in the RF drive chain to increase the stability of the RF amplitude, and therefore of the secular frequencies, is already planned, and future measurements should clarify the levels of frequency stability reached. This new RF drive chain is based on a two-stage high-stability voltage reference, parts LT3042 and LTC6655 from Analog Devices, and a low-noise RF amplifier, part BXMP1042 from RFMW. Preliminary data obtained during this upgrade points towards a measurement-noise limited voltage stability of

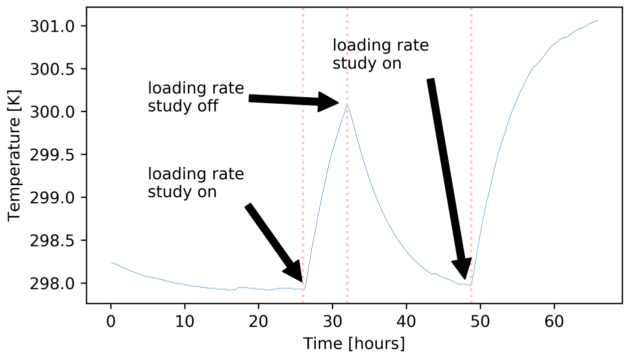

(a)

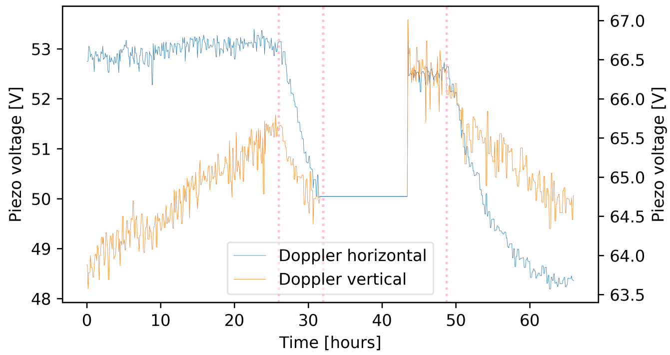

(b)

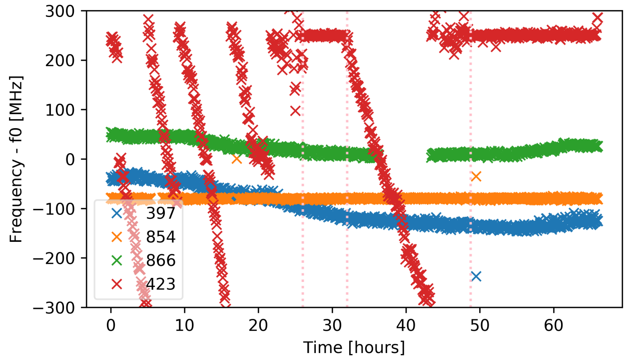

(c)

Figure 5.18: Long term stability of the parameters in the experiment, as explained in sect. 5.4. This data is taken interleaved with the loading rate study, sect. 5.3. (a) During times when the oven is on, the temperature of the pillbox rises, here as measured via the cryostat. For ease of reference, the times that the loading rate study is switched on and off are marked in pink. Outside the loading rate study, ions are only loaded when needed, i.e. on average every 2 hours. (b) The beam pointing is servoed every few minutes, and plotted here. The horizontal Doppler piezo voltage is plotted on the left axis, the vertical one on the right. (c) The wavemeter readings of the lasers. As explained in the text, the 854 nm laser is servoed to the wavemeter, i.e. the difference to the wavemeter as plotted here is constant. The wavemeter is showing some drift on the 397 and 866 nm lasers, however as they are servoed onto the ion, this shows the wavemeter's intrinsic drift. The 423 nm laser is a group shared resource, and only locked to the wavemeter frequency when trying to load; otherwise it is free to drift. Therefore its frequency is showing as constant during the loading rate study, and is showing slow drifts outside of it, with some sudden jumps to the nominal frequency when trying to load.

$\approx 10\,\text{ppm}$ pk-pk, which would be good enough for high-fidelity gates. It also stressed the importance of thermally insulating the SMA cables used for connecting the RF drive chain to the vacuum system.

E_x The micromotion compensation field E_x can be seen in fig. 5.19b. During the first 24 hr period, in which no loading is done, the changes in micromotion compensation field are within measurement noise. There is some concern that the 378 nm laser used for loading charges the trap, leading to instabilities[14]. The work function of Au is $\lambda_{Au} = 240\,\text{nm} \ll 378\,\text{nm}$, i.e. pure gold cannot be charged with a 378 nm laser, but any surface contaminants might have a lower threshold. During the periods of loading, the compensation field E_x changes by ${}^{dE_x}\!/_{dt} = 5\,{}^{V/m}\!/_{12\,\text{hrs}}$, which could be due to the temperature, charging effects, or a combination of both. However, ions are observed to stay trapped at compensation fields as high as $E_x = \pm 1500\,{}^{V}\!/\text{m}$. Therefore, these small changes in compensation field due to temperature or charging should not affect the loading in this trap.

[14]Notably, this effect can lead to a vicious cycle; the ion getting lost due to charging-related instabilities and reloading the ion creating more charging. Also, firing the oven increases its outgassing rate, raising the system pressure, and leading to more background gas collisions, which can contribute to this cycle.

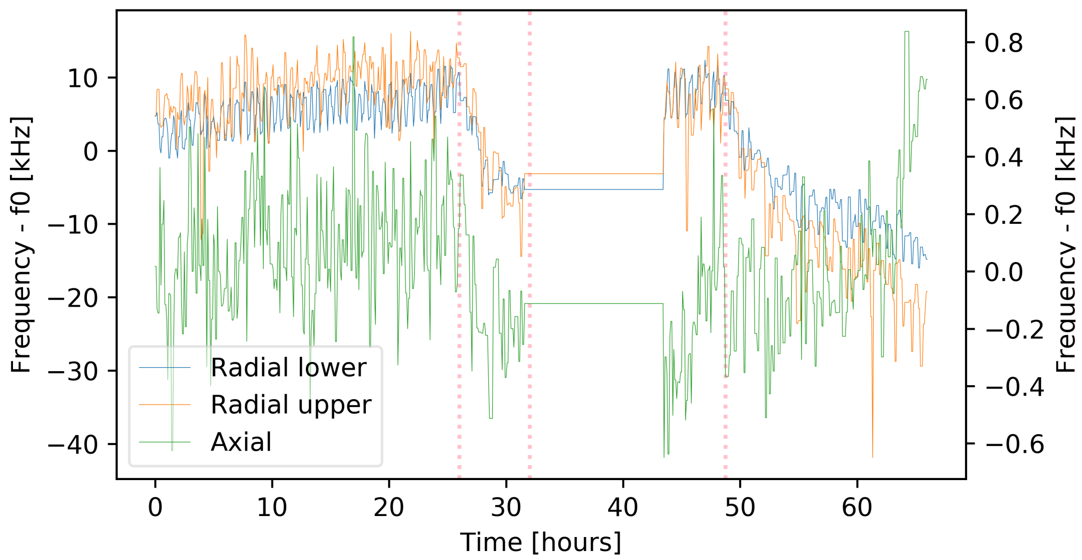

(a)

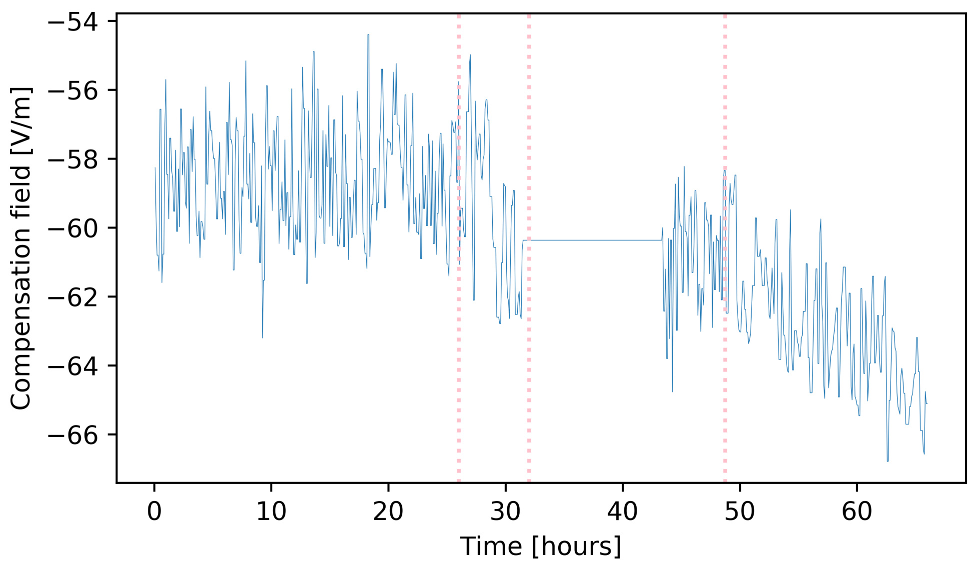

(b)

