A scalable demonstrator for trapped-ion quantum computing using modules connected by electric fields

Nicholas Ian Johnson

Submitted for the degree of Doctor of Philosophy

University of Sussex, Brighton, United Kingdom.

October 2021

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

Signature:

Nicholas Johnson

Abstract

This thesis describes the development of a modular demonstrator device for scalable quantum computing using ion traps connected by electric fields. The experiment was configured to control the positions of two independent surface-electrode ion trap modules, whose electrodes when aligned create a potential suitable for ion transport between the modules. In principle this method can be used to construct a modular ion-trap processor of arbitrary size. An analysis of the requirements for reliable, low-loss transfer of ion qubits between the modules is presented. The experiment was used to trap ^{174}Yb ions and to transport ions between modules as a proof-of-principle quantum link. The ion trap geometry incorporates a junction to enable reordering and transport of ions between spatially separated zones, and the experiment is designed for operating microwave-driven quantum logic gates using a magnetic field gradient. To increase the maximum attainable current for on-chip magnetic gradient coils, a scalable cryogenic cooling system was developed with high cooling power and an extensible design. The cooling system was connected to the demonstrator experiment and two additional ion trap experiments to prove its capability to scale in size. The minimum temperature achieved was 40 K under no active heat load, and 70 K for a heat load of 111 W. This permits large gradients for high-fidelity logic gates, an order of magnitude reduction in ion heating rate, and lowers vacuum chamber pressure for increased ion lifetimes. The demonstrator experiment is a step towards the realisation of a large-scale quantum computer based on trapped ions.

Acknowledgements

I first want to thank Winfried Hensinger for giving me the opportunity to study for a PhD at Sussex. I must also thank Nathan Chong who encouraged me to apply to be a doctoral student in the first place. Over the last four years, I have received generous help and support from every member of the Ion Quantum Technology research group, but I am especially grateful to Raphaël, Tomas, Anton, Sam and David who have always been patient and willing to spare their time to talk me through things. A huge thanks also to Reuben for all the tea and therapy when times were tough. To Harry, you have been a kind and generous friend, and a humble and patient landlord! To Zak, my fellow Northerner, you were the only one who actually enjoyed (or understood?) my meandering stories. To Mitch, we have had the best kinds of conversations about philosophy and life, and I look forward to many more...and I am still waiting for those Chaos Space Marines! To Raphaël, thank you for your dedication and commitment, we have come a long way in such a short time and it has been a pleasure serving with you! To Mariam and Falk, I remain totally in awe of what you have achieved in just a few short months. I must give a heartfelt shout out to Mum and Pops, and Chris and Li Min, who have supplied much love and support during my research, and were always rooting for me to succeed. Finally, I have been extremely lucky to have the unshakeable support of my husband, Tianyang, who has cooked me sandwiches and taught me how to persevere and to believe in myself.

Publications

A scalable helium gas cooling system for trapped-ion applications

F. R. Lebrun-Gallagher, N. Johnson, M. Akhtar, S. Weidt, D. Bretaud, S. J. Hile, A. Owens, F. Bonus and W. K. Hensinger

Quantum Sci. Technol.

(Submitted Dec 2021)

A matter link for quantum computing modules

M. Akhtar, F. Bonus, N. Johnson, F. R. Lebrun-Gallagher, M. Siegele, S. Hong, S. Weidt, S. J. Hile, S. A. Kulmiya and W. K. Hensinger

(In preparation)

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The advent of the electronic digital computer in the middle of the 20th Century, and the theoretical framework that emerged alongside it, brought about a transformation of society that is impossible to overstate. Today, its impact reaches into every aspect of the modern world. An enormous variety of applications have been made possible by the technologies developed for information processing, including global telecommunications, the internet, international financial trading, advanced science and engineering research – the list goes on. The original motivations behind the construction of the first generation of computers were concerned with answering the most urgent practical problems of the time: the Colossus at Bletchley Park was built to assist with the cryptanalysis of the Lorenz cipher during World War II [1]; the ENIAC was designed in 1945 to calculate trajectories of artillery munitions for the United States Army [2]. A century earlier, Charles Babbage had invented mechanical computers to automate the generation of mathematical reference tables, in an effort to avoid human errors. His Analytical Engine, designed in 1834 but never built, was inspired by the mechanisms of the textile looms used in the factories of the industrial revolution. It was digital and programmable, with a system based on punched cards for reading in data and controlling programs [3]. The early pioneers of computer science could not have foreseen the far-reaching consequences of the techniques and technologies they were creating. The same is almost certain to be true of the quantum technologies being developed now.

The last two decades have seen rapid and encouraging progress in quantum computing research. The elementary building blocks of quantum computation, including quantum logic gates and even small quantum algorithms, have been demonstrated using a variety of different physical implementations. A number of national funding initiatives have been established, such as the United Kingdom's National Quantum Technologies programme [4],

the European Union Quantum Flagship [5] and the National Quantum Initiative in the US [6]. Possibly the most encouraging signal of the pace of recent development is the interest shown by private enterprises such as Google, IBM, Microsoft and others, who now have their own quantum computing research programmes [7, 8, 9]. Academic research groups have given rise to start-up companies such as Ion-Q and Quantum Circuits Inc. in the United States, Q-CTRL in Australia, and Riverlane in the UK, seeking to commercialise quantum computing technologies. This is notable because while universities and government research often spearhead the early breakthroughs in science, only industry can take forward applications to real world problems of interest, and at scale. Applications are being explored in a wide range of fields such as machine learning [10], computational finance [11] and simulations of molecular interactions [12, 13] for medical and materials research. Implementations using trapped ions and superconducting circuits are arguably the forerunners in the race to develop a scalable architecture for building a universal device – one capable of performing any quantum algorithm. In many respects, the present experimental landscape in quantum information science bears a remarkable similarity to the early years of conventional computer development, when implementations were laboratory-based and room-sized. Today we are beginning to see these experiments make the transition from proof-of-principle devices to nascent commercial technologies, with the potential to again revolutionise the world we live in. And we might expect that, like the first computers, the most impactful uses for quantum computers in the future have yet to be predicted.

In this chapter I will explain the basic concepts of classical and quantum information processing, followed by a brief summary of the physical requirements for quantum computation to be implemented in real systems. I provide an overview of research into designs for realising large-scale quantum computing devices with trapped ions, and introduce the scalable demonstrator experiment that was developed during my PhD research.

1.1 Classical information processing

The theoretical basis of digital computation is independent of any particular hardware implementation, and requires that information be encoded in an abstract two-level construct called a binary digit ('bit'), whose distinct states can be labelled as 0 and 1. This formulation is entirely classical and thus the digital model of computation is constrained to operate by the laws of classical physics. The data encoded in the form of bits are processed using a sequence of logic gate operations. Logic gates are abstract mathematical entities that represent the basic Boolean functions of conditional logic, and some examples are shown

in Figure 1.1. A logic gate is universal if it can be applied successively and exclusively to replicate the output of any other logic gate. Both NAND and NOR gates have this property. By building circuits using just one type of universal gate it is possible to execute any finite set of computational steps.

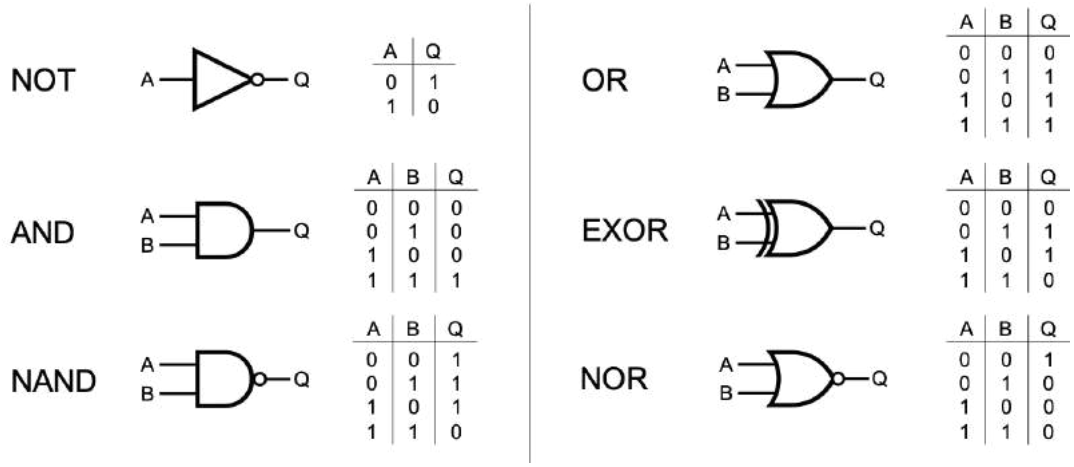

Figure 1.1: Logic gates for digital information processing. The NAND and NOR gates are *universal*, as they can be used to construct any others.

One advantage of digital computation using information encoded in discrete values, compared to analogue computation using continuous variables, is the relative simplicity of performing error correction. The occurrence of erroneous bit flips (from 0 to 1, or vice versa) due to the presence of noise can be more readily detected and tracked through a circuit compared to continuous analogue drifts that can take any value. Error correction on classical bits is implemented using redundancy by storing multiple duplicate values of a bit and taking forward the majority result, assuming that errors are sparse [14]. This method cannot be used for quantum bits owing to the no-cloning theorem [14] which prohibits the copying of an unknown quantum state, and so a different approach to error correction must be employed which carries a considerable computational overhead (see section 1.2.1).

The second half of the last century saw developments in digital computing technology accelerate at a staggering pace, driven by the miniaturisation of the silicon transistor. The transistor is the most fundamental logic element within a computer processor, and it is used to encode a single bit (stored as a potential difference across the device). Since the early 1960s, the number of transistors per integrated circuit has doubled every 18 months or so. This exponential growth in processing power and complexity was famously named 'Moore's Law' after Gordon Moore, Chief Executive Officer of Intel, first observed the trend in 1965 and we are only now starting to see the end of its impressive run. A typical modern computer processor contains a few billion transistors. High performance

supercomputers for scientific and industrial applications operate using many thousands of processors running in parallel. At the time of writing, the world's most powerful computer is the Fugaku supercomputer [15] developed by Fujitsu and installed at the RIKEN Center for Computational Science in Japan. It can achieve a peak performance of 537 petaflops (1 petaflop = 10^{15} floating point operations per second).

Although the rate of progress in classical digital computing continues to be impressive, nevertheless there are problems of scientific and commercial interest that remain computationally hard. These are problems with exponential computational complexity, and therefore the number of computational steps required to reach a solution grows exponentially with the input size. Examples of this type of problem include finding the prime factors of an integer, protein folding calculations and estimations of molecular energy levels. Nature itself is not classical but quantum, and the best-known classical algorithms for calculating quantum many-body dynamics exhibit exponential computational complexity. As a result, even the most powerful supercomputers are unable to calculate, in any reasonable amount of time, exact solutions of the Schrödinger equation for an entangled system of N spin-$\frac{1}{2}$ particles, when $N{>}50$ or so [16, 17]. The physics of quantum systems present a new opportunity for a different approach to computation intended to purposefully exploit phenomena that are quantum mechanical in nature.

1.2 Quantum information processing

The theoretical advantages of using the laws of quantum mechanics as a basis for computation were first explored in the 1980s. In 1982, Richard Feynman proposed that a quantum computer, based on the properties of a suitable quantum mechanical system, could be used as a universal digital simulator, capable of representing any other quantum system, and that any classical computer would require exponential resources to do the same [18]. The central idea was to replicate an interaction Hamiltonian by applying a set of operators to a spin-$\frac{1}{2}$ lattice, where each site in the lattice is a two-level system that can encode a quantum bit (qubit) of information. A qubit is defined by the distinct eigenstates $|0\rangle$ and $|1\rangle$, which form an orthornormal basis. In accordance with the Schrödinger equation, the general quantum state $|\psi\rangle$ of a qubit can be described as a linear superposition of the eigenstates,

$$|\psi\rangle = a\,|0\rangle + b\,|1\rangle\,, \tag{1.1}$$

where a and b are in general complex coefficients that may be treated as probability amplitudes, subject to the normalization condition, $|a|^2 + |b|^2 = 1$. This reflects the property that, in any single measurement, the qubit will always be found to exist in one of the two possible eigenstates $|0\rangle$ or $|1\rangle$. As a result of the normalization condition, equation 1.1 can be re-written as

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle\,. \tag{1.2}$$

The complete state space of a qubit can be described by the 2-dimensional surface of a unit sphere, on which the north and south poles represent the $|0\rangle$ and $|1\rangle$ eigenstates, respectively. This representation is shown in Figure 1.2, and is known as the Bloch sphere.

An ideal qubit would be completely isolated from its environment, but also able to be manipulated with high precision by an external means of control. A system of N qubits has a state space of 2^N dimensions, defined by the tensor product of the individual two-dimensional qubit subspaces. This is referred to as Hilbert space. The computational advantage, or quantum speed-up, made possible by the laws of quantum mechanics stems from the unique properties of Hilbert space that permit non-classical correlations among qubits. In particular, superposition and entanglement have turned out to be powerful physical resources for computation that can, in some cases, be exploited to construct more

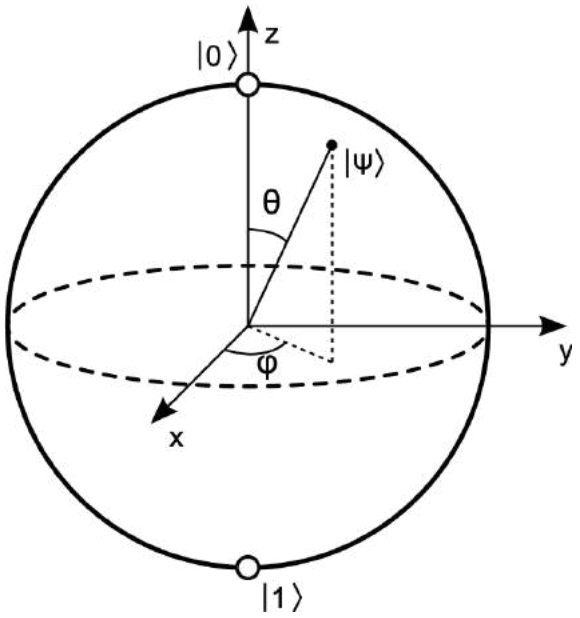

Figure 1.2: The Bloch sphere representation of a qubit. A quantum state can be parameterised by the angles θ and ϕ (plus a global phase that has no observable effect) [14]. The eigenstates $|0\rangle$ and $|1\rangle$ are located at the north and south poles of the sphere, respectively. Image reproduced from [19].

efficient algorithms than is permitted by classical physics. The price of this speed-up is paid for by the exquisite levels of physical isolation and control needed to preserve the fragile quantum coherence of the qubits being manipulated.

1.2.1 Early discoveries

One of the first quantum algorithms was formulated by David Deutsch in 1985 [20] and later generalised in collaboration with Richard Jozsa in 1992 [21]. The Deutsch-Jozsa algorithm solves a specific problem concerning the classification of a function by analysing its output, and does so in exponentially fewer steps than can be achieved by any classical algorithm. This was an important example of a quantum speed-up, although the problem itself does not have any known practical application. Deutsch had previously shown that any unitary operation on N qubits can be decomposed into specific two-qubit and single-qubit operations performed in a sequence [22]. This is referred to as the circuit model of quantum computing because it is a generalisation of the digital logic circuits in classical computing. One example of a quantum logic gate is the controlled-NOT (CNOT) gate that operates on two qubits,

$$U_{\mathrm{CNOT}} = |00\rangle\langle 00| + |01\rangle\langle 01| + |10\rangle\langle 11| + |11\rangle\langle 10|\,. \tag{1.3}$$

The CNOT can be understood to work by taking the first qubit as a control condition for deciding whether or not to perform the state inversion of the second qubit. The CNOT gate combined with single-qubit gates forms a universal set of operations for quantum computation.

In 1994, Peter Shor published a quantum algorithm for computing the prime factors of an integer with a (nearly) exponential speed-up compared to the best-known classical algorithm [23]. This was a problem of immense practical significance owing to the importance of prime factorisation for cryptography and information security. In 1996, Lov Grover published a quantum algorithm capable of searching unstructured datasets using $O(\sqrt{N})$ steps compared to classical algorithms that require $O(N)$ steps, a quadratic speed-up [24]. Around the same time, quantum error correcting codes were being developed by Peter Shor [25], Andrew Steane [26], John Preskill [27] and others. These methods showed that multiple physical qubits can be used to encode the quantum state of a 'logical qubit' that is resilient to the effects of noise. By introducing additional 'ancilla' qubits and quantum logic gate operations, it was shown that errors up to a limiting threshold can be detected and corrected. This was an important discovery because real physical systems suffer from

inaccuracies in the parameters used to characterise the state of a qubit, owing to imperfect initialisation, couplings to environmental noise, and errors that accumulate along the chain of gate operations. One of the first experimental realizations of a quantum logic gate between qubits was implemented in 1995 by David Wineland's research group at the US National Institute for Standards and Technology, using trapped Be^+ ions to demonstrate a CNOT gate [28]. The experiment put into practice theory developed by Cirac and Zoller earlier that same year [29].

The importance of these results for computer science, mathematics and physics led to an explosion of interest in quantum computing. Research efforts intensified to find further examples of quantum algorithms, alongside increasing experimental efforts to develop physical implementations.