Figure 5.19: Long term stability of the parameters in the experiment, as explained in sect. 5.4. This data is taken interleaved with the loading rate study, sect. 5.3. (a) Measurements of the ion's motional frequencies, as explained in the text. The mean frequencies f_0 are 1.132 MHz, 6.645 MHz, and 6.932 MHz for the three modes. For each point, a motional frequency experiment is run, see sect. 5.2.6. As can be assumed by looking at fig. 5.9, determining the centre frequency of the signals has about $2\pi \cdot 1$ kHz of shot noise, limiting the accuracy of this method of determining the motional frequencies. The axial frequency is plotted on the right axis and drifts much less than the radial ones. (b) Measurements of the ion's micromotion compensation, as explained in the text. For each point, a set of micromotion correlation experiments are run, see sect. 5.2.8.

5.5 Microwave field across the trap

One of the chief concerns about the design of the trap is the suitability of the magnetic field produced by the microwave currents for two-qubit gates. A high magnetic field gradient leads to a faster gate, which is intrinsically desirable. It also leads to a lower chance of a quantum of motion entering the information bus[15]. One of the primary goals of this thesis is to increase the two-qubit gate speed by a factor of 10.

To verify how well the simulations of the trap agrees with the reality, a set of experiments is run at room temperature. They are a combination of experiments explained in previous sections. Firstly, the rabi flopping times of the σ-polarised $4S_{1/2}|4,+4\rangle \rightarrow 4S_{1/2}|3,+3\rangle$ and the π-polarised $4S_{1/2}|4,+3\rangle \rightarrow 4S_{1/2}|3,+3\rangle$ are measured, as described in sect. 5.2.7. These rabi flopping times can be directly converted to the amplitude of the microwave magnetic field, using eq. 5.30. Secondly, the ion is moved by setting a new compensation field, which can be converted into an ion position, as explained in sect. 5.2.9. These two steps are repeated to get the amplitude of the microwave magnetic field as a function of ion position.

Before these measurements are run, both trim coil currents are calibrated to align the magnetic field vector with the σ-polarised 397&393 nm lasers by maximising the speed of the shelving scheme, see sect. 2.7. Also at each point, several calibration sub-experiments are run; namely a servo to adjust the main static magnetic field such that the ion is resonant to the σ-polarised transition frequency $f_{4\rightarrow 3}$, as defined in sect. 3.12, another servo that adjusts the π-polarised transition frequency $f_{3\rightarrow 3}$ to be resonant with the ion[16], and the thresholding experiment of sect. 5.2.4[17].

[15] Assuming the same motional heating rate.

[16] This is necessary due to the AC Zeeman shift, otherwise all microwave transitions in the $4S_{1/2}$ manifold would be resonant with their calculated values as the magnetic field is calibrated.

[17] The further the ion is from its micromotion compensated position, the stronger the lasers couple to the trapping field $f_{RF} = 63.98\,\mathrm{MHz}$. This additional Doppler shift greatly reduces the fluorescence induced by the Doppler cooling lasers, necessitating new thresholds to be

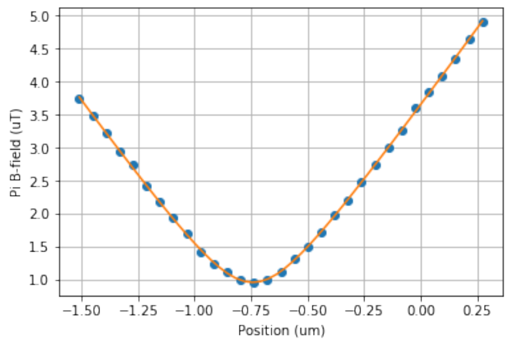

(a)

Figure 5.20: Measured microwave magnetic field $B_{MW,\pi}(y)$ as a function of position of the ion. The ion is displaced by applying a compensation field via the DC electrodes, here converted to a length displacement via a scaling derived in sect. 3.13. At each point, the π-time τ_{pi} of the rabi flops of the π-polarised $4S_{1/2}|4,+3>$ to $4S_{1/2}|3,+3>$ transition is measured; from this a magnetic field is calculated using eq. 5.30.

The compensation field can be converted to a displacement in two ways. Firstly, by using the imaging system, as described in fig. 4.3. Secondly, as the motional frequencies of the trap are known, see sect. 5.2.6, they can be used in a quadratic pseudo-potential approximation to predict the displacement of the ion due to a small electric field, i.e. that due to the compensation field. Both methods yield a scale of $0.78 \cdot 10^{-3}\,\mu\mathrm{m}/\left(\mathrm{V/m}\right)$. The zero position is the one which is micromotion compensated, see sect. 5.2.8. It is found to be at a field of -69 $^{V}/_{\mathrm{m}}$, i.e. a displacement of $-69\,^{V}/_{\mathrm{m}} \cdot 0.78\,\mathrm{nm}/\left(^{V}/_{\mathrm{m}}\right) = 54\,\mathrm{nm}$, which is a small displacement given the scale of the trap.

The amplitude component of the microwave magnetic field in the π direction has the form $B_{MW} = B(y)\cos(\omega t + \phi)$[18]. Its position dependent component $B(y)$ is given by eq. 2.5

$$B(y) = \frac{\hbar\Omega(y)}{R_{11}} \tag{5.30}$$

using the reduced Planck constant $\hbar$, the Rabi frequency as a function of position

established. It also leads to a loss of the ion beyond $\sim 2\,\mu$m from the compensated position.

[18] This section uses the axes defined in sect. 2.9.4.

$\Omega(y)$, and the transition matrix element R_{11}. The magnetic field is calculated from the measured Rabi frequencies via equation 5.30 with the help of a Python toolkit, [25], and plotted in fig. 5.20.

Due to the phase shift θ between the finite length of the "U" bend of the microwave line, the current in the two wires is not the same; this gives rise to a finite magnetic field minimum B_0. As the distance to the wires is big compared to the achievable displacement of the ion, a linear approximation of the field is used:

$$B(y) = B_0 + e^{i\theta}(y - y_0) \tag{5.31}$$

Taking the absolute value of this complex field gives

$$\begin{aligned} |B(y)| &= B(y) \cdot \bar{B}(y) = \\ &= \sqrt{B_0^2 + (B' \cdot (y - y_0))^2 + 2 \cdot B_0 \cdot B' \cdot (y - y_0) \cdot cos(\theta)} \end{aligned} \tag{5.32}$$

Fitting the results to equation 5.32 yields a magnetic field minimum of $B_0 = 1.0\,\mu\text{T}$ at a distance of $x_0 = 740\,\text{nm}$ and an asymptotic maximum magnetic field gradient B' of $4.6\,{}^{T}\!/\text{m}$. This measurement uses the setup with two synths at a nominal microwave input power of $p_{in} = 21.7\,\text{dBm}$, see sect. 3.12, while the simulations assume $p_{sim} = 30\,\text{dBm}$. The magnetic field minimum and gradient scale as the amplitude of the magnetic field, i.e. as $10^{\frac{p_{sim}-p_{in}}{20\,dB}}$. Using this to scale the measured quantities to p_{sim}, the magnetic field minimum B_0 becomes $\tilde{B}_0 = B_0 \cdot 10^{\frac{p_{sim}-p_{in}}{20\,dB}} = 2.6\,\mu\text{T}$, and the magnetic field gradient becomes $\tilde{B}' = B' \cdot 10^{\frac{p_{sim}-p_{in}}{20\,dB}} = 12.0\,{}^{T}\!/\text{m}$.

The trap axes are tilted by 15° to allow for Doppler cooling all axes with one laser beam path. The effect of this tilt is neglected and the above calculations assume a 0° tilt, but this would only change the scaling by a small amount.