1.2.2 The DiVincenzo criteria

The fundamental physical requirements for quantum computation were summarised by David DiVincenzo in 2000 [30], who identified that any implementation must satisfy the following criteria:

1. A scalable physical system with well characterised qubits,
2. The ability to initialise the state of the qubits to a simple fiducial state, such as $|000\rangle$,
3. Long relevant coherence times, much longer than the gate operation time,
4. A universal set of quantum gates,
5. A qubit-specific measurement capability,
6. The ability to interconvert stationary and flying qubits,
7. The ability to faithfully transmit flying qubits between specified locations.

When selecting (or engineering) a physical system to use as the basis for quantum computing, we need a supply of qubits that are well isolated from their environment, are individually addressable, and which permit the initialisation, manipulation and readout of the qubit state. The coherence times must be long enough to permit gate operations to be performed, and ideally many gates in succession, before errors due to decoherence become too large. To be capable of running any algorithm, including the necessary overheads of error-correcting codes, the quantum system must be scalable to arbitrary numbers of

qubits. The final two criteria on the list above (points 6 and 7) relate to *quantum communication* between qubits, and they are necessary conditions for realising a quantum network, which is a key requirement for one of the proposed scalable architectures discussed in section 1.5. In the years since the DiVincenzo criteria were identified, various combinations of the individual requirements have been demonstrated in a range of different physical systems including semiconductor qubits [31], nitrogen-vacancy defects in diamond [32], superconducting circuits [33] and trapped ions [34].

1.3 Near-term devices

Current high-profile devices such as the IBM Q System One, which offers 20 qubits accessible via the cloud, and the 72-qubit Bristlecone quantum processor being developed by Google, suffer from relatively noisy qubits and do not implement quantum error correction. This limits the maximum number of useful gate operations that can be carried out one-after-another (often referred to as the circuit depth), because errors accumulate along the chain of computations. This prohibits the size and complexity of quantum algorithms that can be run by such devices. Therefore, these systems cannot be considered scalable but are nevertheless useful steps towards larger-scale quantum computing, and which can be used to explore the bounds of what is possible with current technology. Recently, hybrid quantum-classical techniques for optimisation using variational quantum eigensolvers [35] have been used to calculate molecular energy levels, using hardware developed by IonQ [36] and Rigetti [37]. An alternative approach to optimisation that does not use the circuit model of quantum computing is adiabatic quantum annealing [38], which is the method employed by D-Wave. This technique is a form of analogue simulation that attempts to encode a specific difficult-to-solve problem into the many-qubit state of the annealer, and then evolves the system to its ground state which represents an optimal solution to the problem under study. D-Wave reports control of 2048 qubits in its 2000Q system, although it is unclear how much of a quantum speed-up can be achieved using noisy qubits for annealing or even how any such speed-up might be appropriately measured, but this remains an active area of research [39]. It is likely that, in the near term, devices such as those described above will contribute to our understanding of the value of low-depth quantum algorithms using imperfect qubits and what advantage, if any, might be achievable for computations without error correcting measures [40].

1.4 Trapped ions

Trapped ions are one of the most promising technologies for quantum computing because they have demonstrated all of the fundamental requirements identified by Divincenzo for universal quantum information processing [30]. Single-qubit and two-qubit logic gates have been demonstrated with fidelities of 99.9999 % [41] and 99.7 % [42], respectively, with gate times of the order of a few hundred microseconds. Coherence times of over 10 minutes have been achieved in trapped-ion qubits [43], making them ideal for use as long-lived quantum memories. Ions within a trap interact via their mutual Coulomb repulsion, and their collective motional modes can be exploited to mediate entanglement between ion qubits within a one-dimensional chain [29, 44]. This method has been used to generate a fully entangled Greenberger-Horne-Zeilinger state of 14 qubits [45], and a similar implementation has demonstrated an exquisite level of control over 20 trapped ion qubits with multipartite entanglement observed within constituent sub-chains [46]. Reducing the motional excitation of ion qubits improves multi-qubit gate fidelity by reducing heating of collective modes [44], and methods have been developed to cool ions to their ground state of motion [47, 48]. Sympathetic cooling of one ion by another has been implemented using ions of different atomic species [49, 50]. This method applies Doppler cooling to a dedicated cooling ion species within the same potential well as the data ion, to avoid disrupting the qubit state of the data ion during logical operations.

Recent developments in quantum error correcting codes have shown that quantum logic operations can be sustained with a per-operation error threshold close to 1 % using only nearest neighbour single-qubit and two-qubit gates between physical qubits arranged in a two-dimensional configuration [51]. Implementation of these surface codes [52], as they are known, may be achieved by transporting ions through a two-dimensional trap [34], as described in the following section. The number of physical qubits required to encode a logical qubit depends strongly on the error rate per physical qubit, with fewer necessary for lower error rates. As an example, to encode a single fault-tolerant logical qubit it may take in the order of 10^4 ions based on the assumption of an error rate close to 0.1 % (a fidelity of 99.9 %) [52]. It follows that practical algorithms involving error correcting codes may need to use millions of trapped ions to solve real-world problems. The outstanding challenge is to find a scalable approach to quantum processing that will permit reliable and accurate control over many ions in parallel.

1.5 Scalable quantum computing with trapped-ion qubits

A number of proposals have been developed for constructing a large-scale, universal quantum computer based on ion traps. These architectures provide logical connectivity between all qubits [30] and use modularity to address the scalability challenge.

Monroe *et al* [53, 54] have proposed a scalable architecture based on a quantum network of distributed ion-trap processors. In this approach, microfabricated ion-trap modules are coupled via photonic interconnects. This requires an efficient method to convert quantum information from ion qubits to photons and back, with the photons acting as the flying qubits described by DiVincenzo (see section 1.2.2). Entanglement between ions within a module is performed using deterministic phonon-mediated interactions, whereas probabilistic entanglement between modules is achieved by interference of fluorescence photons from each module [54]. The attractiveness of this approach is that optical fibres may in principle be used to connect over long distances in the same manner that conventional telecommunication networks operate today. The challenges of this approach arise from the low efficiency with which fluorescence photons can be coupled into the optical fibres and the high absorption losses within the fibres at the UV wavelengths of the emitted photons which limits the length of the fibre link and results in low entanglement rates. Losses may be reduced by quantum frequency conversion to increase the photon wavelength, which was recently demonstrated to permit photons from a calcium ion to be transmitted over 10 km through a fibre [55]. At present, entanglement rates between modules are of the order of a few tens of Hz in practice [56].

Kielpinski *et al* have described a scalable architecture based on the idea of a quantum charge-coupled device (QCCD) in which ions are transported between interaction zones within a two-dimensional trap array [57]. Logic gates are performed in parallel on small numbers of ions at specific locations of the trap. This avoids the limitations of motional mode crowding in large ion chains [58], and the weakening of interactions between ions located increasingly far apart [46]. The trap geometry incorporates junctions through which ions can be transported and re-ordered. Other zones permit ion loading and qubit state readout.

The fundamental ion transport operations required by this architecture include shuttling of ions along a line between adjacent zones, transport through junctions, separation of two or more ions into individual potential wells, and recombination (or merging) into a single potential well. These operations have all been previously demonstrated by a number of research groups. Shuttling of ions within a linear trap [59], through a T-junction [60] and

X-junction [61, 62] were first achieved using individually-constructed multi-substrate ion traps. However, microfabricated ion traps [63] are a more attractive option for scalability since many identical copies of the device can be manufactured accurately. In particular, surface-electrode traps [64] have been developed in which the rf and control electrodes all lie within a single layer, thus simplifying their construction. Ion transport through a surface electrode X-junction trap has been demonstrated with less than a single phonon of motional excitation [65]. In general, ion transport can be achieved by applying time-varying voltage waveforms to segmented control electrodes to translate a potential well along the RF nill of a trap (a more detailed description of this procedure is provided in Chapter 7). Ion transport may be achieved adiabatically with minimum excitation of the ion's motion (as in [62, 65]) as required for high-fidelity gates, so long as the time taken for transport is much greater than the period of the ion's secular oscillation within the well. However, it is desirable to achieve fast, non-adiabatic ('diabatic') ion transport to minimise the time spent moving ions between locations and maximise the attainable clock speed of the quantum processor. In this case a method must be employed to remove the energy gained during transport, such as by managing the forces on the ion due to the transport potential [66] or by sympathetic cooling. Diabatic transport of an ion has been demonstrated over a distance of 280 µm in 3.6 µs - just five motional oscillations within the trap - and with an energy increase of 0.1 motional quanta [66]. Diabatic transport of a mixed-species ion chain near the ground state of motion has also been achieved [67]. Fast diabatic separation of a chain of ions into two distinct potential wells has been demonstrated, and in the case of a two ion chain this was achieved in 55 µs [68]. Diabatic ion swapping, in which the positions of neighbouring ions are interchanged by rotation, has been shown [69, 70], and re-ordering of a pair of ions has also been performed by sending one ion into a junction to allow the other to overtake [71]). In a recent demonstration of the QCCD-type approach, a combination of the fundamental operations described above were implemented within a multi-zone linear trap using $^{171}Yb^{+}$ qubits sympathetically cooled by $^{138}Ba^{+}$ ions [72].

Lekitsch *et al* [34] have developed the QCCD concept further by designing an architecture based on ion-trap modules connected by electric fields. A simplified illustration of one such module is shown in Figure 1.3. The top surface of each module is an array of interconnected X-junction ion traps, with integrated electronics below the surface to provide control of trapping potentials. Using this approach, a large-scale device can be constructed by physically aligning closely-spaced modules to form a tiled array of ion traps. The trap potential provides a path for ion transport between adjacent modules.

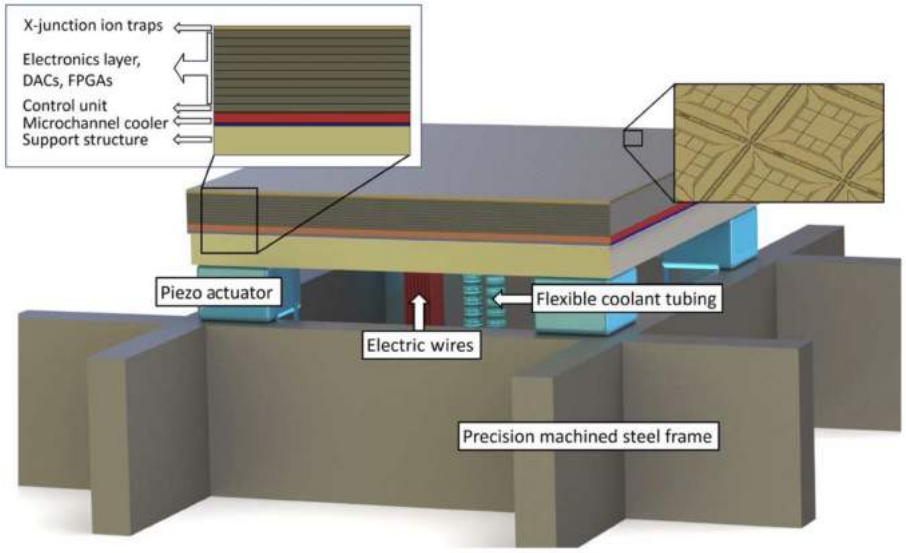

Figure 1.3: Lekitsch *et al* [34] propose a scalable architecture for quantum computing based on ion trap modules connected by electric fields. A single module incorporates microfabricated ion trap chips with integrated electronics for shaping trapping potentials, piezo actuators for chip alignment, and a heat sink for cryogenic cooling.

An additional fundamental building block for ion transport is needed for this architecture, namely the reliable shuttling of ions between separate microchip ion traps. In this way, quantum information is distributed by ions moved through the trap array, and local phonon-mediated interactions occur between ions shuttled over centimetre distances. Entanglement rates depend on the secular frequencies of ions within the trap, which are of order 1 MHz, and on the speed of ion transport. The structure can in principle scale up to arbitrary size by connecting together many modules to support very large numbers of qubits.

Another key feature of this approach is the use of long-wavelength microwave and radio frequency radiation in combination with an applied magnetic field gradient [74] to drive Mølmer-Sørensen type gates [44] between ions. This is made possible by preparing qubits using two hyperfine states of $^{171}Yb^{+}$, which have different magnetic moments so that in the presence of a magnetic field gradient there exists a state-dependent force on the ion that imparts sufficient momentum to couple transitions between the internal and motional states of the ions, as required by [44]. Integrated current wires below the trap surface generate a large magnetic field gradient at specific zones of the trap. In addition to quantum gates, the magnetic gradient scheme enables individual ions to be addressed by means of the Zeeman splitting between the qubit states, which varies with the magnitude of the field at different positions within the trap. Figure 1.4 shows how ions may be positioned at three different zones (blue, green and red) where the qubit transition frequencies are tuned into or out of

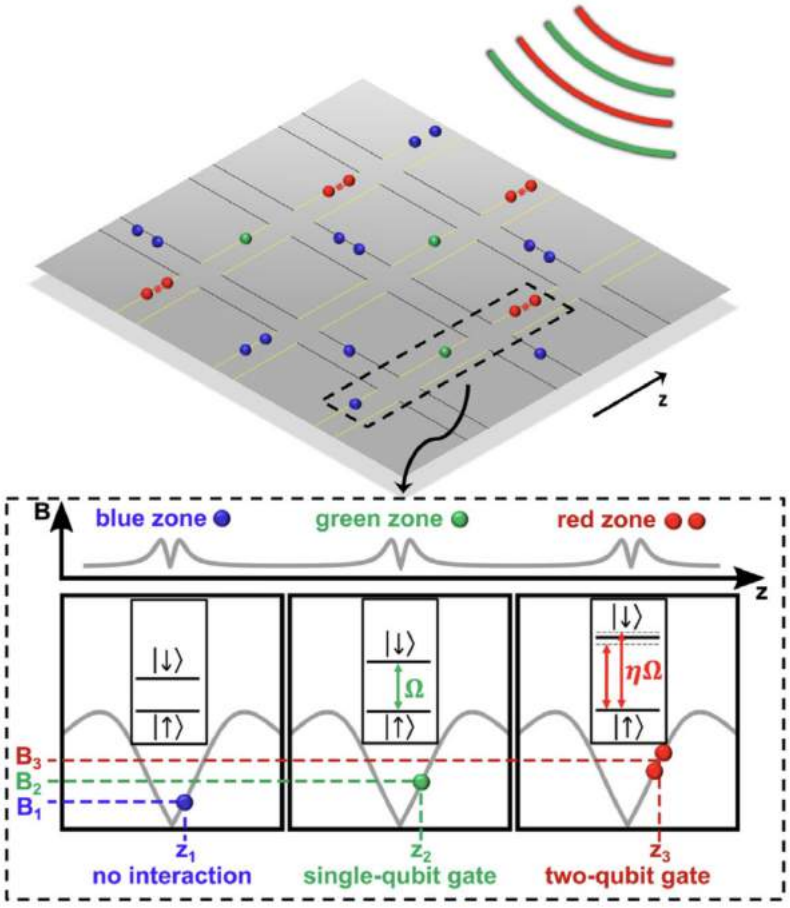

Figure 1.4: The approach described by Lekitsch *et al* [34] uses globally applied microwave fields to drive gates and address ions in a two-dimensional trap array. The top image shows the trap geometry. The bottom image shows the non-uniform magnetic field, B (grey line), which varies with position along a linear section of the trap (z axis). Control electrodes adjust the trapping potential to position ions within the field, therefore tuning qubit transitions into or out of resonance with applied microwave fields via the Zeeman effect. For ions in the blue zone, qubit transition frequencies are off resonance so no interaction occurs, whilst the green and red zones can be chosen to operate single- and two-qubit gates respectively. Image reproduced from [73].

resonance with globally applied microwave fields. This permits the scheme to replace large numbers of precisely-aligned Raman lasers with only a handful of global fields, drastically simplifying the engineering overhead of building a many-qubit device. Recently, the group has implemented a two-qubit gate using the above scheme, with a high fidelity in line with laser-based approaches [73]. The two-dimensional ion trap architecture is also compatible with quantum error correction, using a surface code that relies only on nearest neighbour interactions between ions [52].

This is the vision being pursued by the Ion Quantum Technology research group at Sussex, and my PhD project was to construct a scalable demonstrator experiment as a step towards this long-term goal. The project is described in the next section.

1.6 The scalable demonstrator experiment

The aim of the work described in this thesis was to construct a demonstrator experiment to investigate the requirements for connecting two independent surface-trap microchips using electric fields, and to provide all of the necessary elements to operate the device as a two-module quantum processor using the microwave gate scheme detailed above. As this was a new experiment nearly all components of the experimental set-up needed to be considered and developed[1] (responsibility for some components of the demonstrator project lay with other members of the research group, and these are detailed in the contributions section at the end of this chapter).

The component sub-systems that make up the scalable demonstrator experiment are as follows:

1. Surface-electrode ion trap chips featuring a precisely fabricated edge (referred to as 'edge-traps'), with electrodes that extend to the edge to permit close alignment between chips

2. A piezo positioning system for alignment of two edge-traps to form a path for reliable ion transport between them

3. Electronics for delivering radio frequency (RF) and static control potentials to the ion trap electrodes, including time varying shuttling potentials, with filtering to reduce noise near the ion secular frequency

4. Atomic sources with fast response times for efficient loading of the trap with ^{174}Yb and ^{171}Yb isotopes

5. Laser and microwave fields for Doppler cooling, ionisation, state preparation and detection using $^{171}Yb^{+}$ hyperfine qubits

6. Magnetic gradient coils to enable microwave-driven quantum logic gates using the scheme described in the previous section

7. Broadcast system for coherent microwave and RF fields to drive gates using the magnetic gradient scheme

8. Ion imaging and state detection system

[1] With the exception of external cavity diode lasers which were already available in the laboratory.

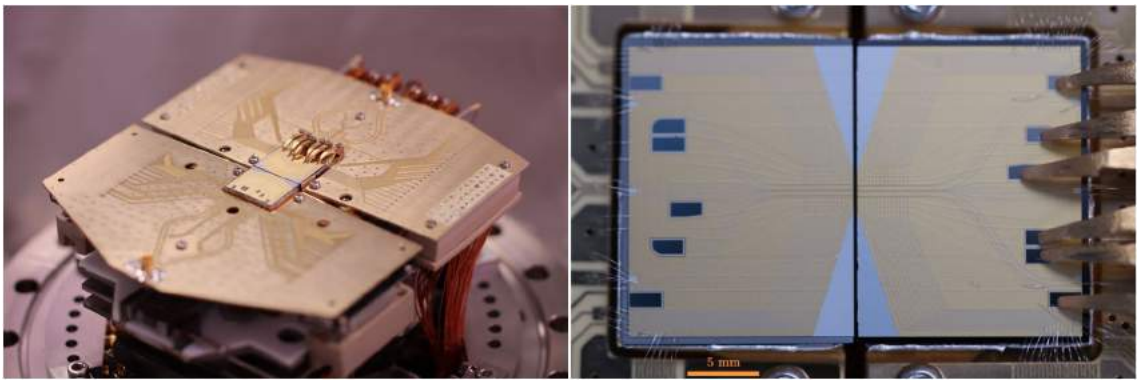

Figure 1.5: The scalable demonstrator experiment in which two ion trap quantum computing modules (left image) are connected by electric fields. The ion-trap microchips are separated by a small distance (right image) which is controllable in three dimensions using piezo stages to align the electrodes at the edges of the traps.

9. A scalable cooling system capable of being extended to any number of ion-trap modules, with an operating temperature chosen to provide a significant reduction of the ion motional heating rate

10. Vacuum system compatible with all of the sub-systems above, and capable of ultra-high vacuum pressures to support long ion lifetimes

Figure 1.5 shows the two-module processor mounted to the open vacuum chamber of the demonstrator experiment. The edge-trap chips and printed circuit boards are visible on the top surface of the device. A close-up of the edge traps shows the gold-plated surface electrodes, the small separation between chips (black vertical line), and just visible are the linear RF rails of each chip which extend horizontally on either side of the gap at the centre of the image (in the setup shown there is no X-junction).

1.7 Thesis summary

The structure of the thesis is as follows.

Chapter 2 provides the theoretical background on ion-trapping, including a description of the linear Paul trap and microfabricated ion traps. The Ytterbium ion is introduced as a hyperfine qubit, and the atomic energy level structure is discussed. Then the techniques used for ion trapping are explained, including two-step photoionisation and Doppler cooling, along with qubit initialisation and state detection. Dressed states are explained and I provide a description of how gates are performed in the dressed state basis.

Chapter 3 discusses the two-module ion-trap processor. I describe the surface trap chips and the requirements for module alignment to permit reliable transport between chips with

low motional excitation of the ion. I analyse the results of simulations showing the pseudopotential barrier in the gap region between chips. Then the effect of trap separation on ion dynamics is discussed, looking at the motional heating rate due to the pseudopotential gradient. The results of this section inform the requirements of the positioning system for module alignment.

Chapter 4 explains the wider experimental set-up of the demonstrator experiment. Here I describe the development of an ultra-high vacuum system to provide isolation for ion-trap quantum computing experiments, and the methods used to generate low-noise RF and DC potentials for ion trapping. I describe the laser set-up required for ionisation, Doppler cooling and state detection, the microwave and RF broadcast setup for coherent manipulation of ion qubits, and the scalable cryogenic cooling system.

Chapter 5 discusses the construction and characterisation of the two-module ion-trap processor and the preparation for experiments using the first generation of edge-trap microchips. I describe the electrical testing of the traps, and the procedures used to treat and resolve issues that arose with the traps prior to their installation into the experiment. I characterise the performance of the module positioning system by measuring its travel range and precision. I present the results of alignment tests with the ion trap modules.

Chapter 6 describes the trapping of $^{174}Yb^{+}$ using the second generation ion traps with integrated magnetic gradient coils. microchips, and the results of ion transport experiments in which $^{174}Yb^{+}$ ions were transported between independent trap modules. I perform tests of the RF behaviour of the microchips during the alignment procedure.

The thesis concludes with a summary and discussion of how the demonstrator informs the prospects for scaling-up ion trap processors to large numbers of qubits using modules connected by electric fields.

1.8 Chapter contributions

The construction of the demonstrator experiment was a collaborative project. The following colleagues contributed to the work described in this thesis:

- Chapter 3: Raphaël Lebrun-Gallagher developed the high stability mounting structure, module pre-alignment stage and imaging system for the experiment. Simulations performed by Zak Romaszko and Martin Siegele informed my analysis of the pseudopotential barrier at the interface between chips, and the requirements for module alignment.

- Chapter 4: Reuben Puddy, Zak Romaszko, Martin Siegele and Seokjun Hong developed the ion-trap microchips. Raphaël Lebrun-Gallagher designed the on-chip magnetic gradient coils and constructed the laser setup with Mariam Akhtar and Falk Bonus. Raphaël Lebrun-Gallagher and I co-developed the scalable cooling system.

- Chapter 6: Mariam Akhtar, Falk Bonus, Raphaël Lebrun-Gallagher and I conducted trapping runs with $^{174}Yb^{+}$ ions. Falk Bonus generated DC potentials for ion trapping and transport using trap simulations by Martin Siegele and software written by Sam Hile.

Chapter 2

Ion-trapping fundamentals and the Ytterbium hyperfine qubit

2.1 Introduction

In this chapter I provide a description of the methods used to confine atomic ions using ion traps, and the techniques that are employed to manipulate ion qubits for quantum information processing experiments.

Ion traps can be categorised into two basic types, each having a different principle of operation that allows for the confinement of charged particles in three dimensions. The first is the radio frequency (RF) quadrupole trap, also referred to as the Paul trap [75], which uses a combination of static and time-varying electric fields to generate an effective trapping potential. The second is the Penning trap [76], which combines an electrostatic potential with a strong homogeneous magnetic field to confine charged particles. The experimental work described in this thesis is based on the Paul trap. To generate a three-dimensional confining potential it is not sufficient to create a static distribution of electric charge, for example by applying a steady potential difference across the electrodes of a trap device, because no stable equilibrium is permitted by Earnshaw's theorem [77]. Using solely static charges, it is only possible to form saddle points in the potential, and not minima in three dimensions of space. This is a result of Gauss's Law in free space which states that the divergence of an electric field, $\boldsymbol{E}$, must be equal to zero,

$$\boldsymbol{\nabla} \cdot \boldsymbol{E} = 0. \tag{2.1}$$

If we relate this to the electric force on a point charge, q, arising from a potential, U, then

we have,

$$\boldsymbol{F} = q\boldsymbol{E} = -\boldsymbol{\nabla} U, \tag{2.2}$$

and we see that,

$$\nabla^2 U = \frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2} = 0. \tag{2.3}$$

Consequently, there must always be a direction along which a charged particle is not confined by the potential. The Paul trap overcomes this instability by using a time-varying electric potential to generate an effective minimum in three dimensions when averaged over time, as discussed in the next section.

2.2 The Paul trap

To illustrate the general principle of operation of a Paul trap, I will describe a variant known as the linear Paul trap, whose electrode structure is shown in Figure 2.1.

The design uses four rod-like electrodes and two 'endcap' electrodes for trapping a string of particles. Each pair of rod electrodes on opposite diagonals are connected together, and a potential $U + V_{RF}\cos(\Omega t)$ is applied to one pair, while to the other pair a potential $U - V_{RF}\cos(\Omega t)$ is applied, where U is a static offset voltage and V_{RF} is the amplitude of the oscillating potential. This generates an oscillating saddle potential, or quadrupole, that is confining in one direction but anti-confining in the perpendicular direction. A static potential V_{DC} is applied to the two endcap electrodes to confine along the axial direction (z-axis). The combined electric potential generated by the trap in three dimensions is the sum of the RF and DC potentials,

$$\Phi(x, y, z, t) = \Phi_{RF}(x, y, t) + \Phi_{DC}(x, y, z, t), \tag{2.4}$$

which individually are described by

$$\Phi_{RF}(x, y, t) = \eta_{RF}\Big(U - V_{RF}\cos(\Omega t)\Big)\frac{(x^2 - y^2)}{2{r_0}^2} \tag{2.5}$$

$$\Phi_{DC}(x, y, z, t) = \eta_{DC} V_{DC}\left(\frac{2z^2 - (x^2 + y^2)}{2{z_0}^2}\right) \tag{2.6}$$

where r_0 and z_0 are the distances from the centre of the trap to the RF and DC electrodes respectively, and η_{RF} and η_{DC} are geometric efficiency factors (having values between 0 and 1) that are equal to unity if the electrodes are perfectly hyperbolic [78]. Figure 2.2 is a plot of $\Phi_{RF}(x, y, t)$, showing the saddle potential in the xy plane at $t = 0$ and $t = 2\pi/\Omega$.

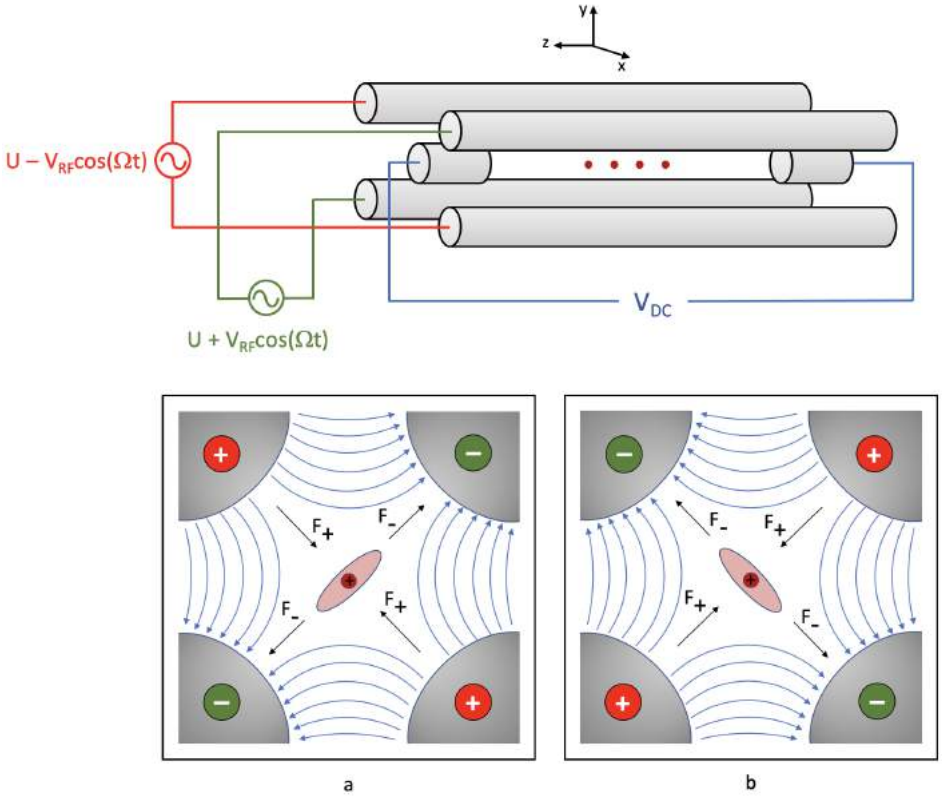

Figure 2.1: A linear Paul trap. The top image shows the three-dimensional layout of the four rod-like RF electrodes and two DC endcap electrodes. A string of ions is represented as red dots. The bottom images show the quadrupole electric field generated by the four RF electrodes at an instant of time (a), and again a half-cycle later after the oscillating fields have reversed direction (b). The central positive charge (dark red circle) represents a single ion located at the field minimum, while the oval structure surrounding it represents a cloud of ions close to the trap centre, being alternatively pushed and pulled by the changing force.

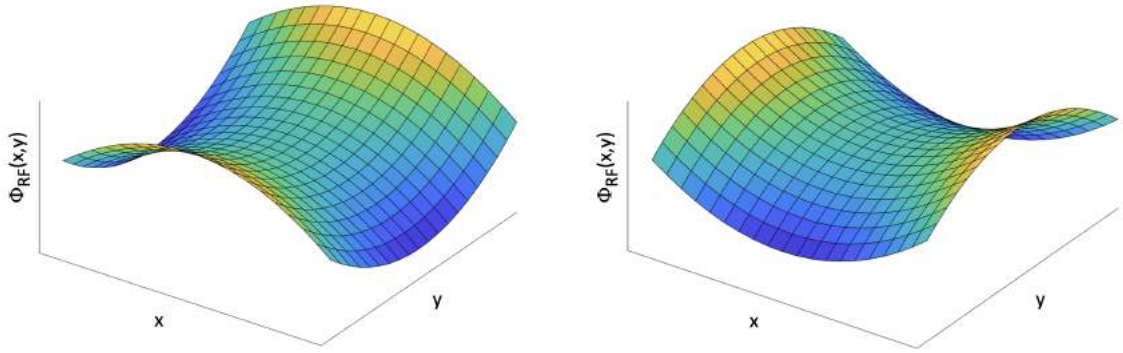

Figure 2.2: Saddle potential generated by the RF electrodes of a Paul trap, plotted at $t = 0$ (left) and a half-cycle later at $t = \pi/\Omega$ (right), when the polarity has inverted [78].

The time-independent pseudopotential approximation considers the time-averaged force on the ion due to the oscillating RF potential, $\Phi_{RF}(x, y)$, as if this force was the result of a static 'pseudopotential', $\Phi_p(x, y)$, defined by

$$\langle \boldsymbol{F}(x, y, t) \rangle = -e\boldsymbol{\nabla}\Phi_P(x, y), \tag{2.7}$$

$$\Phi_P(x, y) = \frac{e}{2m\Omega^2} \left\langle \boldsymbol{E}_{RF}^2(x, y, t) \right\rangle . \tag{2.8}$$

where $\langle\rangle$ denotes the time average performed over the period $2\pi/\Omega$, e is the fundamental electronic charge and m is the mass of the ion. Since the electric field can be written $\boldsymbol{E}_{RF} = -\boldsymbol{\nabla}\Phi_{RF}$, and using Equation 2.5, we arrive at

$$\Phi_P(x, y) = \frac{e}{4m\Omega^2} \frac{\eta_{RF}^2 V_{RF}^2}{r_0^4}(x^2 + y^2). \tag{2.9}$$

The effective confinement due to the pseudopotential is harmonic, that is, quadratic in x and y (for ion positions close to the potential minimum and assuming perfectly hyperbolic RF electrodes). When combined with the static potential Φ_{DC} due to the endcap electrodes, the Paul trap creates an effective confining potential with quadratic minima in all three dimensions.

2.2.1 Ion motion within a Paul trap

The motion of an ion in the potential described by equation 2.4 can be decomposed into an axial oscillation in the z-direction and radial motion in the xy plane. The axial motion is the simpler case to consider, as it is defined by the static potential generated by the two endcap electrodes. We can apply equation 2.2 to calculate the electric force on a particle with charge, e, and mass, m, to arrive at the differential equation,

$$\frac{\partial^2 z}{\partial t^2} = -\frac{2e\eta_{DC}V_{DC}}{m{z_0}^2} z, \tag{2.10}$$

which is simple harmonic motion in the z-direction with angular frequency,

$$\omega_z = \sqrt{\frac{2e\eta_{DC}V_{DC}}{m{z_0}^2}}. \tag{2.11}$$

This is the axial secular frequency of the trap. Next we consider the more complicated radial motion of the ion in the xy plane, which is described by the well-known Matthieu equation [79],

$$\frac{\mathrm{d}^2u}{\mathrm{d}\xi^2} = -[a_u - 2q_u\cos 2\xi]u, \tag{2.12}$$

where u represents the spatial coordinate, either x or y, and $\xi = \frac{\Omega t}{2}$. The two stability parameters are

$$q_x = -q_y = \frac{2e\eta_{RF}V_{RF}}{m\Omega^2{r_0}^2}, \tag{2.13}$$

$$a_x = -a_y = -\frac{2e\eta_{DC}V_{DC}}{m\Omega^2{z_0}^2}, \tag{2.14}$$

and their relative values determine whether or not the motion of the ion remains stable within the trap. Under the conditions $q \ll 1$ and $a \ll q$ (which is the case for $\Omega/2\pi$ in the MHz range and $V_{DC} \ll V_{RF}$), and restricting consideration to the x-axis for clarity, the motion in the x-direction is approximately,

$$x(t) \approx x_0\cos(\omega_x t)\left[1 + \left(\frac{\sqrt{2}\omega_x}{\Omega}\right)\cos(\Omega t)\right], \tag{2.15}$$

where ω_x is the characteristic radial frequency of the ion, which is referred to as the secular frequency,

$$\omega_x = \frac{e\eta_{RF}V_{RF}}{\sqrt{2}m\Omega{r_0}^2}. \tag{2.16}$$

We can see from equation 2.15 that the trajectory of the ion is a combination of a slower oscillation at the secular frequency, ω_x, and a faster oscillation at the trap drive frequency, Ω, applied to the RF electrodes. The secular frequency is characteristic of the shape and magnitude of the quadrupole potential. The faster oscillation is termed micromotion and it is usually an unwanted effect that we wish to minimise during quantum information experiments. Owing to the ion's secular motion away from the position of the potential minimum, it is impossible to eliminate micromotion entirely.