As stated in sect. 2.3, the goal of this experiment is to decrease the gate time by a factor of ten over the previous time of $t_{g2} = 3.25\,\text{ms}$, which used a field

	Measured	Simulated
B'	$12\,{}^{T}\!/\mathrm{m}$	$28\,{}^{T}\!/\mathrm{m}$
$B_{0,\pi}$	$2.6\,\mu\mathrm{T}$	$6\,\mu\mathrm{T}$
$B_{0,\sigma}$	$125\,\mu\mathrm{T}$	$300\,\mu\mathrm{T}$
x_0	$-740\,\mathrm{nm}$	$-300\,\mathrm{nm}$

Table 5.3: Comparison of measured to simulated magnetic field properties, normalised to 1 W. The first three quantities are the π-polarised magnetic field gradient B', the π-polarised magnetic field minimum $B_{0,\pi}$, and the σ-polarised magnetic field minimum $B_{0,\sigma}$. These all scale linearly with the applied microwave amplitude. The deduced scaling factor is $C_{amp} = 2.35 \pm 0.05$, i.e. there is a difference in power by a factor of $C_{pwr} = C_{amp}^2 = 5.5 \pm 0.2$.

gradient of $B'_{prev} = 7\,{}^{T}\!/\mathrm{m}$ with a microwave power of 1 W [supplementary p.1, 26]. The ratio of matrix elements r is 1.425 705, see sect. 2.9.4; therefore, accounting for the effect of the different qubits, the field gradient that this experiment needs to achieve to match the previous time is $\frac{B'_{prev}}{r} = 5\,{}^{T}\!/\mathrm{m}$. Therefore, using the same microwave power $P = 1\,\mathrm{W}$, the gate time becomes $t = t_{g2} \cdot \frac{B'_{prev}}{r \cdot \tilde{B}'} = 1.3\,\mathrm{ms}$, a factor of 2.4 increase in speed.

To achieve the factor-of-ten gate speed increase, the magnetic field gradient should be $B'_{target} = 10 \cdot \frac{B'_{prev}}{r} = 50\,{}^{T}\!/\mathrm{m}$. As shown in tab. 5.3, the microwave field gradient is predicted to be $B'_{meas} = 12\,{}^{T}\!/\mathrm{m}$ at a microwave power of 1 W; therefore to achieve the magnetic field gradient $B'_{target} = 50\,{}^{T}\!/\mathrm{m}$, the required power would be $P_{target} = \left(\frac{50\,{}^{T}\!/\mathrm{m}}{12\,{}^{T}\!/\mathrm{m}}\right)^2 \cdot 1\,\mathrm{W} = 17\,\mathrm{W}$. This is quite high, but should be experimentally achievable by e.g. using a low duty cycle of the microwaves. However, as explained below, the required power P_{target} should decrease drastically at cryogenic temperatures.

A comparison of the measured to the simulated magnetic field properties can be found in tab. 5.3. The discrepancy of measured to simulated power could be explained by the combination of a few factors. Firstly, the trap was measured at room temperature, but simulated at cryogenic temperatures, e.g. using the cryogenic resistivity of Au. The higher resistivity of Au at room temperature would decrease the quality factor of the resonator, decreasing the magnetic field.

Secondly, the trap resonator length was selected based on performance at LN_2-temperatures, see sect. 5.1, i.e. the resonator frequency would be slightly wrong at room temperature. Thirdly, the microwave input power p_{in} is measured on the air side cable entering the vacuum system, i.e. any losses in the cabling and the microwave PCB also show up as a decrease in the magnetic field.

5.6 Initial cryogenic experiments

For initial testing of the cryogenic system, the pillbox without an installed ion trap was cooled, and $\sim 20\,\mathrm{K}$ was reached[19]. To test the limits of the cryogenic system, the temperature was further reduced, and 4.2 K was successfully reached, and no vacuum leaks were found apart from a brief spike in pressure – presumed to be a transient virtual leak – of the vacuum system during cool-down. At this temperature the horizontal orientation of the cryostat became problematic; as liquid helium accumulated on the inside of the cryostat, it also flowed towards the region meant for the gaseous helium exhaust. That region is only sealed with nitrile "o"-rings, which then leaked liquid helium. This is not a problem during the targeted $\sim 20\,\mathrm{K}$ operation, as there is no accumulation of liquid helium. Future improvements of the system will include a $5°$ angle of the cryostat with respect to a horizontal orientation, keeping any accumulation of liquid helium contained. This might also decrease the liquid helium consumption rate at $\sim 20\,\mathrm{K}$.

The results of the magnetic field gradient measurements of the previous section were obtained at room temperature, and a discrepancy of the simulation to the measurement could be due to temperature. To test the effect of the first two factors, the measurement was attempted at cryogenic temperatures. Unfortunately, during the first cool-down to $\sim 20\,\mathrm{K}$ of the pillbox with a trapped ion, two problems arose.

Firstly, the pillbox moved by $\sim 0.5\,\mathrm{mm}$. This is problematic as the travel

[19] As described in sect. 3.3.4.1, the pillbox legs had to be changed from macor to SS.

range of the piezo-mirrors is only $\sim 150\,\mu$m at the ion and the imaging system also needs to be moved. The movement of the pillbox is likely due to the thermal braid being too stiff; the pillbox's legs are bent instead. Presumably, this is also the reason why the ceramic legs, see sect. 3.3.4.1, broke during cool-down; their brittleness would cause them to shatter instead of bend. A new, more flexible braid was designed and is on order. This could also cause the ceramic legs to become viable again.

Secondly, the ion was lost during cool-down at $-50°$C, and subsequent loading attempts failed. Many causes were investigated, and it was narrowed down to the DC electrodes. While none of the DC electrodes showed a short to ground or each other, applying square pulses to the DC electrodes and fitting the rise time to a model of the RC filter, it was inferred that three electrodes have become disconnected, one of them used for trapping in the central region. Fortunately, another set of electrodes suitable for trapping is still available, and trapping subsequently succeeded again. The cause for the disconnected electrodes is still under investigation. As the problem arose during cool-down and is located between the RC filter components and the air side, the most likely suspects are cracks in the PCB, e.g. due to the fragile soldering. The new batch of PCBs should be more resilient against this, as they have fewer components, are easier to solder, and have dielectric on top of the traces, which should protect them against damage.

Both the thermal braid, and the PCBs are planned to be swapped out in the near future, which should then allow for low temperature investigations.

Chapter 6

Conclusion

The overarching aim of this experiment is the successful implementation of an order-of-magnitude improvement over the current highest fidelity microwave based two-qubit gate, in both gate speed and fidelity. While a two-qubit gate has not yet been achieved in this experiment, a number of meaningful interim results suggest that a successful implementation will be accomplished in the near future.

The development of this new apparatus involved the creation of laboratory infrastructure, including a new vacuum bakeout system, for use in multiple quantum information experiments across the physics department. As shown over the course of sections 4.1 to 4.3, a number of improvements were made to existing manufacturing processes, including new insights into the creation of more reliable electroplating processes. The eutectic bonding described in sect. 4.2.2 provides a new and promising technique for ion trap chips; it creates a robust vacuum-compatible joint with high thermal conductivity and good reliability under thermal cycling.

An experimental system to house the ion trapping chips was designed and implemented. The ion-to-electrode distance exhibited in our experiments is among the smallest ever attempted in trapped ions. It introduces or exacerbates issues such as small trap depth and the requirement of good pointing stability. The fact that we have successfully trapped ions demonstrates that we have navigated these issues successfully. Despite the small ion height, a single ion lifetime of 2 hours at room temperature was observed, which is sufficient not to increase experimental difficulty or present problems further down the line. Additionally, a number of ion diagnostics and servos have been implemented, on which future progress will be based. The cryogenic system was able to cool the ion trap down to the liquid helium boiling point of 4.2 K, and after the described upgrade should cause no obstacle for experimental progress at cryogenic temperatures of ~ 16 K.

Rabi-flopping has been observed in single trapped ions of ^{43}Ca in the ground state hyperfine manifold $4S_{1/2}$. This has been used to calibrate the size and direction of the microwave magnetic field as a function of position, and thus obtain the magnetic field gradient. The results suggest that the implementation of the order-of-magnitude speed improvement are just about possible at room temperature and should be well-achievable at cryogenic temperatures.

Further improvement works to the system currently underway include a more flexible braid and an angled arrangement of the cryostat, to allow for cryogenic results to be obtained down to 4.2 K at a potentially reduced helium consumption, a new set of easier to solder PCBs which could solve the problems of electrode connection reliability during cooldown as well as providing a thermal sensor to corroborate the readings from the cryostat thermal sensor. Additional improvements include a new high-stability RF drive chain and a feedback & feedforward circuit stabilising the current for the main magnetic field coils, potentially lowering the two-qubit error rate, an 850 nm laser beam path, leading to a higher contrast readout, and the inclusion of ^{88}Sr in the oven tubes for potential two-species gates.

Owing to the absence of decoherence associated with photon scattering, the microwave method offers, in principle, the prospect of the highest fidelity two-qubit gate in trapped ion experiments, and possibly in any kind of scalable implementation of quantum information processing. This in any case is the next goal.

This fidelity, along with the speed improvement, is the main DiVincenzo criterion, #4, which is not yet fulfilled to a standard needed for quantum CCD based, localised quantum computing. The results presented within this work, supplemented by future experimental results gained once current improvements are completed, should therefore pave the way towards a near-term, small to medium scale quantum computer.

Bibliography

[1] , 2016. Saes rial vacuum, saes group.
URL http://saesrial.com/wp-content/uploads/2017/02/SAES-RIAL-in-Atom-Trap-2016_rev-LM.pdf

[2] Alcock, C. B., Itkin, V. P., Horrigan, M. K., jul 1984. Vapour pressure equations for the metallic elements: 298–2500k. Canadian Metallurgical Quarterly 23 (3), 309–313.
URL https://doi.org/10.1179/cmq.1984.23.3.309

[3] Allcock, D. T. C., 2011. Surface-electrode ion traps for scalable quantum computing.

[4] Amils, R. I., 2016. Thermal conductivity of silver loaded conductive epoxy from cryogenic to ambient temperature and its application for precision cryogenic noise measurements. Cryogenics.
URL http://dx.doi.org/10.1016/j.cryogenics.2016.03.001

[5] AspenCore, I., 2014. Passive low pass filter.
URL https://www.electronics-tutorials.ws/filter/filter_2.html

[6] Ballance, T. G., Goodwin, J. F., Nichol, B., Stephenson, L. J., Ballance, C. J., Lucas, D. M., 3 2018. A short response time atomic source for trapped ion experiments. Review of Scientific Instruments 89 (5), 053102.
URL http://dx.doi.org/10.1063/1.5025713

[7] Beukema, G. P., 2015. Vacuumtechnology course.

[8] Britton, J., 2019. Sinara hardware github.
URL https://github.com/sinara-hw/meta/wiki

[9] Bruzewicz, C. D., Chiaverini, J., McConnell, R., Sage, J. M., jun 2019. Trapped-ion quantum computing: Progress and challenges. Applied Physics Reviews 6 (2), 021314.
URL http://dx.doi.org/10.1063/1.5088164

[10] CoorsTek, I., 2016. Advanced Alumina Materials & Manufacturing Processes.

[11] Coyne, D., 4 2006. Polyimide Hydrocarbon Outgas Rate and Temperature Limits for the Isolation System Actuator. Laser Interferometer Gravitational Wave Observatory, ligo-t04003201-d Edition.

[12] CRC, 2015. CRC Handbook of Chemistry and Physics, 96th Edition. CRC Press.

[13] Cremonini, A., Mariotti, S., Roda, J., 10 2012. An overview on packaging of microwave electronic devices operating in a cryogenic environment. Cryogenics 52 (10), 445–451.
URL http://dx.doi.org/10.1016/j.cryogenics.2012.04.020

[14] Cross, T., 2016. Technology quarterly, after moore's law.
URL https://www.economist.com/technology-quarterly/2016-03-12/after-moores-law

[15] Cryotronics, L. S., 2016. Temperature Measurement and Control Catalog.
URL https://www.lakeshore.com/Documents/LakeShoreTC_l.pdf

[16] DAsaro, L. A., 1980. Electroless gold plating on iii-v compound crystals. Journal of The Electrochemical Society 127 (9), 1935.
URL http://dx.doi.org/10.1149/1.2130040

[17] Demtroeder, W., 2008. Laser Spectroscopy. Springer Berlin Heidelberg.

[18] DiVincenzo, D. P., 2000. The physical implementation of quantum computation. Fortschritte der Physik 48 (9-11), 771–783.
URL https://arxiv.org/abs/quant-ph/0002077

[19] Feynman, R. P., jun 1982. Simulating physics with computers. International Journal of Theoretical Physics 21 (6-7), 467–488.
URL http://dx.doi.org/10.1007/bf02650179

[20] Griot, M., 2010. Gaussian beam optics.
URL http://www.pa.msu.edu/courses/2010fall/phy431/PostNotes/PHY431-Notes-GaussianBeamOptics.pdf

[21] Hanlon, J., 2003. A User's Guide to Vacuum Technology, Third Edition. Wiley-Interscience, iSBN 0-471-27052-0.

[22] Harris, N. S., 1989. Modern Vacuum Practice. McGraw-Hill Book Company.

[23] Harty, T. P., 2013. High-fidelity microwave-driven quantum logic in intermediate-field 43ca+.

[24] Harty, T. P., Allcock, D. T. C., Ballance, C. J., Guidoni, L., Janacek, H. A., Linke, N. M., Stacey, D. N., Lucas, D. M., nov 2014. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. Physical Review Letters 113 (22).
URL http://dx.doi.org/10.1103/physrevlett.113.220501

[25] Harty, T. P., Hanley, R., Wolf, J., 2019. Ionphys: Toolkit for calculating state energies, transition matrix elements, etc.
URL https://github.com/OxfordIonTrapGroup/ion_phys

[26] Harty, T. P., Sepiol, M. A., Allcock, D. T. C., Ballance, C. J., Tarlton, J. E., Lucas, D. M., sep 2016. High-fidelity trapped-ion quantum logic using near-

field microwaves. Physical Review Letters 117 (14).
URL http://dx.doi.org/10.1103/physrevlett.117.140501

[27] Honig, Hook, 1960. Vapor pressure data for some common gases. RCA Review 21.3, 360–368.

[28] Horowitz, M. A., mar 2019. Quantum Computing. National Academies Press.
URL http://dx.doi.org/10.17226/25196

[29] House, M. G., sep 2008. Analytic model for electrostatic fields in surface-electrode ion traps. Physical Review A 78 (3).
URL http://dx.doi.org/10.1103/physreva.78.033402

[30] Imreh, G., 2008. Implementing segmented ion trap designs for quantum computing.

[31] Inlek, I., Crocker, C., Lichtman, M., Sosnova, K., Monroe, C., jun 2017. Multispecies trapped-ion node for quantum networking. Physical Review Letters 118 (25).
URL http://dx.doi.org/10.1103/physrevlett.118.250502

[32] Instruments, N., 2003. Tx-line: Transmission line calculator.

[33] Kim, W., Wang, Q., Jung, K., Hwang, J., Moon, C., 2004. Application of au-sn eutectic bonding in hermetic RF MEMS wafer level packaging. In: 2004 1st IEEE Consumer Communications and Networking Conference (IEEE Cat. No.04EX745). IEEE.
URL http://dx.doi.org/10.1109/isapm.2004.1288016

[34] Kramida, A., Ralchenko, Y., 1999. Nist atomic spectra database, nist standard reference database 78.
URL http://dx.doi.org/10.18434/t4w30f

[35] Li, D., Shang, Z., She, Y., Wen, Z., 5 2017. Investigation of au/si eutectic wafer bonding for MEMS accelerometers. Micromachines 8 (5), 158.
URL http://dx.doi.org/10.3390/mi8050158

[36] Loschnauer, C., 2020. Two-qubit near-field microwave gates on 43ca+ (unpublished dphil thesis).