Stable ion trajectories exist only for certain values of the stability parameters q_u and a_u described above, and outside of this parameter space the ion cannot remain bound within the trap potential. Here it is useful to define a measure called the trap depth, which is the difference in energy between the potential minimum and the escape energy of the trap. A trapped ion remains bound within the trapping potential so long as its kinetic energy is less than the trap depth. The length of time for which an ion remains bound is referred to as the ion lifetime. The ion lifetime is reduced by processes that cause the ion to gain

kinetic energy, and these are the topic of the next section.

2.2.2 Motional heating of trapped ions

A number of different phenomena within the environment of the trap can cause the kinetic energy of a trapped ion to increase. This kinetic energy gain is called motional heating, and the heating rate is defined as the mean number of motional quanta that the ion acquires per unit time, $\dot{\bar{n}}$. Two distinct processes contribute to the ion heating rate, so that we can write the total heating as a sum of two terms,

$$\dot{\bar{n}} = \dot{\bar{n}}_{ext} + \dot{\bar{n}}_{anom}, \tag{2.17}$$

where $\dot{\bar{n}}_{ext}$ arises from the presence of noise on external voltage sources, and $\dot{\bar{n}}_{anom}$ is 'anomalous heating'.

The first term represents noise on the RF pseudopotential due to imperfect voltage sources, which causes motional heating to occur if the ion is positioned away from the RF electric field null [80]. At the exact position of the electric field null heating is negligible, but a pseudopotential gradient arises within a junction [81] which causes a non-zero field at the ion position. This is also the case in the gap region between separate traps in the two-module processor (as discussed in Chapter 3). Motional heating couples most strongly when the noise frequency is resonant with the ion secular frequency. If we consider ion positions along the z-axis with axial secular frequency ω_z, and voltage noise given by a voltage spectral density, S_{V_N}, the heating rate for noise around $\Omega_{RF} + \omega_z$ is given by

$$\dot{\bar{n}}_{ext} = \frac{e^4}{16m^3\hbar\omega_z\Omega_{RF}^4}\left[\frac{\partial}{\partial z}E_0^2(z)\right]^2 \frac{S_{V_N}(\Omega_{RF}+\omega_z)}{V_0^2}, \tag{2.18}$$

where $E_0(z)$ is the electric field in the z direction and V_0 is the peak RF potential applied to the trap [61]. It can be seen that in the ideal situation, where the ion is positioned on a continuous RF nill along the trap axis, the term $\frac{\partial}{\partial z}E_0^2(z)$ is zero since there is no gradient in the trapping field.

Assuming that charge fluctuations occur on the electrode surfaces owing to the finite impedance of the electrodes - an effect known as Johnson noise - then we would expect the heating rate to exhibit a $1/d^2$ dependency [80], where d is the distance to the nearest electrode surface. However, experiments have measured an approximate $1/d^4$ dependency [82, 83, 84], indicating that other phenomena on the surface must contribute strongly to the motional heating of the ion, and leading to the name 'anomalous heating' to refer to

this unexpected result.

Attempts have been made to characterise the effect of electrode surface roughness [85, 86] on the heating rate, and while the underlying physical mechanism is not yet fully understood, the results suggest that very small scale surface features such as single atomic monolayers of contaminants and and adsorbents are responsible for generating electric field noise. Experiments have demonstrated that cleaning the ion-trap surface using argon ion beam bombardment [87] or pulsed laser light [88] can improve the measured heating rate by two orders of magnitude. The heating rate has also been shown to decrease at lower temperatures [89], and in Ref. [90] it was estimated that cooling the ion trap to a temperature of 70 K would reduce the heating rate by a factor of 10 [90]. The above considerations are particularly important for microfabricated ion traps because of the relatively close proximity of the ion to electrode surfaces, which can result in significant heating of the ion given the strong dependence of the heating rate on ion-electrode distance.

2.3 Microfabricated ion traps

Microfabrication techniques permit the construction of miniaturised ion traps with electrode dimensions of a few tens or hundreds of micrometres, by depositing layers of metal and insulating material onto a substrate such as silicon [91, 92]. The advantage of microfabricated ion traps is the relative ease with which identical copies can be constructed with precise features using well-established manufacturing processes. Ion trap microchips with complex geometries have been demonstrated, including Y-junctions [93] and X-junctions [65] for reordering and routing of ions between trapping zones, and on-chip features have been built using micro-electromechanical systems (MEMS) technologies, such as apertures for ion-loading [93], integrated microwave circuitry for driving gates [94, 95], and on-chip micro-optical elements such as mirrors for fluorescence collection [96]. Therefore, microfabricated ion-traps are attractive as a basis on which to construct a large-scale quantum computing architecture using two-dimensional trap arrays [49, 34].

2.3.1 Three-dimensional and two-dimensional ion trap designs

Microfabricated ion traps fall into two categories: three-dimensional (3D) or two dimensional (2D) depending on the configuration of their electrodes. In 3D traps there are multiple layers of electrodes, so that these traps more closely resemble a miniaturised version of the macroscopic Paul trap pictured in Figure 2.1. Two or three layers of electrodes are separated by insulating material, as illustrated in the first two images of Figure 2.3.

Typically a three-layer structure results in a larger trap depth for a given applied RF voltage, with finer control over the axial (z-axis) positioning of the ion [91], while a two-layer structure has the advantage of being simpler to fabricate. In 2D traps, all of the electrodes are fabricated in a single layer on the surface of a substrate, as shown in image (c) of Figure 2.3, which is why this design is also referred to as a surface-electrode ion trap or 'surface trap'. Surface traps have the advantage of being relatively simple to fabricate compared to 3D designs [91]. For the demonstrator experiment, surface traps have been developed with up-to-edge electrodes to permit their trapping regions to be closely aligned to allow ion transport between separate chips. These are described in Chapter 3.

The basic configuration of electrodes on a linear surface trap is known as a 'five wire' design [91] because of the two RF tracks and three lines of DC electrodes, as shown in Figure 2.4. More complex variations on the five wire design can be fabricated to allow for finer control over the position of the ions above the trap surface, and to enable fine adjustments to the trap potential. An example is given on the right hand side of Figure 2.4. Segmentation of the inner ground electrode between the two RF tracks permits the separation of neighbouring ions by splitting the potential into two adjacent wells. After separation, ions may be transported to other zones on the trap for specific operations such as gates or state detection. Additionally, dedicated DC rotation electrodes are usually added to increase the effectiveness of laser Doppler cooling. These are located on either side of each RF track and carry static voltages to rotate the principal axes of the pseudopotential, creating an angle with respect to the surface normal. This rotation is needed for surface electrode traps because Doppler cooling requires that the wavevector of the

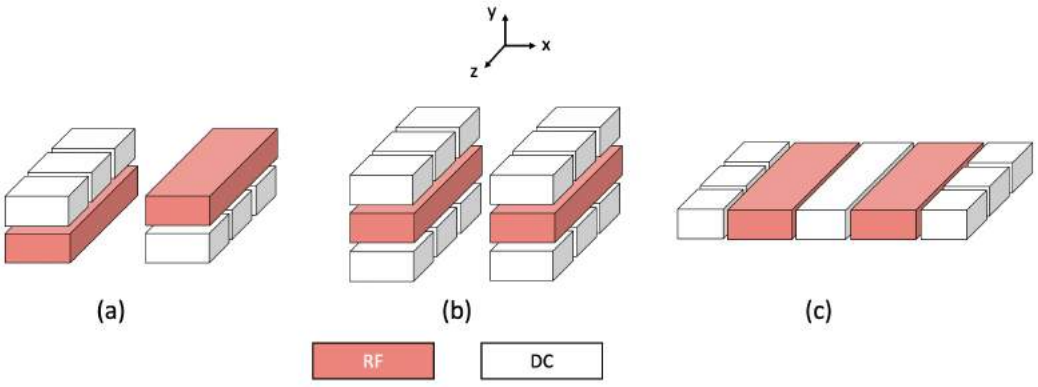

Figure 2.3: Examples of 3D and 2D ion trap geometries. In 3D traps, multiple layers of electrodes surround the ion. Images (a) and (b) show examples of two layer and three layer designs. Image (c) shows a 2D or planar ion trap geometry in which all of the electrodes are fabricated in a single surface layer. For this reason they are also known as surface traps [91].

applied laser beam has components in all three principal axes of the trap in order to be effective in damping the motion of confined ions. For the laser beam to access the ion without reflecting off the surface of the trap, the laser is typically constrained to travel parallel to the trap surface, and cannot be angled in such a way as to achieve a significant vertical component (the direction normal to the surface). A rotation of ~ 10 degrees can be achieved by applying voltages of +1 V and -1 V to the rotation electrodes [90].

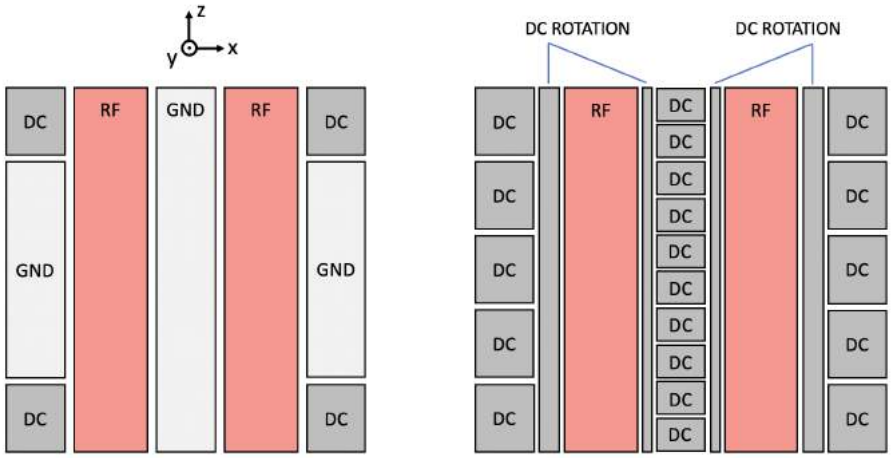

Figure 2.4: (*Left*) A basic 'five wire' electrode layout for a linear surface trap. (*Right*) Modification of the five wire design to achieve additional control over the trapping potential. Inner segmented DC electrodes permit fine control over ion transport and separation/recombination operations. Rotation electrodes carry a DC bias voltage to adjust the principal axes of the trap potential to enable Doppler cooling of all motional modes using a single laser directed parallel to the trap surface.

To reduce the amplitude of micromotion, the minimum of the static harmonic potential produced by the DC endcap electrodes must overlap precisely with the minimum of the quadrupole potential. To achieve this overlap requires careful choice of the voltages applied to the DC electrodes, which can be determined by simulation of the trap potential based on the electrode configuration and dimensions. This process of aligning the quadrupole and static potentials generated by the trap is known as micromotion compensation.

The electric potential above a surface trap microchip is shown in Figure 2.5. For the trapping potential illustrated in the figure, no rotation has been applied and so the y-axis and surface normal are parallel in this case. The ion is trapped at the pseudopotential minimum in the xy plane, which is located within the closed contour in the plot shown in Figure 2.5. The ion is confined in the z-direction by DC endcap electrodes. The ion height above the chip surface depends on the separation between the two RF tracks and the relative dimensions of the RF and central ground electrode.

For a surface trap, the direction of weakest confinement is always along the y-axis,

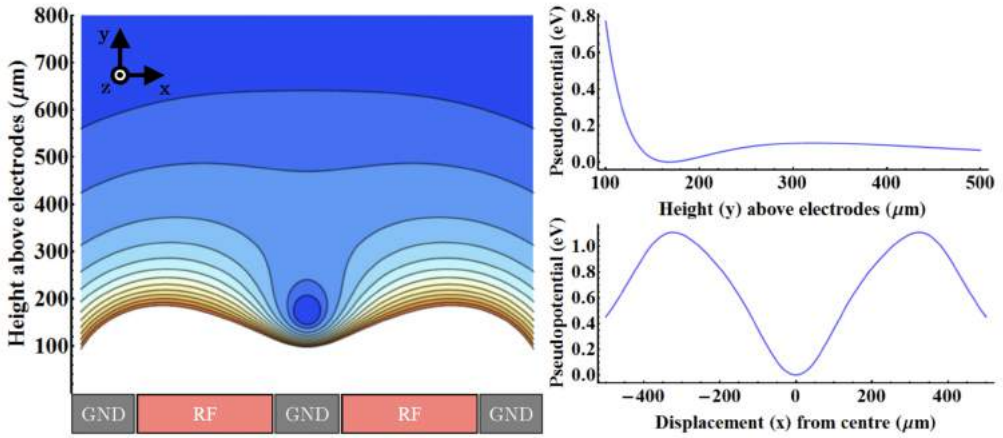

Figure 2.5: Electric potential above a surface-electrode ion trap microchip. (*Left*) The ion is trapped above and between the RF electrodes, within the ellipse formed by the closed contours. (*Right*) The weakest confinement is in the vertical ($+y$) direction. Image adapted from [90].

away from the chip surface. This can be seen in Figure 2.5, by comparing the two graphs on the right of the figure. The energy barrier in the y-direction is $\sim 0.15\,\text{eV}$, while in the x-direction it is $\sim 1.2\,\text{eV}$ [90].

2.4 The Ytterbium ion

This section describes the energy level structures of ytterbium ^{174}Yb and ^{171}Yb, and the techniques by which ytterbium atoms can be ionised, trapped and cooled to permit quantum information experiments with single ions.

2.4.1 Photoionisation

The outermost electron in neutral ytterbium has an ionisation energy of 6.25 eV [97], which may be provided by a photon of wavelength 198.3 nm to excite the transition from the $^1S_0$ ground state to the continuum. Laser light at this wavelength is difficult to generate, and it is more convenient to supply this energy by using the ECDL lasers available within the research group, in a two step ionisation process as illustrated in Figure 2.6. First, a laser with wavelength close to 399 nm is used to promote the electron from the $^1S_0$ ground state to the $^1P_1$ state. From there, the transition to the continuum requires a wavelength less than 397.1 nm and so we can use a 369.5 nm laser (also used for Doppler cooling and state detection, see next section) to ionise the atom. In practice this two-step process is easily achieved by overlapping the two laser beams when illuminating the atom beam from an atomic source (a more detailed description of the method is provided in Chapter 4). Exact

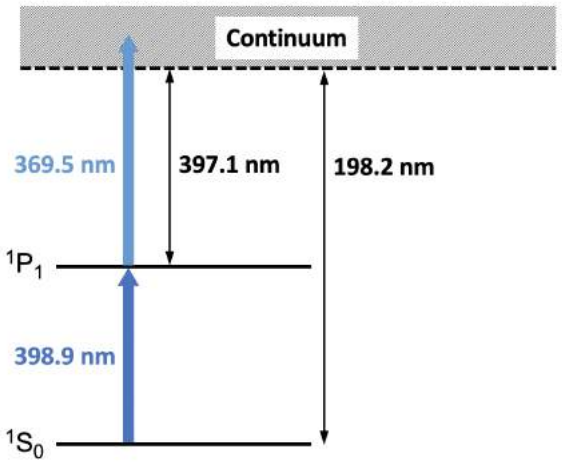

Figure 2.6: Two-stage ionisation of the Ytterbium atom. The exact wavelength of the transition from $^1S_0$ to $^1P_1$ (near 399 nm) depends on the isotope, and this provides a means to select between different Ytterbium species (in our case, ^{171}Yb and ^{174}Yb). A second laser, with wavelength 369 nm, is used to excite the electron from the $^1P_1$ level to the continuum, ionising the atom.

wavelengths are provided in [98].

2.4.2 Doppler cooling

As we recall from section 2.2.2, ions in a trap gain kinetic energy over time in a process known as motional heating, which results from electric field noise in the vicinity of the trapping potential. If left untreated, motional heating can result in the ion acquiring sufficient energy to escape from the trap. To improve the ion lifetime within the trap, a method is needed to remove kinetic energy from the ion, and this can be achieved using the technique of laser Doppler cooling [99]. By selecting an appropriate atomic transition and detuning the laser towards the red sideband, any ion with a velocity component opposite in direction to the wavevector of the laser will experience a reduced effective detuning due to the Doppler effect. The opposite shift occurs for ions moving away from the laser, in which case the effective detuning increases. This leads to a velocity dependent photon scattering rate, and a force on the ion that depends on the magnitude and direction of its velocity, slowing the ion's motion.

For a two-level particle with energy separation, ω_0, moving with velocity, $\boldsymbol{v}$, the effective detuning, $\delta_{eff} = \delta - \boldsymbol{k}.\boldsymbol{v}$, where δ is the unshifted detuning and $\boldsymbol{k}$ is the wavevector of the laser [100]. For light with intensity, I, we can define the saturation parameter as $s_0 = I/I_{sat}$, where

$$I_{sat} = \frac{\hbar\Gamma{\omega_0}^2}{6c}, \tag{2.19}$$

with Γ being the linewidth of the atomic transition and c the speed of light. The rate of photon scattering is therefore,

$$\gamma_{scatt} = \frac{s_0(\Gamma/2)}{1 + s_0 + (2\delta_{eff}/\Gamma)^2}, \tag{2.20}$$

and the force on the particle is

$$F_{scatt} = \hbar k \gamma_{scatt}. \tag{2.21}$$

Since Doppler cooling relies upon the difference in the scattering rates due to the velocity of the particle along the direction of the laser, its effectiveness reduces as the motional energy of the particle is reduced. This is called the Doppler limit, and the minimum temperature achievable with Doppler cooling is called the Doppler temperature,

$$T_D = \frac{\hbar\Gamma}{2k_B}, \tag{2.22}$$

where k_B is the Boltzmann constant.