[37] Lucas, D., Ramos, A., Home, J., McDonnell, M., Nakayama, S., Stacey, J.-P., Webster, S., Stacey, D., Steane, A., jan 2004. Isotope-selective photoionization for calcium ion trapping. Physical Review A 69 (1).
URL http://dx.doi.org/10.1103/physreva.69.012711

[38] Martinez, W. E., Gregori, G., Mates, T., 3 2010. Titanium diffusion in gold thin films. Thin Solid Films 518 (10), 2585–2591.
URL http://dx.doi.org/10.1016/j.tsf.2009.07.187

[39] Matijasevic, G. S., Lee, C. C., Wang, C. Y., 2 1993. Au-sn alloy phase diagram and properties related to its use as a bonding medium. Thin Solid Films 223 (2), 276–287.
URL http://dx.doi.org/10.1016/0040-6090(93)90533-u

[40] Molmer, K., Sorensen, A., mar 1999. Multiparticle entanglement of hot trapped ions. Physical Review Letters 82 (9), 1835–1838.
URL http://dx.doi.org/10.1103/physrevlett.82.1835

[41] Negnevitsky, V., Marinelli, M., Mehta, K. K., Lo, H.-Y., Flühmann, C., Home, J. P., nov 2018. Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register. Nature 563 (7732), 527–531.
URL http://dx.doi.org/10.1038/s41586-018-0668-z

[42] Noertershaeuser, W., Blaum, K., Icker, K., Müller, P., Schmitt, A., Wendt, K., Wiche, B., may 1998. Isotope shifts and hyperfine structure in the 3d 2dj to 4p 2pj transitions in calcium II. The European Physical Journal D -

Atomic, Molecular and Optical Physics 2 (1), 33–39.
URL http://dx.doi.org/10.1007/s100530050107

[43] Ospelkaus, C., Langer, C., Amini, J., Brown, K., Leibfried, D., Wineland, D., aug 2008. Trapped-ion quantum logic gates based on oscillating magnetic fields. Physical Review Letters 101 (9).
URL http://dx.doi.org/10.1103/physrevlett.101.090502

[44] Peacock, R. N., 1980. Practical selection of elastomer materials for vacuum seals. J. Vac. Sci. Technol. 17.
URL http://dx.doi.org/10.1116/1.570380

[45] Pino, P., 2007. Intermateability of sma, 3.5mm and 2.92mm connectors. Microwave Journal.

[46] Products, A., 2013. Cryogenic liquid containers.
URL http://www.airproducts.com/~/media/files/pdf/company/safetygram-27.pdf

[47] Radebaugh, R., 2019. Nist cryogenic material properties.
URL https://trc.nist.gov/cryogenics/materials/materialproperties.htm

[48] Ren, F., Abernathy, C. R., Pearton, S. J., Wisk, P. W., 3 1994. Thermal stability of ti/pt/au nonalloyed ohmic contacts on InN. Applied Physics Letters 64 (12), 1508–1510.
URL http://dx.doi.org/10.1063/1.111874

[49] Roos, C., jan 2008. Ion trap quantum gates with amplitude-modulated laser beams. New Journal of Physics 10 (1), 013002.
URL http://dx.doi.org/10.1088/1367-2630/10/1/013002

[50] Ruth, H. N., 2002. Technical note 02-053: The thermal conductivity of macor at intermediate cryogenic temperatures. Tech. rep., University of Southamp-

ton.
URL http://tnweb.jlab.org/tn/2002/02-053.pdf

[51] Sepiol, M. A., 2016. A high-fidelity microwave driven two-qubit quantum logic gate in 43ca+.

[52] Simon, N. J., December 1994. Cryogenic properties of inorganic insulation materials for iter magnets: A review.
URL http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/35/058/35058184.pdf

[53] Slichter, D., Allcock, D., 2015. Trap fabrication update: new process flow and magtrap samples.

[54] Stephenson, L. J., Nadlinger, D. P., Nichol, B. C., An, S., Drmota, P., Ballance, T. G., Goodwin, J. F., Schafer, V. M., Lucas, D. M., Ballance, C. J., 2019. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. Unpulished paper.

[55] Tarlton, J. E., 2018. Probing qubit memory errors at the 10e-5 level.

[56] Technology, E., 2009. Epo-tek tech tips.
URL http://www.epotek.com/site/technical-material/technical-tips.html

[57] Tommaseo, G., Pfeil, T., Revalde, G., Werth, G., Indelicato, P., Desclaux, J. P., aug 2003. The gj-factor in the ground state of ca+. The European Physical Journal D - Atomic, Molecular and Optical Physics 25 (2), 113–121.
URL http://dx.doi.org/10.1140/epjd/e2003-00096-6

[58] Wolf, A. L., van den Berg, S. A., Gohle, C., Salumbides, E. J., Ubachs, W., Eikema, K. S. E., sep 2008. Frequency metrology on the 4s 2s0.5 to 4p 2p0.5

transition in 40ca+ for a comparison with quasar data. Physical Review A 78 (3).
URL http://dx.doi.org/10.1103/physreva.78.032511

[59] Wolf, A. L., van den Berg, S. A., Ubachs, W., Eikema, K. S. E., jun 2009. Direct frequency comb spectroscopy of trapped ions. Physical Review Letters 102 (22).
URL http://dx.doi.org/10.1103/physrevlett.102.223901

[60] Young, W., August 1996. LIGO Vacuum Compatible Materials List. Laser Interferometer Gravitational Wave Observatory, ligo-e960050-a Edition.

[61] Zhou, L., 6 2000. Low resistance ti/pt/au ohmic contacts to p-type GaN. Applied Physics Letters 76 (23), 3451–3453.
URL http://dx.doi.org/10.1063/1.126674

Chapter 7

Appendix

7.1 Electroplating Process Development

Electroplating process development proved difficult. The development is done in-house and wafers are processed in the clean rooms at the Clarendon Laboratory, University of Oxford. The goal of this process development was to produce a mirror-like finish of the gold, as this was associated with a low heating rate in a previous experiments, [p. 145, 51] and [53].

Seed layers In order to improve the adhesion of the gold layer, the sapphire wafer is first coated in a layer of titanium. The sputterer used is a K675X from Emitech[1]. After creating a high vacuum with a turbo pump, argon is bled into the chamber, creating a pressure of $\approx 7 \cdot 10^{-3}$ mbar. The titanium is cleaned for 30 s with a current of 150 mA, in order to remove the oxide layer that has formed after coming into contact with ambient air while changing wafers. The sputterer's thickness measurement is calibrated for a planetary motion holder. However, a static holder is sufficient for

[1]It was later discovered that the adhesion of the gold to the sapphire was insufficient; wire bonding tests pulled the gold off in a similar way to the problems described in 4.3.1. The gold layers also failed the "scotch tape test". The problem was tracked down to the titanium adhesion layer; the sputterer was assumed to be miscalibrated. Starting with wafer number 30, tests are prepared using evaporating service run by Rick Makin, Physics Department, University of Oxford, and pass the "scotch tape test" and are able to be wire bonded onto.

a flat, featureless wafer. This improves the sputtering rate by a factor of $\approx$ 2 over the measured thickness. Thus, the titanium layer's thickness is set to 5 nm, in order to get $\approx$ 10 nm, and the gold layer's thickness is set to 50 nm, to get $\approx$100 nm. Between the layers of titanium and gold, the argon leak is stopped, the chamber is pumped to a high vacuum again, and finally argon is again bled into the chamber. After processing wafers, the sputterer is shut down while under vacuum, in order to minimize oxidation of the titanium.

Stirring All instructions for electroplating solutions recommend stirring of the solution. The stirring plate used for the runs 1 through 4 is a UC151 stirrer from Stuart, with a maximum stirring speed of 2000 rpm, set by an analogue "1-9" dial, thus the corresponding stirring speeds in tab. 7.1 are estimated and assume the dial is linear. This limits the accuracy with which stirring speeds can be reproduced. Subsequent runs use an SI-0301 MagStir Genie programmable magnetic stirrer from Scientific Industries. It is set to reverse stirring direction every 40 s, a ramp rate of 300 rpm/s, a pause between directions of 5 s, and to its maximum stirring strength. The periodic reverse in stirring direction helps to avoid anisotropic thicknesses across the wafer. The size of the stirring bars is denoted as $s = 4.5\,\mathrm{mm} \cdot 12\,\mathrm{mm}$, $m = 6\,\mathrm{mm} \cdot 30\,\mathrm{mm}$, and $l = 10\,\mathrm{mm} \cdot 60\,\mathrm{mm}$.