For Doppler cooling to be successful in damping the motion of an ion in three dimensions, the orientation of the laser beam must be chosen so that there is a component of the wavevector, $\boldsymbol{k}$, in all three of the principal axes defining the motion of the ion (x, y and z as shown in Figure 2.4). Additionally, in comparison to the analysis given above, ytterbium ions have more than two energy states and it is necessary to excite fluorescence continuously without losing population to other nearby states, which can occur due to off resonant excitation, spontaneous decay or collisions with background gases. To achieve this, we are required to set up a 'closed' Doppler cooling cycle using additional laser wavelengths that are chosen to excite transitions from adjacent states and return the population back to the main fluorescence transition.

Natural ytterbium consists of the stable isotopes with nucleon numbers 168, 170, 171, 172, 173, 174 and 176. For quantum information experiments the isotope ^{171}Yb is chosen because of its suitability both as a hyperfine qubit and for its magnetic-field sensitive Zeeman states (that are compatible with the magnetic gradient scheme used to enact quantum logic operations as described in section 2.5). These properties arise from the nuclear spin which results in atomic energy levels that have an associated magnetic moment [101]. The hyperfine structure makes this isotope a more demanding candidate for Doppler cooling because of the additional transitions that must be addressed by radiation fields in order to set-up a closed cooling cycle. For this reason, experiments begin by trapping ^{174}Yb$^+$ because of its zero nuclear spin and lack of hyperfine structure, making Doppler

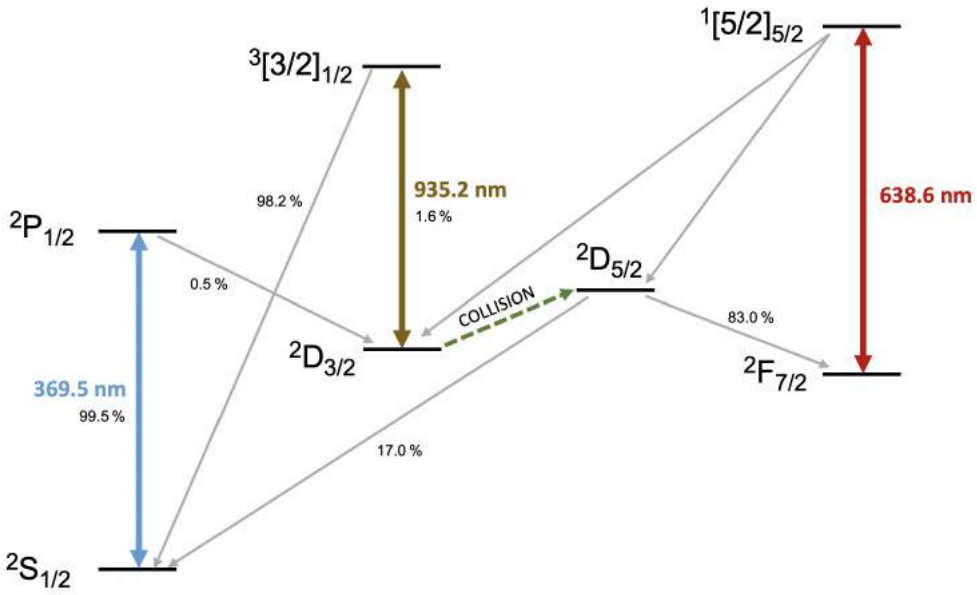

Figure 2.7: Energy level diagram for $^{174}Yb^+$ ion showing the atomic transitions and laser wavelengths required for Doppler cooling. The principal transition is between the $^2S_{1/2}$ and $^2P_{1/2}$ states, which is driven by a laser with wavelength close to 369 nm. The other wavelengths are needed to return lost population back to the $^2S_{1/2}$ ground state, from where the cooling cycle repeats.

cooling relatively simple in comparison to $^{171}Yb^+$. Using $^{174}Yb^+$ we can confirm that the ion trap is operating as expected and optimise the trapping parameters, before moving to $^{171}Yb^+$ for qubit manipulation experiments.

The atomic transitions and laser wavelengths required for Doppler cooling of $^{174}Yb^+$ ions are shown in Figure 2.7. All transitions can be driven by commercially available ECDL lasers. The main transition is between the $^2S_{1/2}$ and $^2P_{1/2}$ states, but there is a small probability (0.5 %) that the $^2P_{1/2}$ state will decay into other 'dark' states, which disrupt the process of photon absorption and fluorescence on the main transition that is required for Doppler cooling. Therefore, additional lasers are needed to excite transitions from these dark states to repopulate the main cooling transition. A laser with wavelength 935.2 nm is used to return population from the $^2D_{3/2}$ state, but occasionally background gas collisions with the ion cause transitions to the $^2D_{5/2}$ state, which requires a laser at 638.6 nm to close the cooling cycle. Again, the exact wavelengths are provided in [98].

For $^{171}Yb^+$, the nuclear spin of $1/2$ complicates the Doppler cooling process because of the requirement to address the hyperfine transitions. The energy levels and radiation fields needed to set-up a closed cooling cycle for $^{171}Yb^+$ are shown in Figure 2.8. The $^2S_{1/2}$ and $^2P_{1/2}$ states each split into two hyperfine levels with angular momentum $F = 0$

and $F = 1$. With the 369 nm laser tuned to drive fluorescence on the $^2S_{1/2}$ $F = 1 \leftrightarrow {}^2P_{1/2}$ $F = 0$ transition, off-resonant scattering can populate the adjacent hyperfine levels, which then decay to states not resonant with the laser. Therefore additional frequencies must be applied to return the population back to the cooling transition. The frequency splitting in the $^2S_{1/2}$ manifold is $\sim$ 12.64 GHz, and can be addressed with microwave radiation. The frequency splitting in the $^2P_{1/2}$ manifold is $\sim$ 2.1 GHz, which can be addressed by modulating the frequency of the laser using an electro-optic modulator (EOM). The D and F states likewise exhibit hyperfine splitting with transition frequencies in the GHz or MHz range. These transitions can be addressed if the laser beam intensity is set to be much larger than the saturation intensity, which results in power broadening of the transitions and removes the need for additional radiation fields [98]. In this way, we can use the 369 nm, 935 nm and 638 nm lasers to excite transitions for all of the sub-levels within the P, D and F states to repopulate the $^2S_{1/2}$ state and thus create a closed cycle for Doppler cooling.

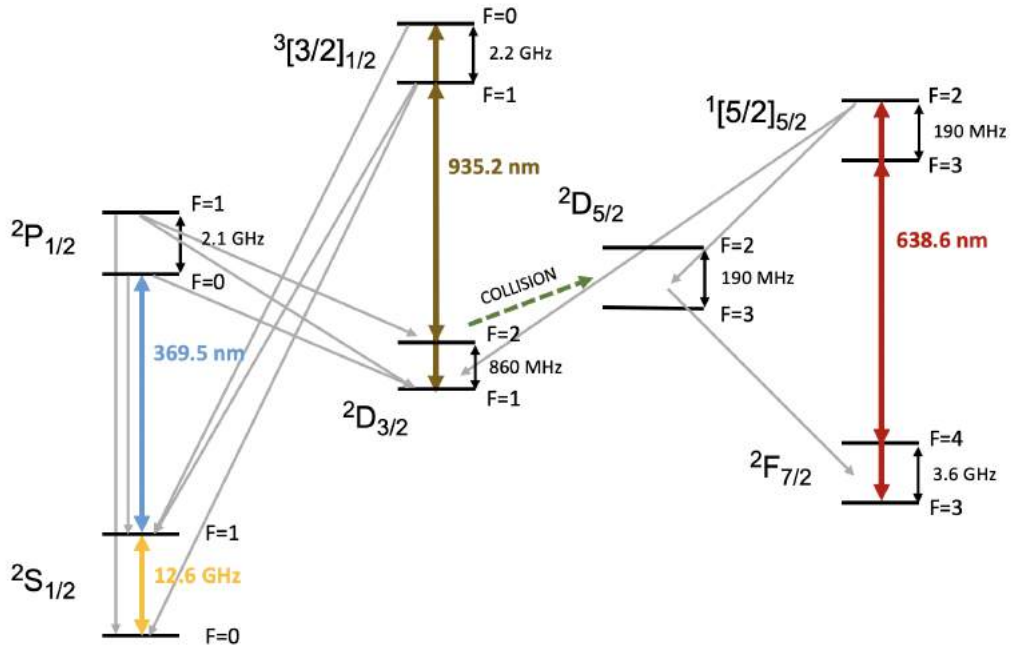

Figure 2.8: Atomic transitions required for Doppler cooling of the $^{171}Yb^+$ ion. Hyperfine splitting of the energy levels occurs for this isotope due to a nuclear spin of 1/2.

2.5 The Ytterbium hyperfine qubit

The $^2S_{1/2}$ ground state manifold of $^{171}Yb^+$ is shown in Figure 2.9. The ground state exhibits hyperfine splitting [101] owing to the angular momentum coupling between the nuclear spin of $1/2$ and the total electronic spin of $1/2$. There are two possible ways for the spins to combine, giving rise to two hyperfine levels with total angular momentum of $F = 0$ and $F = 1$. In the absence of a magnetic field, the projected magnetic states $|0'\rangle \equiv |F = 1, m_F = 0\rangle$, $|+1\rangle \equiv |F = 1, m_F = +1\rangle$ and $|-1\rangle \equiv |F = 1, m_F = -1\rangle$ in the $F = 1$ sub-manifold are degenerate in energy. In a magnetic field, the $F = 1$ sub-manifold undergoes Zeeman splitting and the three states acquire a difference in energy. The direction of the applied magnetic field defines a quantization axis with respect to the polarisation of the radiation used to drive transitions. The two states with $m_F = 0$ are labelled $|0\rangle \equiv |F = 0, m_F = 0\rangle$ and $|0'\rangle \equiv |F = 1, m_F = 0\rangle$, and we refer to them as the 'clock' states. They are only second-order sensitive to magnetic fields [73], and therefore are naturally protected from ambient magnetic field fluctuations. The frequency of the clock transition in a weak magnetic field is $\omega_c \approx \omega_0 \approx 2\pi \times 12.642\,812\,1\,\text{GHz}$. A qubit formed between the clock states has a long coherence time owing to the reduced spontaneous emission probability [101], but the lack of magnetic moment prohibits such a qubit from being compatible with the magnetic gradient scheme for driving multiqubit logic operations [74] (described in section 2.6.2) as required by our scalable architecture for quantum computing [34], that was introduced in Chapter 1. The magnetic states $|+1\rangle$

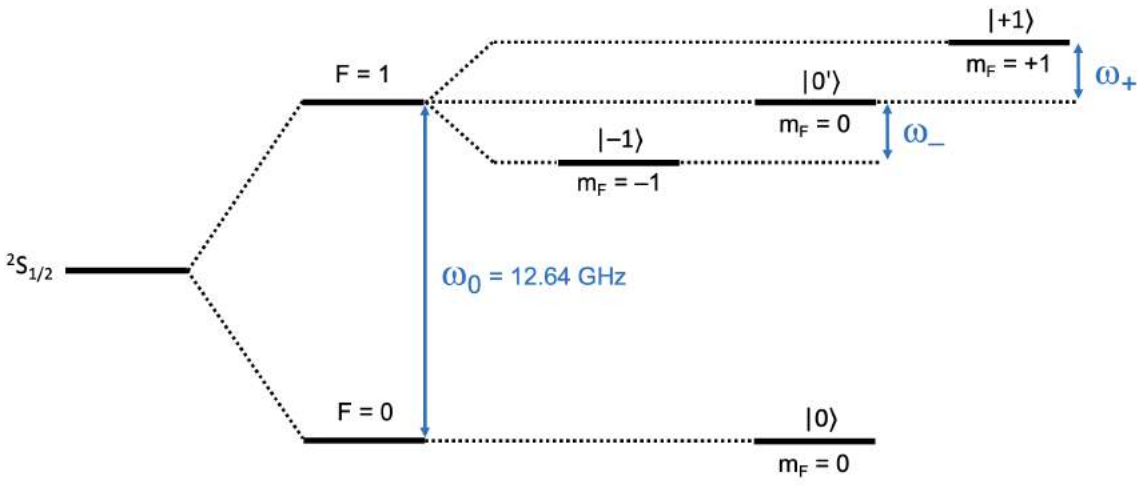

Figure 2.9: The hyperfine manifold of the $^{171}Yb^+$ ion. In the absence of an external magnetic field, only the two hyperfine levels $|0\rangle$ and $|0'\rangle$ are observed, and the $F = 1$ states are degenerate in energy. In magnetic field is applied to the ion, the degeneracy is lifted owing to the magnetic moment of the atom in the field (the Zeeman effect) and thus the states $|0'\rangle$, $|+1\rangle$ and $|-1\rangle$ differ in energy.

and $|-1\rangle$ are separated from the state $|0'\rangle$ by the splitting frequencies ω_+ and ω_- which are linearly dependent on the magnitude of the applied magnetic field. For weak magnetic fields, $\omega_+ \approx \omega_- \approx \mu_B B/\hbar$, and in a field of 10 G the splittings have values of around 14 MHz [102].

It is possible to set-up a magnetically sensitive qubit between the hyperfine ground state $|0\rangle$ and either of the two magnetic states $|+1\rangle$ or $|-1\rangle$. These states are referred to as the 'bare' atomic states, in contrast to the microwave 'dressed' states that are described later in this chapter. The advantage of constructing a qubit using the bare states is the comparative ease with which the qubit can be addressed and manipulated with microwave radiation. The disadvantage is the short coherence time of the qubit due to its sensitivity to ambient magnetic field fluctuations, which causes dephasing of the qubit state. This disadvantage can be overcome by the use of microwave dressed states that are protected from the effects of magnetic field noise. A method for creating these states is given in section 2.6.3.

2.5.1 Initialisation and state detection

For quantum information processing, it is essential to be able to prepare qubits in an initial state with high fidelity, and to have a reliable method for readout of the quantum state after coherent manipulation of the qubit.

Typically for trapped ion qubits, the initial state can be prepared by optical pumping, using lasers resonant with a well-chosen set of atomic transitions, and from which the ion can decay into dark states that are not coupled by the lasers. The population will remain trapped in these states if they are sufficiently de-coupled from the frequency of the lasers. For $^{171}Yb^+$ this can be achieved by using the same set-up that was implemented for Doppler cooling, with the three lasers of wavelength 369 nm, 935 nm and 638 nm as shown in Figure 2.8. By turning off the 12.64 GHz microwave field applied between the $^2S_{1/2}$ hyperfine states, and setting the 369 nm laser to be resonant with the $^2S_{1/2}$ $F = 1 \leftrightarrow {}^2P_{1/2}$ $F = 1$ transition, the state of the ion can decay into the $^2S_{1/2}$ $F = 0$ hyperfine state which is not directly coupled to any of the lasers. We observe that the population transfers into this state over time, where it remains effectively trapped. After sufficient time has passed, the qubit will have been prepared with very high accuracy (typical infidelity $<10^{-4}$ [102]) into the $|0\rangle \equiv \left|{}^2S_{1/2}, F = 0\right\rangle$ state.

A method for quantum state detection is required for qubits constructed within the $^2S_{1/2}$ manifold. We can discriminate between the state $|0\rangle \equiv \left|{}^2S_{1/2}, F = 0\right\rangle$ and the

excited states within the $^2S_{1/2}, F = 1$ manifold (any of the states $|+1\rangle$, $|0'\rangle$ or $|-1\rangle$, that we can define to be our qubit state $|1\rangle$), using the following measurement procedure. The three lasers are applied to the same transitions that were used for initialisation, with the microwave field turned off. This results in fluorescence if the ion state is within the $F = 1$ manifold (the qubit state $|1\rangle$), otherwise the ion remains dark and we conclude that it is in the state $|0\rangle$. Transitions between the states within the $^2S_{1/2}$ manifold have the polarisation constraints [101] shown in Figure 2.10. To ensure that transitions are excited from all three of the $F = 1$ states, the polarisation of the 369 nm laser with respect to the quantization axis (defined by the B-field direction) must have equal amounts σ^- (right circular), σ^+ (left circular) and π (linear). This can be achieved by placing $\lambda/2$ and $\lambda/4$ waveplates into the beampath of the laser before the ion-trap. We then apply the laser for a fixed amount of time and count the number of fluorescence photons collected by a photomultiplier tube (PMT). If this is repeated many times, a histogram may be plotted similar to the one shown in Figure 2.11. By setting a threshold number of photons, we can decide whether the ion

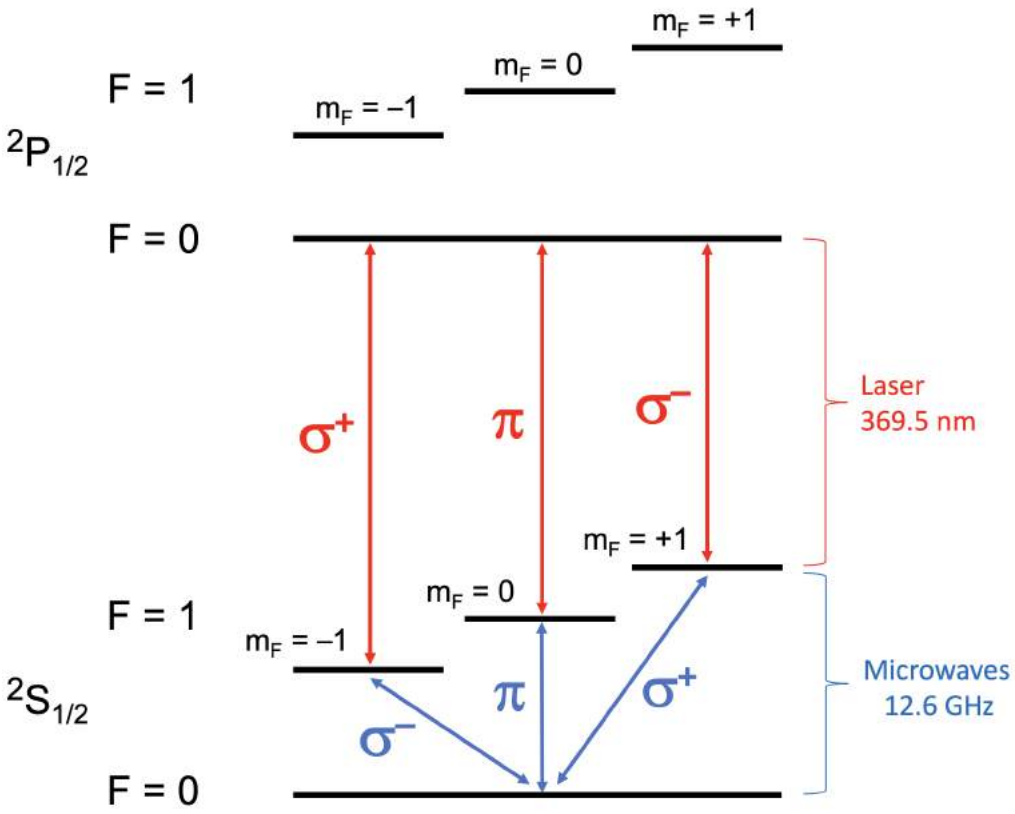

Figure 2.10: Polarisations required to excite transitions between the $^2S_{1/2}$ manifold and the $^2P_{1/2}, F = 0$ state of $^{171}Yb^+$. The transitions shown in red are excited by a laser near 369.5 nm with equal amounts of left circular (σ^+), right circular (σ^-), and linear (π) polarisation with respect to the quantisation axis set by an applied magnetic field. The hyperfine transitions shown in blue are excited by applying microwaves with frequency near 12.6 GHz.

was measured to be in the 'dark' state $|0\rangle \equiv \left|^2S_{1/2}, F = 0\right\rangle$ or the 'bright' state $|1\rangle \equiv \left|^2S_{1/2}, F = 1, m_F = +1, 0, -1\right\rangle$. If the number of photons detected is equal to or below the threshold then the state is measured to be 'dark', otherwise if the number of photons is above the threshold then the state is measured to be 'bright'. The main limitation of this technique is caused by off-resonant coupling of the 369 nm laser to the hyperfine states within either of the $^2S_{1/2}$ or $^2P_{1/2}$ levels, causing errors in detection and increasing the overlap between the histograms in Figure 2.11. The likelihood of this coupling can be reduced by shortening the detection time for which the laser is turned on. However, since the fraction of photons collected is typically small, there may be insufficient photons detected to distinguish the dark and bright states, along with significant background noise due to laser scatter from the chip surface. This is not a fundamental problem, and in future experiments it can be overcome by increasing the collection efficiency of the detection optics, for example by increasing the diameter of the objective lens or positioning the lens closer to the ion to increase the photon count for a given detection time. In reference [103], specialised optics were used to achieve a detection fidelity of 0.99915(7) for the hyperfine levels of $^{171}Yb^+$. Additionally, microfabrication techniques have been used to integrate a mirror onto the surface of an ion trap to reflect fluorescence photons back towards the detection optics and increase collection efficiency for faster state detection [96].

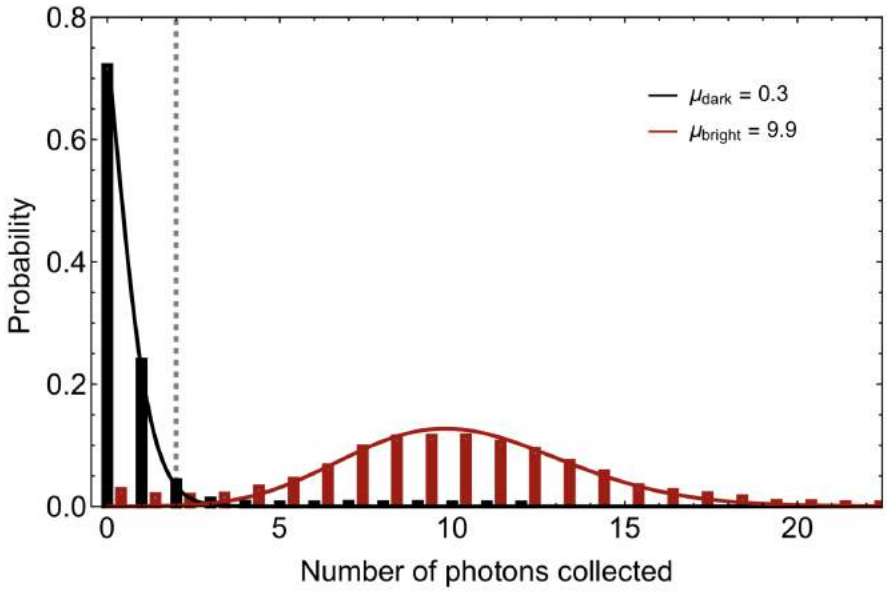

Figure 2.11: Histograms showing the number of photons collected by a PMT after performing many repetitions of the state detection procedure for ions prepared in either the dark state $|0\rangle$ (black) or the bright state $|0'\rangle$ (red). The results are used to select a discrimination threshold for state detection, in this case n=2 photons. Measurements equal to or below the threshold are recorded as detecting the dark state, while measurements above the threshold are recorded as the bright state. Image reproduced from [102].

For the situation where two ions are closely spaced within a trap (having typical separations of a few micrometres), the state detection fidelity is lower when using a PMT because there is a greater degree of overlap between the photon number histograms. For more than two ions, the fidelity decreases further, so this method is not scalable. This can be remedied by using a CCD camera in place of the PMT, in which case the ions can be spatially resolved and a separate photon count recorded for each ion. However, for large numbers of ions in an array this method becomes limited by the resolution and pixel size of the CCD sensor. In the scalable quantum computing architecture proposed in [34], ions on a surface trap array may be separated and transported to detection zones for readout of the quantum state away from other ions. The group has produced microchips that feature multiple trapping zones and segmented electrodes to enable ion separation and transport, and I will describe these chips and the microfabrication procedures used to make them in Chapter 3.

2.6 Quantum logic using long wavelength radiation

Qubits formed within the hyperfine levels of trapped ions are characterised by transition frequencies in the microwave region of the electromagnetic spectrum. Coherent manipulation may be achieved by directly addressing these transitions with microwave radiation, or alternatively by using laser light by exploiting Raman transitions [104]. Laser driven quantum logic operations have achieved single-qubit [105, 106] and two-qubit gates [106, 107], with fidelities that are well above the threshold required for fault-tolerance. However, complex methods are required to ensure frequency stability, and beams need to be precisely aligned to $\approx 1\,\mu\text{m}$ accuracy to address individual ions. Microwaves have advantages over lasers in that they are simple to generate, have high frequency stability and can naturally cover a wide area, making them attractive for scalable quantum computing implementations that may require millions of ions. Microwave-driven single qubit gates have been demonstrated with gate errors of less than 10^{-6} [41], and it would be advantageous to apply microwave radiation to drive two-qubit gates, thus replacing lasers for quantum logic operations. However, the logic gate schemes developed for trapped ions require a coupling of the internal spin state with the external motional degrees of freedom of the ions that relies upon photon momentum [29, 44], which for microwave radiation is too weak to be effective. This limitation was addressed by Mintert and Wunderlich [74] who showed that a static magnetic field gradient can be used to provide this spin-motion coupling by exploiting transitions between states with different magnetic moments, to thus permit single-qubit

and two-qubit gates to be driven by long wavelength radiation. I will briefly review the photon scheme and then explain the magnetic gradient scheme.

2.6.1 Spin-motion coupling using photon momentum

To perform quantum logic between trapped ion qubits, a method is needed for coupling their internal states in order to perform conditional logic between qubits. The direct spin-spin coupling between adjacent ions is very small, so it cannot be used directly for this purpose. However, ions within a trap interact via their mutual Coulomb repulsion, and it is possible to exploit their collective motional modes of oscillation as an effective 'data bus' to transfer quantum information between ions. This is known as spin-motion coupling and it can be used to perform gates between multiple ion qubits. This approach was first suggested by Cirac and Zoller [29] for two ions in their motional ground states, and extended by Molmer and Sorensen [44] to ions in thermal motion.

Spin-motion coupling is driven by the momentum kick imparted to an ion by a photon, and it requires that the ion experiences a significant variation in the electromagnetic field over the distance defined by it's spatial wavefunction. The strength of the coupling between the internal spin state of the ion and its motional state is indicated by the Lamb-Dicke parameter (LDP),

$$\eta = \frac{2\pi}{\lambda}\sqrt{\frac{\hbar}{2m\omega_z}}, \tag{2.23}$$

where λ is the wavelength of the applied radiation, m is the mass of the ion and ω_z is the ion secular frequency within the harmonic potential of the trap. For the 12.6 GHz hyperfine transition in $^{171}Yb^{+}$ with a typical secular frequency $\omega_z/2\pi \approx 500\,\text{kHz}$ and a 369 nm Raman laser, $\eta \approx 0.2$, whereas for microwaves $\eta \approx 10^{-6}$ and the coupling is negligible [102]. This can be overcome by applying a static magnetic field gradient to the ion, and using qubits formed from states that have different magnetic moments. The resulting state-dependent force can produce a sizable Lamb-Dicke parameter for long wavelength radiation.

2.6.2 Spin-motion coupling using a magnetic field gradient

For qubits formed between two states with different magnetic moments (for example the states $|0\rangle \equiv \left|^2S_{1/2}, F=0\right\rangle$ and $|+1\rangle \equiv \left|^2S_{1/2}, F=1, m_F=+1\right\rangle$) the presence of a magnetic field gradient causes a state-dependent force on the ion, which provides the momentum transfer needed to also drive transitions between the quantized motional states of the ion. This results in an increased spin-motion coupling characterised by an *effective* LDP,

$$\eta_{eff} = \frac{\mu_B \partial_z B}{\sqrt{2}\hbar\omega_z} z_0, \tag{2.24}$$

where $z_0 = \sqrt{\hbar/2m\omega_z}$, μ_B is the Bohr magneton and $\partial_z B$ is the magnetic field gradient in the z-direction [74]. This effective LDP is directly proportional to the gradient of the applied magnetic field, and independent of the wavelength of the applied radiation. This provides a mechanism to generate spin-motion coupling using long wavelength radiation.

The drawback of the magnetic gradient scheme is the requirement for qubit states to have different magnetic moments, which mandates the use of magnetic field sensitive states which are highly sensitive to magnetic field noise, and therefore qubits suffer from short coherence times. This leads to lower gate fidelities and limits the number of gate operations that may be performed in sequence. Instead of using the bare atomic states, we can engineer a magnetically sensitive qubit that is compatible with the magnetic gradient scheme, and which is protected from the dephasing effects of magnetic field noise. This can be achieved by creating superposition states known as dressed states, by applying microwave fields to the ion as discussed in the next section.

2.6.3 Microwave dressed states

A dressed state is an eigenstate of the Hamiltonian describing the dynamics of an ion being driven by resonant electromagnetic radiation. Qubits prepared using the dressed states can have coherence times several orders of magnitude longer than qubits formed from the bare atomic states [108]. Quantum gates using microwave dressed states were first demonstrated in Ref. [109].

To prepare the dressed state system, two microwave fields with equal Rabi frequencies are applied to excite the $|0\rangle \leftrightarrow |-1\rangle$ and $|0\rangle \leftrightarrow |+1\rangle$ transitions, as shown in Figure 2.12. Transforming to the interaction picture and performing the rotating wave approximation (RWA), the system is described by the Hamiltonian,

$$\hat{H}_{MW} = \frac{\hbar\Omega_{MW}}{2}(|+1\rangle\langle 0| + |-1\rangle\langle 0| + H.c.), \tag{2.25}$$

where H.c. is the Hermitian conjugate of the bra-ket terms. There are three eigenstates, or dressed states, of this Hamiltonian, given by

$$
\begin{aligned}
|D\rangle &= \frac{1}{\sqrt{2}}(|+1\rangle - |-1\rangle), \\
|u\rangle &= \frac{1}{2}|+1\rangle + \frac{1}{2}|-1\rangle + \frac{1}{\sqrt{2}}|0\rangle, \\
|d\rangle &= \frac{1}{2}|+1\rangle + \frac{1}{2}|-1\rangle - \frac{1}{\sqrt{2}}|0\rangle.
\end{aligned} \tag{2.26}
$$

Writing the Hamiltonian in the dressed state basis gives,

$$
\hat{H}_{MW} = \frac{\hbar\Omega_{MW}}{\sqrt{2}}(|u\rangle\langle u| + |d\rangle\langle d|). \tag{2.27}
$$

The $|u\rangle$ and $|d\rangle$ states are separated from the $|D\rangle$ state by energy differences of $\hbar\Omega_{MW}/\sqrt{2}$ and $-\hbar\Omega_{MW}/\sqrt{2}$, respectively. The energy level diagram in the interaction picture is shown on the right hand side of Figure 2.12. A qubit constructed from the magnetic insensitive state $|0'\rangle$ and the dressed state $|D\rangle$ satisfies the requirements of the magnetic gradient scheme because the two states carry different magnetic moments, and the qubit is also resilient to the effects of magnetic field noise. To see why this is so, we can treat the magnetic field fluctuations as a perturbation of the form,

$$
\hat{H}_p = \hbar\lambda_0(t)(|+1\rangle\langle+1| - |-1\rangle\langle-1|), \tag{2.28}
$$

where $\lambda_0(t)$ represents an arbitrary time-dependent function.

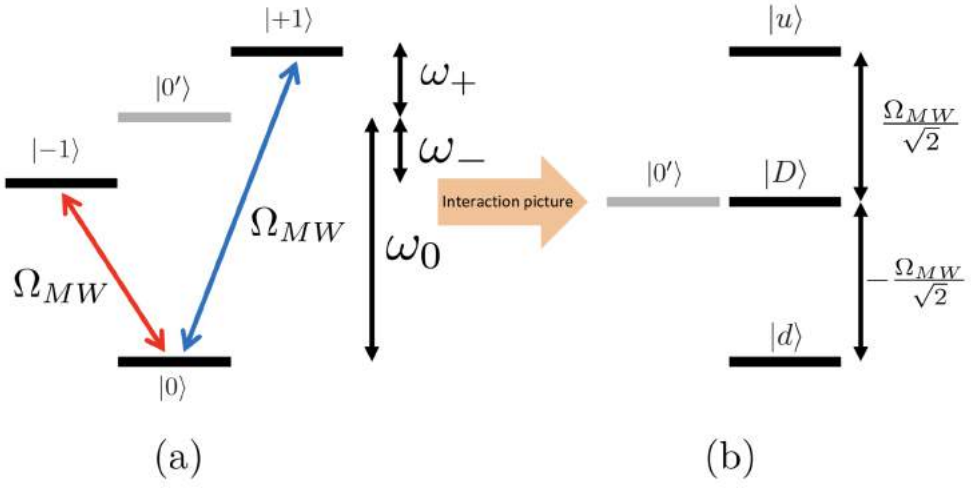

Figure 2.12: Microwave dressed states formed from the bare atomic states of the $^2S_{1/2}$ hyperfine manifold in $^{171}Yb^+$. The energy degeneracy of the $F = 1$ states, $|-1\rangle$, $|0'\rangle$, and $|+1\rangle$ is lifted by an applied magnetic field. Two microwave fields with equal Rabi frequencies Ω_{MW} are applied to the ion (left side of image), which creates three dressed states $|u\rangle$, $|D\rangle$, and $|d\rangle$ in the interaction picture (right side of image). A qubit formed of the $|0'\rangle$ and $|D\rangle$ states is insensitive to magnetic field noise. Figure adapted from [102].

Writing Eq. (2.28) in terms of the dressed states gives,

$$\hat{H}_p = \frac{\hbar\lambda_0}{\sqrt{2}}(|D\rangle\langle u| + |D\rangle\langle d| + H.c.). \tag{2.29}$$

Therefore, magnetic field fluctuations will act to excite transitions between the $|u\rangle$, $|D\rangle$ and $|d\rangle$ states, but the energy separation of $\hbar\Omega_{MW}/\sqrt{2}$ will prohibit any population transfer between the states, unless the frequency is near to $\Omega_{MW}/\sqrt{2}$. This is an important feature of the dressed states, that provides stability against magnetic field noise. The group has performed experiments using a qubit formed by the $|0'\rangle \leftrightarrow |D\rangle$ transition, demonstrating a coherence time of 650 ms, an increase of more than two orders of magnitude compared to qubits constructed from the atomic bare states, which typically have coherence times of 1 ms [73].

The analysis in this section does not need to consider the second-order effects of magnetic field fluctuations on the energies of the bare states because the contributions to the energy-level shifts are small for low magnetic field strengths [108], which is the regime in which our experiments operate. However, for fast and high fidelity gates using the magnetic gradient scheme described in [74] it is advantageous to generate large magnetic field gradients which pose a risk of exposing ions to strong fields outside of the desired low-field regime. To counter this, external compensation coils or on-chip compensation structures are used to reduce the magnetic field strength at the position of the trapped ions, while preserving the highest possible gradient for driving quantum gates.