Anode The anode used for runs 1-5 is a 99.9% pure, Ø0.06 mm wires, 90 mm wide, and 120 mm long woven platinum mesh, part PT008710 from Goodfellow Cambridge Ltd, see fig. 7.1a. In run 1, it was electrically contacted with a Cu frame built into the PP frame. This caused the electroplating solution to turn green, a sign of copper contamination. Subsequent runs with this anode electrically contact it above the liquid level via a strip of Ti, 99.7% pure, 0.89 mm thick, part 11305878 by Fisher Scientific UK Ltd. However, the PP frame of this anode has spaces for liquid to be trapped

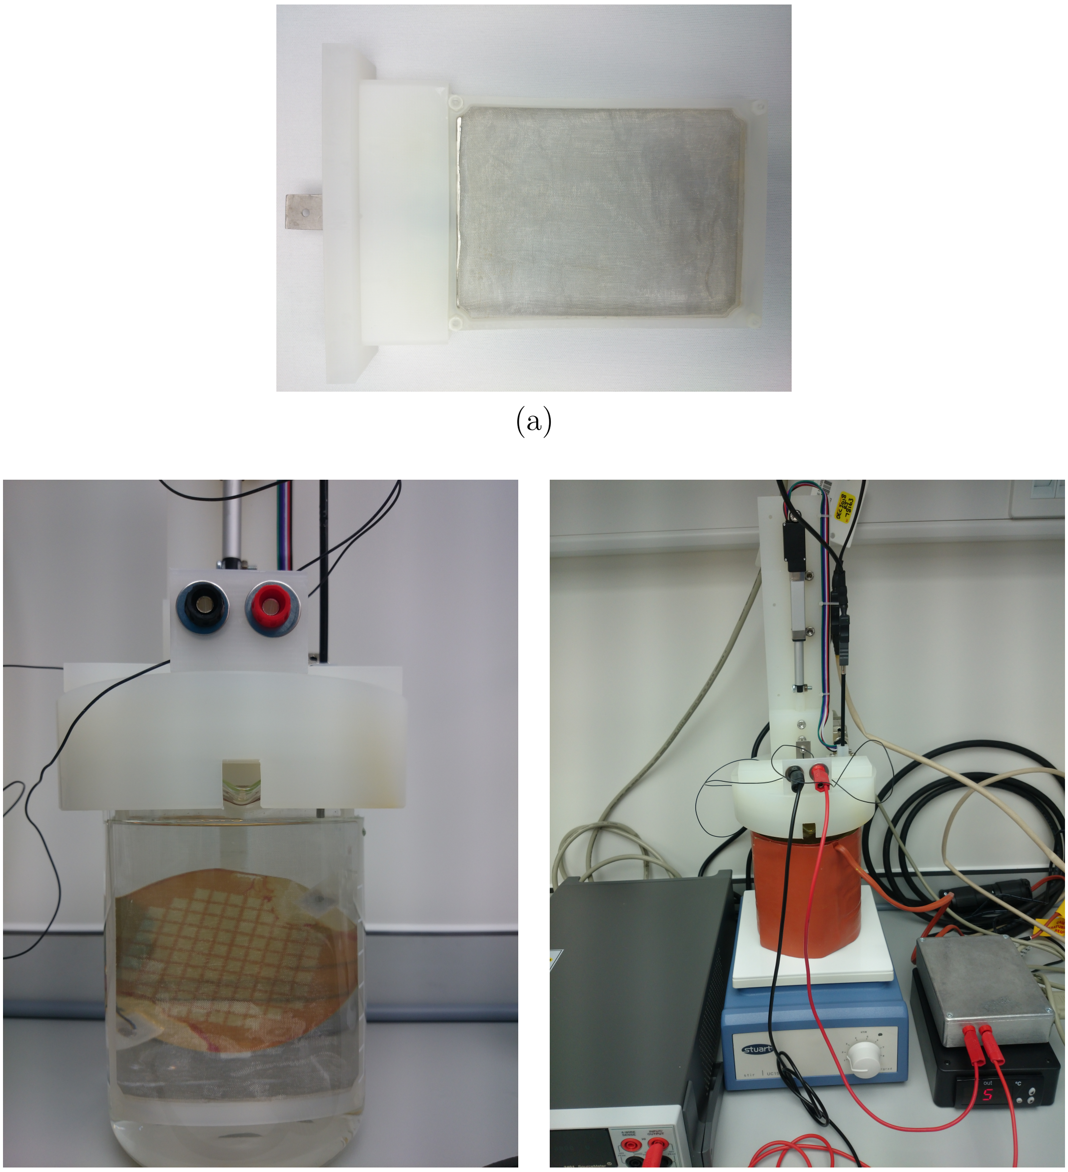

(a)

(b)

(c)

Figure 7.1: Electroplating setup, with the anode (a), the front of the electroplating setup (b), and a wired overview of the entire setup (c). At this stage in development, the ion traps used had a form factor of 5 mm x 5 mm, hence these wafers feature a square pattern instead of the later rectangular one.

in, in particular the space reserved for the Cu contact frame. Starting with wafer 9, and in an effort to increase the cleanliness of the system, the anode is replaced with a mesh of 1 mm thick titanium with a 1 μm plating of platinum, part TiPTA-01-MS-010 from William Gregor Ltd. For a similar reason, all screws of the electroplating setup are changed from SS to PP after run 1.

Solution The types of electroplating solutions used are TSG-250 from Transene, ECF64D & ECF60 from Metalor, and Gold-SF from Metakem. Using the TSG-250, two bottles of their unused solution had different colours, calling their quality control into question, potentially impacting reproducibility. Also, the plated thickness, as measured by a wafer prober, was found to be thicker than intended. There was an inconsistency on Transene's datasheet; the plating speed was specified in two different ways, one equivalent to a gold deposition rate of $\gamma_1 = 0.25\,{}^{\mu m}/_{\min}$ at a current density of $\rho_1 = 4.3\,{}^{mA}/_{\mathrm{cm}^2}$, the other one equivalent to $\gamma_2 = 0.085\,{}^{\mu m}/_{\min}$ at the same current density $\rho_1 = \rho_2$. The company has since corrected this mistake, but initial wafers were plated using a current density $\frac{\gamma_1}{\gamma_2} \approx 3$ times the intended one.

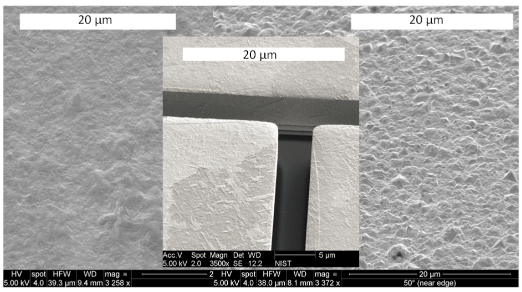

Figure 7.2: Comparison of electroplated wafers. From left to right: a run resulting in fairly shiny Au, the target quality of (patterned) Au [53], and a run resulting in red, burnt-looking Au.

Pattern While the electroplating process should largely be unaffected by the patterning[2], some of the wafers *are* patterned before electroplating. This is

[2]This is assuming the effect of the patterning on the exposed conductor area is taken into account, and the current reduced to keep the current density constant.

done to emulate the final electroplating as closely as possible. Masks from JD Photodata Ltd are used, consisting of 127 mm · 127 mm · 2.2 mm quartz, onto which a layer of chrome is printed, at a tolerancing of 0.2 μm. However, in wafer 19, the exposed photoresist used for patterning was not completely washed off during development. By eye, the exposed areas looked to be free of photoresist, but a thin layer remained and only narrow regions *between* the exposed and unexposed photoresist were completely free of photoresist. This caused the subsequent electroplating to produce much thicker layers in these gaps than intended, while the areas with the thin layer of photoresist remained pristine. As the plating in the gaps framed these pristine areas, this initially made it look like the plating succeeded in the same mirror-finish quality of the seed layer of Au. SEM images then clarified the situation.

Procedure The wafer and anode are rinsed in deionised water and then inserted into the heated electroplating solution, after which the current supply was connected and turned on. Two improvements on this procedure are made; firstly, in order to both reduce bubble formation by increasing the wettability of the wafer, and to clean it, an PT7160 RF O_2-plasma barrel etcher from Quorum Technologies is used. The pressure inside the barrel is first reduced to $\approx 2 \cdot 10^{-2}$ mbar, after which O_2 is leaked into the chamber, rising the pressure to $\approx 7 \cdot 10^{-1}$ mbar. The etching power is set to $^{667}/_{1000}$, corresponding to 100 W, and is used for 1 min. Secondly, there were concerns over the electrical potential of the cathode with respect to the anode while the current supply was still unconnected. To alleviate this concern, the current supply is switched to supply a voltage of 0.01 V, before inserting the wafer. This means the plating of wafer is starting as soon as the wafer is inserted, albeit with a much smaller current density.