2.6.4 Gates in the dressed state basis

A simple method for performing single qubit gates in the dressed state basis was proposed by Webster *et al* [110]. This method takes advantage of the unequal splittings, $\omega_+ \neq \omega_-$, due to the second order Zeeman effect, as shown in Figure 2.12. Arbitrary rotations on the Bloch sphere can be achieved for a dressed state qubit formed between the $|0'\rangle$ and $|D\rangle$ states by applying a single RF field with Rabi frequency, Ω_{RF}, resonant with the $|0'\rangle \leftrightarrow |+1\rangle$ transition, and by selecting the appropriate phase, ϕ_{RF}, and detuning. If the Rabi frequency, $\Omega_{RF} \ll |\delta\omega|$ then the transition $|0'\rangle \leftrightarrow |-1\rangle$ will be off resonant and can be ignored.

We can write the Hamiltonian as, $H = H_{\mu W} + H_{RF}$, where the RF terms are:

$$
\begin{aligned}
H_{RF} &= \frac{\hbar\Omega_{RF}}{2}(e^{-i\phi_{RF}}\left|+1\right\rangle\left\langle 0'\right| + H.c.) \\
&= \frac{\hbar\Omega'_{RF}}{2}(e^{-i\phi_{RF}}\left|D\right\rangle\left\langle 0'\right| + H.c.) \\
&\quad + \frac{\hbar\Omega'_{RF}}{2\sqrt{2}}(e^{-i\phi_{RF}}(\left|u\right\rangle + \left|d\right\rangle)\left\langle 0'\right| + H.c.)
\end{aligned}
\tag{2.30}
$$

in the dressed state basis where $\Omega'_{RF} = \Omega_{RF}/\sqrt{2}$ and we have used the RWA.

Under the condition that $\Omega_{RF} \ll \Omega_{MW}$ then the transitions $|0'\rangle \leftrightarrow |d\rangle$ and $|0'\rangle \leftrightarrow |u\rangle$ are suppressed and the RF field drives single qubit gates between the $|0'\rangle \leftrightarrow |D\rangle$ dressed state qubit only, with a Rabi frequency, Ω'_{RF}. The method can be extended to two-qubit gates, using a similar approach to the one proposed by Timoney *et al* [109].

Within the group, experiments have used microwave dressed states in the hyperfine manifold of $^{171}Yb^{+}$, combined with the magnetic gradient scheme described earlier, to implement a two-qubit entangling gate between ions in a Paul trap with a Bell state fidelity of 0.985(12) [73]. An effective LDP of $\eta_{eff} = 0.0041$ was achieved with an applied magnetic field gradient of $23.6 \pm 0.3\,\mathrm{T/m}$ [73, 102]. The high fidelity of the gate is close to the threshold required for fault tolerance. The fidelity can be improved by increasing the magnitude of the magnetic field gradient applied to the ion, which increases the size of the LDP. By moving to microfabricated ion traps (described in Chapter 3) in which the ions are positioned significantly closer to the current-carrying structures used to generate high magnetic field gradients, a further increase of the LDP is expected to achieve gate fidelities well above the threshold for fault tolerance.

2.6.5 Scalable many-qubit operations

A significant challenge to the realisation of a large-scale quantum computer is the sheer number of physical qubits needed to address real-world problems of interest. As an example, the breaking of RSA encryption may require millions of individual ions [111, 34]. Logical operations performed in parallel therefore require millions of laser fields or microwave fields to drive gates between qubits, which is a considerable obstacle to the implementation of a large-scale quantum computer. However, in the architecture described in [34], which was introduced in Chapter 1, the number of fields only needs to be equal to the number of different types of gate that must be carried out in parallel. This is a drastic simplification of the practical requirements for a many-qubit device. In this architecture, ions within an array of microfabricated traps may be transported into gate zones where a current is applied to generate a magnetic field gradient. Within each zone, the positions of ions within

the gradient can be controlled by DC electrodes so that the local field experienced by each ion effectively tunes the Zeeman splitting of the microwave dressed state qubit into or out of resonance with a handful of global fields applied to the entire array. In this way, many single-qubit and two-qubit gates can be driven at once, while other ions that must be left unchanged can be detuned from the fields so that there is no interaction. Long wavelength radiation is well suited to this approach because of the wide area that may be covered and the inherent frequency stability in comparison to lasers. Finally, the architecture uses a modular approach to address scalability by connecting surface trap chips using electric fields and ion shuttling.

2.7 Summary

This chapter has described the fundamental operating principle of the Paul trap, and how microfabricated surface traps may be used as the elementary building blocks of a scalable quantum computing architecture using trapped ion qubits. I have described the techniques employed for ionisation of neutral ytterbium atoms, and laser Doppler cooling of ions in a trap to reduce their kinetic energy. I have explained how the hyperfine ground state of $^{171}Yb^{+}$ can be used to construct a qubit, and how a magnetic field gradient can be used to couple the internal and motional states of ions to enable quantum logic operations using long wavelength radiation. I have discussed the benefits of microwave dressed states for constructing qubits that are robust to the decoherence effects caused by magnetic field noise, and how gates may be performed in the dressed state basis. At the end of this chapter I explained how the overall scheme may be used to realise a scalable approach to quantum computing, based on a modular architecture with ion-traps connected using electric fields, and logic operations driven by long wavelength radiation in combination with an applied magnetic field gradient.

Chapter 3

The two-module ion trap processor

3.1 Introduction

In this chapter I discuss the alignable ion-trap modules that were designed to operate as a single quantum processor connected by electric fields. I start by providing an overview of the two-module setup. I then describe each of the component systems that make up the modules, beginning with the specialist ion trap chips developed for the experiment, and the microfabrication processes employed in their construction. I provide an analysis of the trapping potential above the surface of the aligned modules in order to evaluate the trap depth, pseudopotential barrier and motional heating of trapped ions due to trap separation in three axes. Following this analysis I introduce the in-vacuum positioning system that was developed for module alignment. I explain how the requirements of the system were met by a combination of solutions including a manually-adjusted pre-alignment stage for initial set-up of the modules, a piezo-based three-axis translation stage for in-vacuum alignment, and an imaging system used for the measurement of module displacements and for qubit state detection. Finally I describe the in-vacuum printed circuit boards of each module that were designed as a compact interface between the ion trap microchips and external voltage and current sources, with on-board filtering to reduce voltage noise.

3.2 The two-module ion trap setup

The main components of the two-module ion trap processor are shown in Figure 3.1. Each module incorporates a specialised surface-electrode ion trap microchip on which the RF rails and DC electrodes extend up to an alignable edge. Electrical connections are made between the chip and an in-vacuum printed circuit board (PCB) that surrounds each chip on three sides, leaving the alignable edge free of obstruction. The chips are bonded to

oxygen-free high thermal conductivity (OFHC) copper blocks which provide an efficient thermal connection to a cryogenic heat sink mounted directly to module 1. To permit free movement, the heat sink connection to module 2 is made using a flexible OFHC copper braid. Cold helium gas is circulated through the heat sink, supplied by a closed-cycle cryogenic cooling system (described in Chapter 4). This is required to dissipate heat generated by on-chip magnetic gradient coils which carry high currents.

To bring the ion trap modules into alignment, the experiment uses an in-vacuum positioning system based on a UHV-compatible 3-axis piezo translation stage. This is used to align module 2 with respect to module 1 which remains stationary. A manually-adjustable pre-alignment stage allows for the modules to be set up with an initial separation between the chips to mitigate for thermal expansion or contraction during baking treatment and

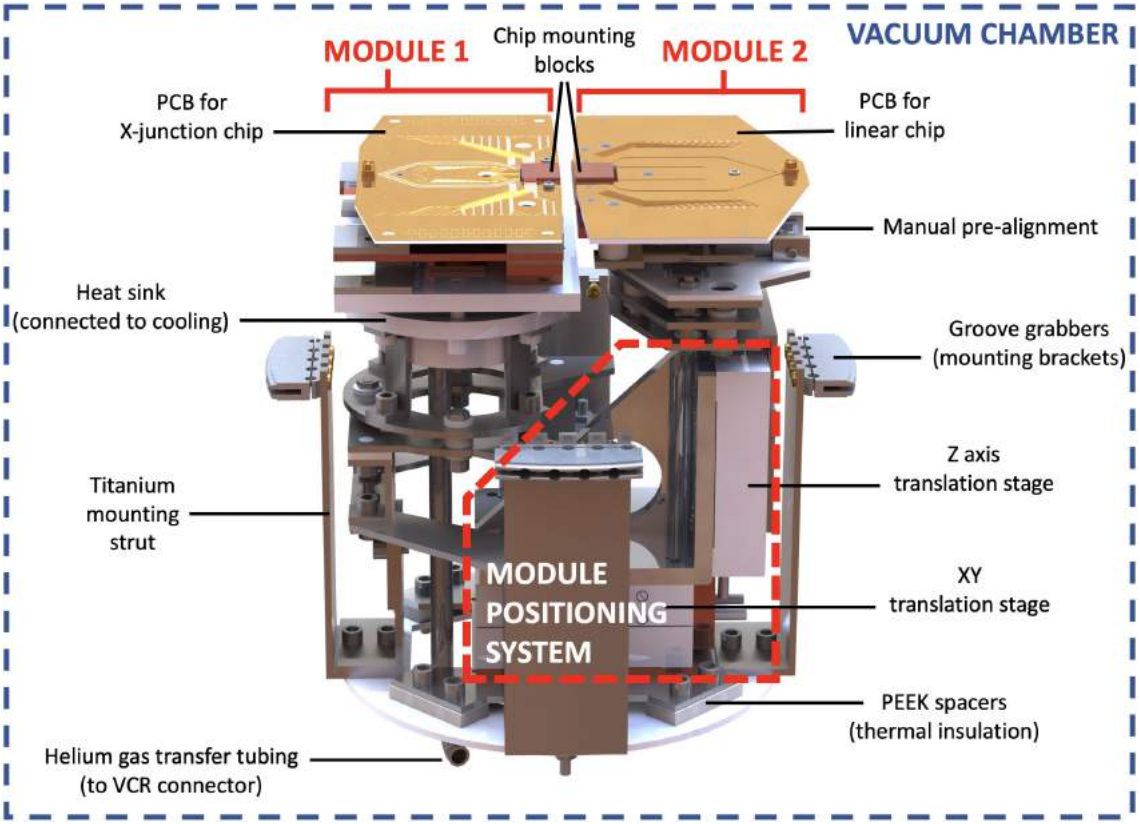

Figure 3.1: The two ion trap modules (and in-vacuum mounting structure) of the scalable demonstrator experiment. Specialist surface trap chips featuring alignable edge-electrodes are attached to copper mounting blocks at the top of each module. Compact PCBs with on-board low-pass filtering deliver trapping potentials to the chips and high-currents for on-chip magnetic gradient coils. Module 1 interfaces directly with a cryogenic heat sink, which connects to Module 2 via a flexible copper braid. Module 2 is mounted onto a piezo-based positioning system. For ease of reference, Module 1 is referred to as the 'heat sink' side of the experiment, while Module 2 is referred to as the 'piezo' side. Multiple stages of thermal insulation are used to isolate the heat sink and the ion trap modules from the walls of the surrounding vacuum chamber which remains at room temperature.

cooling of the experiment that may cause the chips to move into contact. All elements shown in Figure 3.1 are supported within a high stability, in-vacuum mounting structure using a titanium alloy[2] that provides high mechanical strength, low thermal conductivity and a very low coefficient of thermal expansion. This is achieved using multiple levels separated by spacers made of poly-ether-ether-keytone (PEEK). The complete structure was designed to mount onto a CF160 flange within the vacuum chamber described in Chapter 4, using four 'groove-grabber' clamps[3] for a robust mechanical connection. The mounting structure is designed to provide a large thermal resistance between the ion trap chips and the room-temperature walls of the vacuum chamber to reduce conductive heat transfer from the environment.

3.3 Ion trap microchips

Two generations of surface trap microchips were developed for the experiment: the first generation include linear and X-junction geometries with up-to-edge electrodes on one side to permit close alignment with an adjacent chip. The second generation of microchips use a similar electrode geometry to the first generation, but additionally feature magnetic gradient wires embedded into the silicon substrate below the ion-trap layer. This adds the capability for quantum logic operations to be performed using the microwave gate scheme described in Chapter 2. Experiments were performed using both generations of microchips. Details of each generation are described in the following sections.

3.3.1 First generation ion traps

The first generation of microchips were fabricated with linear and X-junction geometries as shown in Figure 3.2 and Figure 3.3. The designs permit ion trapping on multiple zones determined by the potentials applied to segmented DC endcap electrodes positioned either side of the RF rails. Both the linear and X-junction geometries were designed for an ion height of 150 µm, which is determined by the width and separation of the RF electrodes. Each RF electrode is 360 µm wide and the separation between their inner edges is 110 µm. Four rotation electrodes are positioned inside and outside of the RF rails (two per RF rail). A DC potential applied to these electrodes rotates the principal axes of the pseudopotential and allows Doppler cooling of all ion motional modes using a single laser beam directed parallel to the chip surface. The chips are constructed from a surface layer of gold deposited

[2] Grade 5 Ti6Al4V

[3] Kimball Physics 8.00" Groove Grabber Split Axial Clamping, Model no. MCF800-GrvGrb-C01

onto a silicon substrate and patterned by lithography, as detailed in Section 3.3.4.

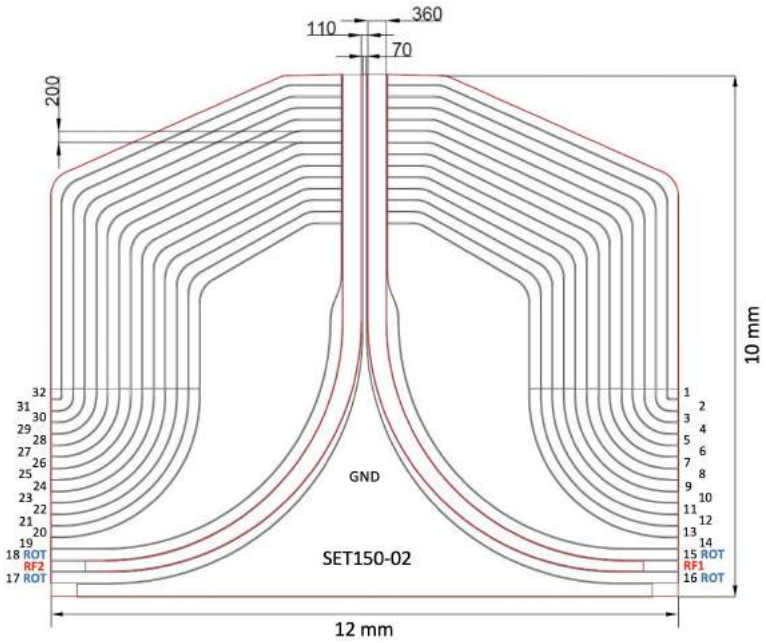

Figure 3.2: Electrode layout for the SET150-02 linear surface-trap microchip. Dimensions shown without units are in micrometres. The design has 32 DC control electrodes and two RF rails with inner and outer DC rotation electrodes to enable Doppler cooling of all vibrational modes using a single laser beam. The top corners of the chip are angled back from the edge to avoid potential obstructions during close alignment of RF electrodes on adjacent chips. Additionally, the silicon substrate at the edge was etched back by 50 µm to ensure that the electrodes overhang slightly.

Additional fabrication considerations constrain the design of the ion traps. All electrical connections to the trap electrodes are routed on the chip surface, which precludes any island-type electrodes, such as inner segmented control electrodes between the RF rails. These types of electrodes provide increased control over the shape of the potential well for operations such as ion swapping, but require advanced microfabrication processes such as through-silicon vertical interconnect access (VIA) connections which pass through the dielectric substrate. In making the decision to exclude VIAs, the complexity of the microfabrication process was reduced. Consequently, ion transport is controlled entirely by outer segmented DC electrodes which can be connected by wire bonds at the edge of the chip. Another consideration was laser access to the ion location, which required that wire bond connections were confined to the bottom section of the chip and away from the trapping region, so that the bond wires are not placed into the path of the laser beam. To reduce RF power loss within the silicon substrate the microchips feature a buried ground plane 8 µm below the surface electrodes to shield the substrate from the RF potential applied to the rails. The top corners of the linear chip were removed (as shown in Figure 3.2) to

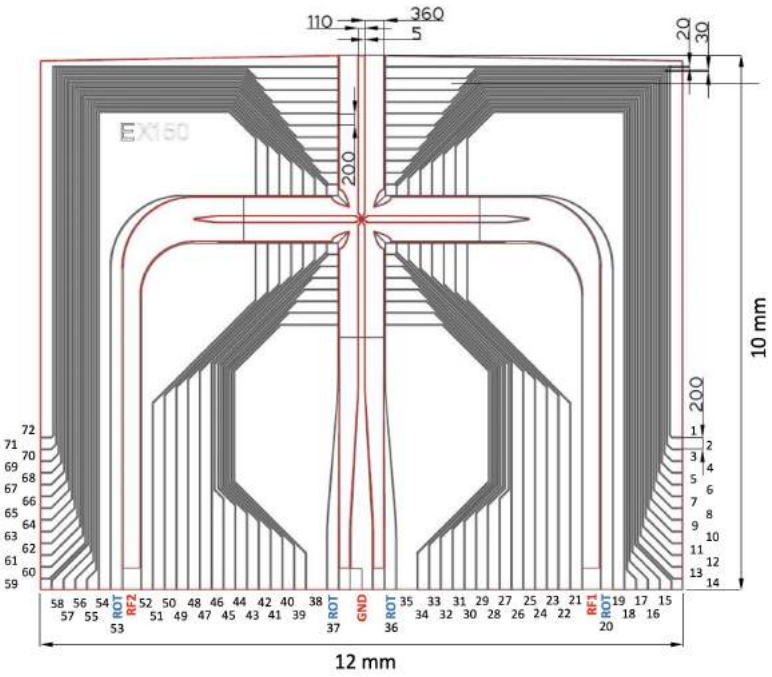

Figure 3.3: Electrode layout for the EX150 surface-trap microchip with X-junction geometry. This trap has two RF rails that split to form the arms of the X-junction, and 72 DC electrodes of which 68 are available for ion shuttling and four are used for principal axis rotation. At the alignable edge of the chip the corners slope away from the centre at an angle of 1 degree to aid alignment of the RF electrodes on adjacent chips. To further avoid obstruction at the edge, the substrate below the electrodes was etched back by 50 µm causing the electrodes to overhang slightly.

allow for close alignment of the RF electrodes of adjacent chips, and avoid problems due to rotational misalignment. The silicon substrate below the RF electrodes at the edge of the trap is etched away by 50 µm to form an undercut below the electrode layer. This effectively shields the ion from stray fields when in the gap region due to charge build-up on the dielectric substrate.