Temperature The temperature controller used initially, part SDC240JC-A from BriskHeat Corp., only has an “on/off” switch, and can only be set to activate heating 1°C from the setpoint, and then tending to overshoot by 1°C. This is problematic, as the temperature stability is a critical parameter for electroplating. Starting with wafer 10, this is switched out with the PID based temperature controller mentioned in sect. 4.1.1, that typically uses a duty cycle of $\approx 5\%$ to maintain the temperature of the beaker. In tab. 7.1, the temperature error is denoted as the periodic under/overshoot up to wafer 9; from then on the amplitude of the remaining temperature oscillations are 0.1°C, and the indicated error is the initial overshoot caused by quickly removing the temperature probe during insertion of the wafer. To reduce the magnitude of this overshoot, the maximum duty cycle of the temperature controller is set to 7% after the electroplating solution reached its temperature. As a future improvement, the electroplating setup should feature the ability to insert a wafer without removing the temperature probe.

Results After each run, the quality of the Au was judged by eye to see how close to a mirror-like finish it looked. For some runs, SEM images were taken and compared to reference SEM images. A major problem in the process development was an overestimation of how mirror-like the Au of the experiments [p. 145, 51] and [53] was. As the SEM images of the better wafers came out slightly rougher than the reference, the judgement by eye was critical. For the TSG-250, the Au quality finally achieved is now judged to be similar to that of the ECF64D used in the final electroplating process.

Run	P	$\rho\,[^{mA}/_{\mathrm{cm}^2}]$	$T\,[°\mathrm{C}]$	t [hrs]	ν [rpm]	l_s	Sol.
1	yes	4.3	$60\pm^2_5$	2.	≈1100	s	TSG-250
2	yes	4.3	$60\pm^2_5$	2.	≈700	s	TSG-250
3	no	2.19	$60\pm^2_5$	3.	≈ 1100	s	TSG-250
4	no	2.19	$60\pm^2_5$	3.	≈ 400	l	TSG-250
5	no	2.19	$60\pm^2_5$	1.5	400	m	TSG-250
6	no	2.19	60 ± 1	1.5	400	m	TSG-250
7	no	2.19	50 ± 1	1.5	400	m	TSG-250
8	no	2.	55 ± 1	1.5	400	m	TSG-250
9	no	2.	60 ± 1	1.5	400	m	TSG-250
10	no	2.	60.0 ± 0.1	1.	400	m	TSG-250
11	no				400	m	TSG-250
12	no				400	m	TSG-250
13	no	2.19	60.0 ± 0.1	1.	400	m	TSG-250
14	no				400	m	TSG-250
15	no	2.19	60.0 ± 0.1	1.	400	m	TSG-250
16	no	1.10	50.0 ± 1	0.83	1600	m	TSG-250
17	no	0.55	50.0 ± 1	0.83	400	m	TSG-250
18	no	2.19	60.0 ± 0.4	1.0	400	m	TSG-250
19	yes	3.92	60.0 ± 0.1	1.0	400	m	TSG-250
20	yes	4.26	60.0 ± 0.5	1.0	400	m	TSG-250
21	yes	3.84	60.0 ± 0.5	1.0	400	m	TSG-250
22	yes	1.96	60.0 ± 0.2	1.0	400	m	TSG-250
23	no	2.00	60.0 ± 0.1	1.0	400	m	TSG-250
24	no	1.54	50.0 ± 0.3	0.7	400	m	ECF 64D
25	no	2.92	50.0 ± 0.6	0.7	400	m	ECF 64D
26	no	1.84	60.0 ± 0.8	1.12	400	m	Gold-SF
27	no	3.00	49.3 ± 0.1	0.35	400	m	ECF 64D
28	no				400	m	ECF 64D
29	no				400	m	ECF 64D
30	no				400	m	ECF 64D
31	yes	1.5	50.0 ± 0.1	0.83	400	m	ECF 64D
32	yes	5.0	50.0 ± 0.1	0.21	400	m	ECF 64D
33	yes	3.0	50.0 ± 0.1	0.35	400	m	ECF 64D
34	yes	0.75	50.0 ± 0.1	1.40	400	m	ECF 64D
35	yes	2.	50.0 ± 0.1	0.53	400	m	ECF 64D
36	yes	1.5	50.0 ± 0.1	0.70	400	m	ECF 60
37	yes	2.	55.0 ± 0.1	0.53	1600	m	ECF 60
38	no	2.16	60.0 ± 0.1	0.53	400	m	TSG-250

Table 7.1: An overview of all 38 process development wafer, whether they were were patterned P, with the current density ρ, the temperature of the electroplating solution T, the electroplating time t, the stirring rotational frequency ν, the size of the stirrer l_s, and the type of the electroplating solution. The wafers without electroplating details were used for SEM images of the seed layers, as a comparison to the SEM images of electroplated wafers.

7.2 PCBs Version 2

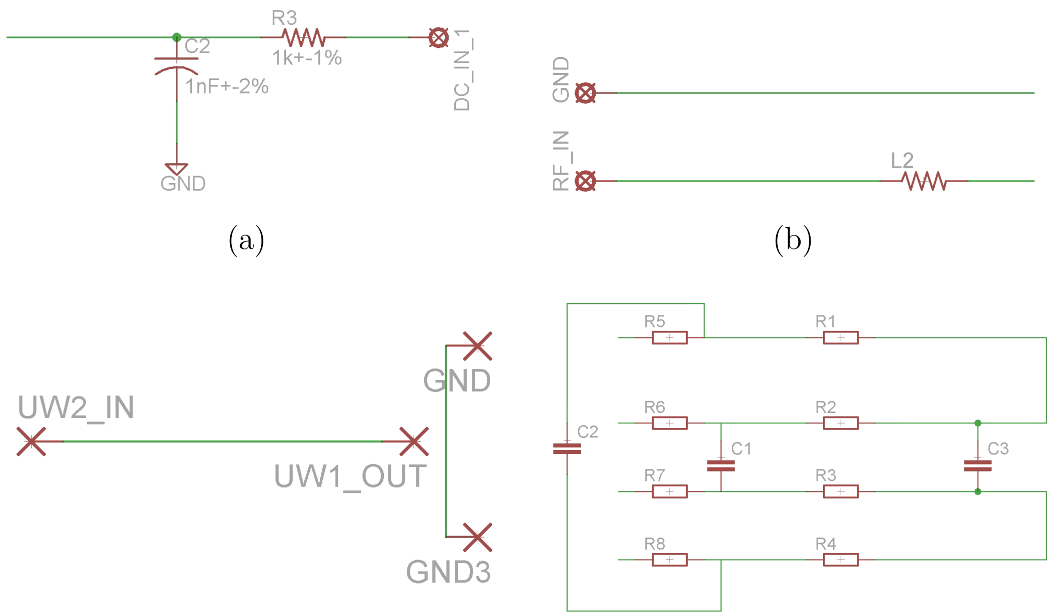

(a) (b)

(c) (d)

Figure 7.3: Schematic overviews of the in-vacuum PCBs. (a) For the DC PCB, all 11 channels have the same schematic, and only one channel is shown. (b, c) The ground pads on the right are the pads used to wire bond the grounding onto. (c) The μW PCB features a spare μW-line, which was initially needed for some gate schemes under consideration, but the final trap design and corresponding gate scheme only needs one μW-line. Created with EAGLE 6.6.0.

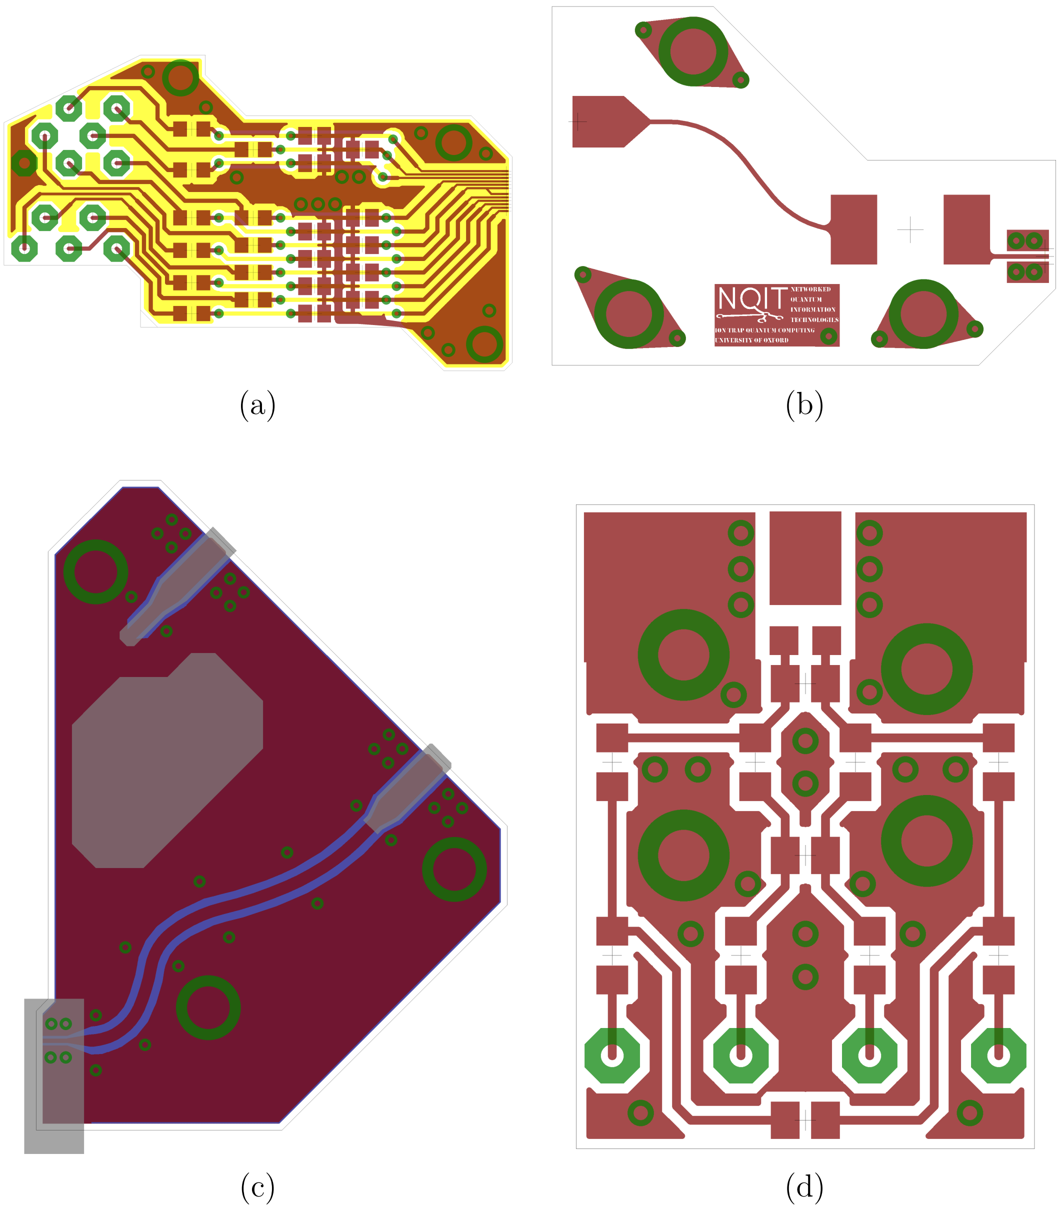

(a) (b) (c) (d)

Figure 7.4: Routing overviews of the in-vacuum PCBs. The layers are: red (top conductive layer), blue (bottom conductive layer), and green (outlines of vias). For photographs of these PCBs on the pillbox, see fig. 4.19c. Created with EAGLE 6.6.0.

7.3 Argon Ion Cleaning

The ion sputtering gun is an IG70 Ionec by OCI Vacuum Microengineering Inc. It has a beam energy of up to 3 keV, a spot size of $5 - 20$ mm. Both can be remote controlled via a IPS3-D digital controller, with built-in USB capabilities. Furthermore, the beam can be diverted with X-Y deflection plates. They are specified to create a diversion of up to± 1 mm, at the working distance between the end of the gun and the sample of 100 mm.

In order to use a modular approach to the Argon Ion cleaning, it is mounted on a multiport CF100 flange, part 409005 from MDC Vacuum Products LLC, which includes six angled CF16 ports, and one CF40 port, the latter of which is used to mount the gun. One of the additional CF16 ports is used for a two-pin electrical feedthrough - part 9412011 from MDC Vacuum Products LLC - one of which is spare. The other pin is electrically connected to a specially designed pillbox lid. This lid is electrically insulated from the rest of the pillbox, and features a smaller opening in the gun's line of sight to the ion trapping chip. The purpose of this design is to first direct most of the argon ion beam into the lid, and to calibrate the current this way. Subsequently, the beam is centred on the chip, through a Ø3 mm hole. By increasing the spot size to 10 mm, the Gaussian profile of the beam is close to a flat-top profile. This has the advantage that all relevant areas on the trap are being cleaned at the same rate. Argon is leaked into the vacuum chamber via a precision leak valve, in order to create an operating pressure of $5 \cdot 10^{-6}$ Torr.

The Argon is supplied by BOC Ltd., and has a specified purity of $> 99.9995\%$. The pressure is reduced with a HP1500 stainless steel single-stage regulator, also from BOC Ltd. It is connected to the leak valve via an internally electropolished SS tube, part RST0081_3000 by Wellington Tube Supplies Ltd.

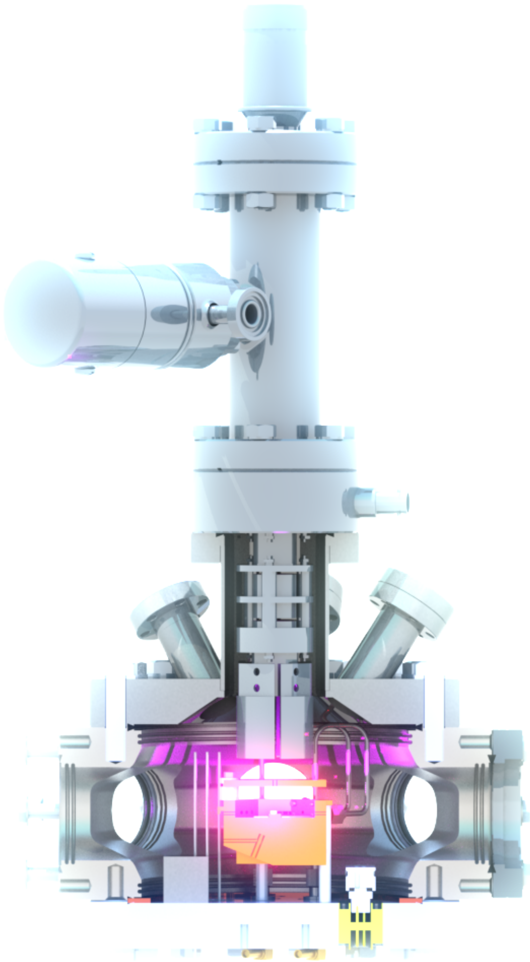

Figure 7.5: Argon ion sputter mounted on the vacuum chamber. Rendered with Solidworks Professional 2016.

7.4 Pressure Drop Rule Of Thumb Derivation

The rule of thumb introduced in sect 4.4.9 is derived here: If pressure drops by $< 10\,\%/\mathrm{day}$, the bake can be stopped, as there is little extra pressure reduction to be gained. Starting with equ. 4.11,

$$P = a \cdot e^{-b \cdot t} + c$$

the pressures on two days, P_1 and P_2, are introduced. They are take at different times, $t_2 = t_1 + \Delta t$, and their pressures are related by $P_2 = \beta P_1$. The pressure drop still to be gained is defined as $d = \frac{P_2 - c}{c}$. Assuming the constant b is known, the drop $d\,(a, c)$ can be derived. Expanding the pressure relation $P_2 = \beta P_1$ gives

$$a \cdot e^{-bt_2} + c = \beta \left(a \cdot e^{-bt_1} + c \right)$$

Solving this equation for a yields

$$ae^{-bt_2} = \frac{c\,(1 - \beta)}{\beta \cdot e^{b \cdot \Delta t} - 1}$$

Therefore,

$$P_2 = ae^{-bt_2} + c = \frac{c\,(1 - \beta)}{\beta \cdot e^{b \cdot \Delta t} - 1} + c$$

This is then inserted into the definition of the pressure drop $d = \frac{P_2 - c}{c}$,

$$d = \frac{1}{c} \left(\frac{c\,(1 - \beta)}{\beta \cdot e^{b \cdot \Delta t} - 1} + c - c \right) = \frac{1 - \beta}{\beta \cdot e^{b \cdot \Delta t} - 1}$$

Taking the constant $b = 0.51\,{}^{1}/_{\mathrm{day}}$ as typical, and inserting the rule of thumb values, $\beta = 0.9$ and $\Delta t = 1$ day, this yields $d = \frac{1 - \beta}{\beta \cdot e^{b \cdot \Delta t} - 1} = \frac{1 - 0.9}{0.9 \cdot \exp(0.51 \cdot 1) - 1} = 20\,\%$.

This is sufficiently low that the bake can be stopped.