3.3.2 Second generation ion traps with magnetic gradient coils

The magnetic gradient traps incorporate microfabricated current carrying wires (CCWs) beneath the trap electrodes (the buried ground plane sits between the electrodes and the of the trap, which are otherwise constructed using the same methods as the single layer traps described above. Figure 3.4 shows the design of an EL125 magnetic gradient chip with a linear trap geometry. Figure 3.5 shows an ME125 magnetic gradient chip with an X-junction design. The magnetic gradient traps were made with both linear and X-junction trap designs with an ion height of 125 µm. The ion height was chosen after comparing simulations of the maximum achievable magnetic field gradient at the ion position, against the expected ion heating rate due to the proximity of the ion to the trap surface. This work

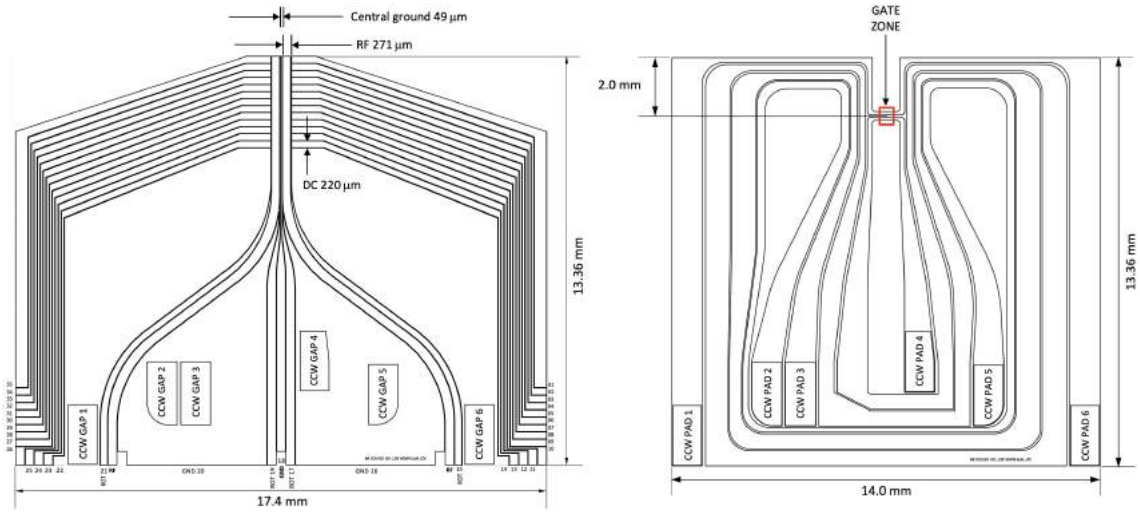

Figure 3.4: EL125 second-generation linear surface trap with microfabricated magnetic gradient coils below the surface electrodes. (*Left*) The trap has 2 RF electrodes separated by a central ground electrode, and 35 DC electrodes including four dedicated to rotation of the principal axis of the trap potential. The ion height is 125 µm. (*Right*) The magnetic gradient coils are designed to generate a magnetic field gradient in excess of $100\,\mathrm{T\,m^{-1}}$ at the ion height in the centre of the gate zone (highlighted by red square), using approximately 10 A of current. The coils are structured to nill the magnetic field at the centre of the gate zone. Six miniature, gold-plated beryllium-copper clamps are used to make the electrical connections to the current carrying wires (CCWs) of the gradient coils.

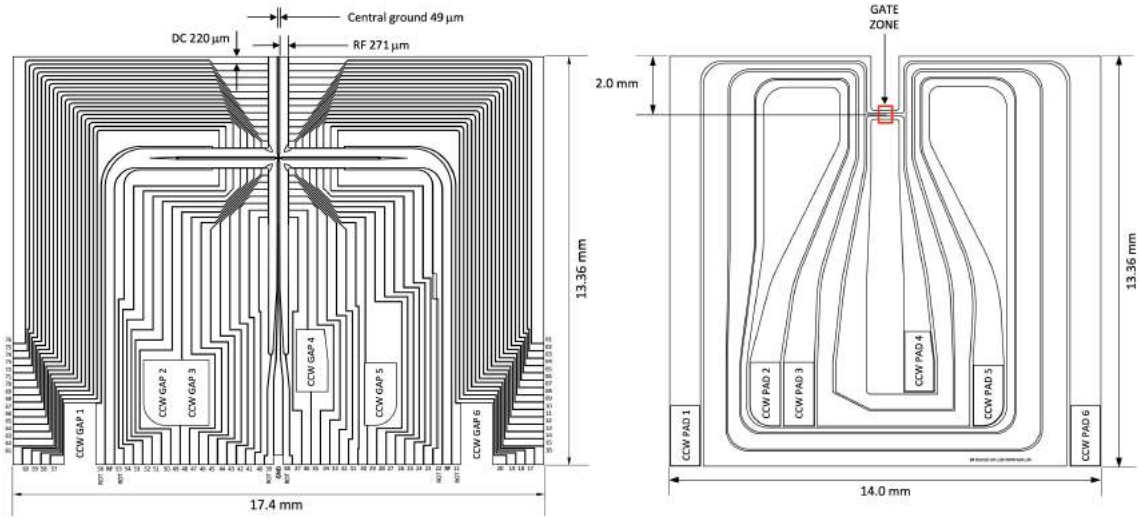

Figure 3.5: ME125 second-generation surface trap with X-junction geometry and integrated magnetic gradient coils. (*left*) The 2 RF electrodes split to form the cross-shaped junction, with a cross-shaped central ground electrode through the centre. The trap has a total of 76 DC electrodes of which 4 are dedicated to rotation of the principal axes of the trap potential. (*right*) The microfabricated magnetic gradient coils are located below the surface electrodes, and are identical to the coils shown in Fig 3.4 for the EL125 chip.

was conducted by Raphaël Lebrun-Gallagher, building upon previous analysis by David Murgia [90]. The coils consist of three separate CCW loops and each loop has two pads for electrical connections (a total of six pads). At the gate zone, two short parallel wires pass anti-parallel currents which generate a high magnetic gradient at the ion height above the chip surface. The gate zone is located on the trap axis at a distance of 2 mm from the alignable edge of the chip. The remaining structure of the coils outside of the gate zone act to nill the magnetic field at the ion height. Simulations performed by Raphaël Lebrun-Gallagher have shown that for a current of 10 A the coils generate a magnetic gradient in excess of $100\,\mathrm{T\,m^{-1}}$ at the ion position above the gate zone.

3.3.3 High current connections for magnetic gradient coils

For experiments using second generation chips, the aluminium bond wires are unsuitable for the high current connections required for the magnetic gradient coils, which operate using currents of ~ 10 A. Therefore a set of gold-plated beryllium copper (BeCu) clamps were constructed for the experiment based on a design by Raphaël Lebrun-Gallagher. The choice of BeCu was driven by the requirements for vacuum compatibility and high electrical conductivity. The clamps are secured in place by M1 screws on the PCB, and have arms that extend onto pads on the surface of the chip to provide an increased cross-sectional area compared to wire bonds and thus a lower resistance to reduce the power dissipated over the connections. A 3 µm gold plating reduces contact resistance at the interface between the clamp and the pad.

3.3.4 Microfabrication procedure

To produce the ion trap chips, layers of gold (Au), silicon oxide (SiO_2) and silicon nitride (Si_3N_4) were deposited onto the surface of a silicon substrate using plasma enhanced chemical vapour deposition (PECVD) and patterned using photolithography and deep reactive ion etching (DRIE). Both generations of microchips share a common upper layer structure that consists of gold surface electrodes above a dielectric layer and a buried ground plane, as shown in Figure 3.6. These upper layers are constructed using the procedure shown in Figure 3.7.

For the magnetic gradient chips, an additional set of steps were required to first fabricate the copper CCWs in the silicon substrate before building the upper layers of the trap. These additional steps are shown in Figure 3.8. A DRIE process was used to etch the dimensions of the CCWs into the silicon wafer, and then copper was electroplated to fill the space.

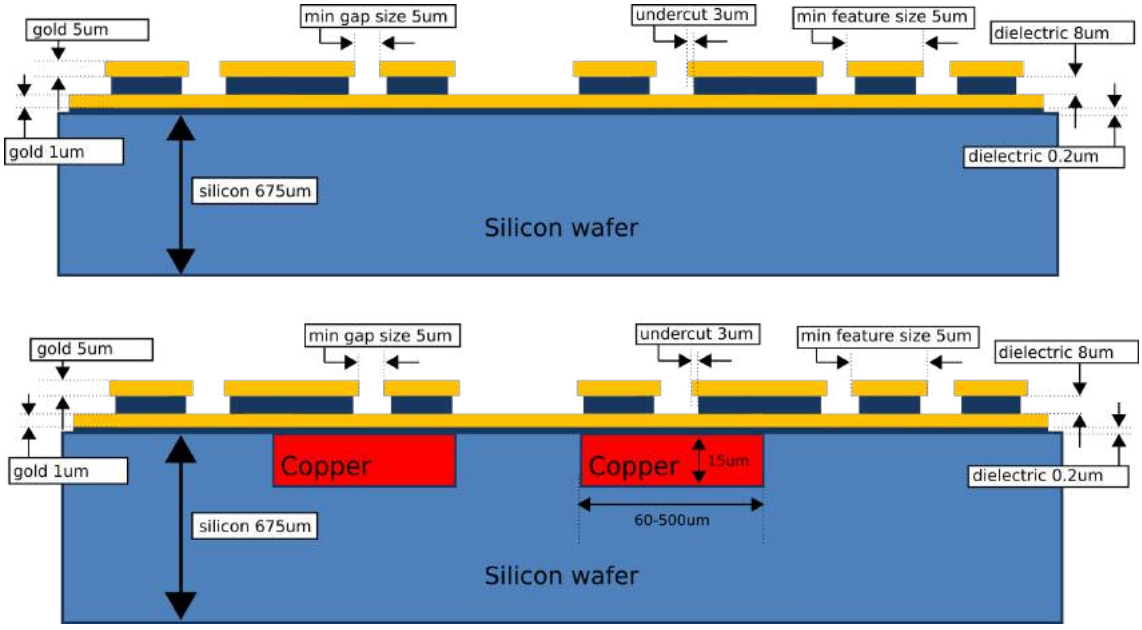

Figure 3.6: Microfabricated layer structure of the ion-trap chips. (*Top*) first generation trap. (*Bottom*) second generation trap with embedded copper wires for generating high local magnetic field gradient for multi-qubit gates.

After this, a chemical mechanical polish was used to remove excess copper and achieve a smooth, flat surface on top of which to fabricate the upper layers of the ion-trap chip.

After microfabrication, the surface trap chips were received on 6 inch silicon wafers, which were then diced using a diamond saw to produce individual chips. To create the clean edge for surface trap alignment, the diced chips were plasma etched using SF_6 and XeF_4. This resulted in a smooth, accurate edge with roughness of $< 1\,\mu m$ as shown in the microscope image of Figure 3.9.

3.3.5 RF breakdown tests

Typically the RF trapping voltages applied to the microchips are required to have an amplitude of Vpeak $\sim 200\,V$ [112]. To understand the maximum RF voltage that the chips can sustain, electrical breakdown tests were performed using an SET150-02 chip, which was fabricated using the techniques described above. To conduct the tests, the chip was wire bonded to a printed circuit board and mounted in a vacuum chamber that was evacuated to a pressure of 10^{-5} mbar. At this pressure electrical arcing through the air is prevented and instead breakdown occurs through the silicon oxide and nitride layers of the chip that insulate the surface electrodes from the underlying ground plane. A voltage signal at 25.25 MHz was applied to the chip, starting with an amplitude of Vpeak $= 50\,V$, and this was increased steadily while imaging the chip using a EMCCD camera to observe

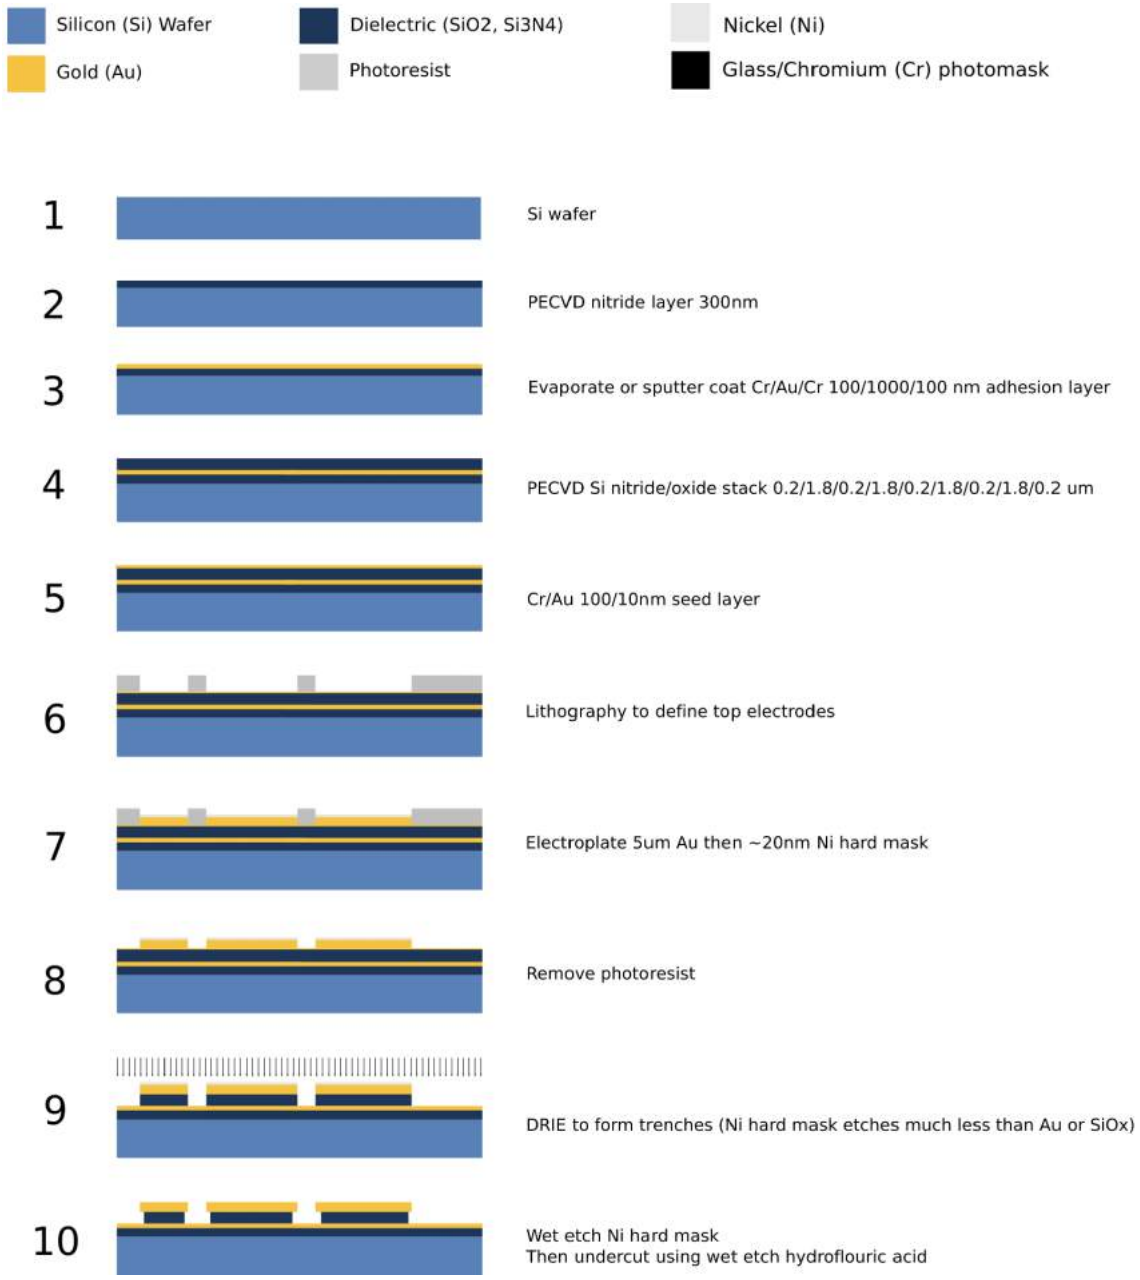

Figure 3.7: Microfabrication workflow used to construct the surface trap microchips (both linear and X-junction geometries).

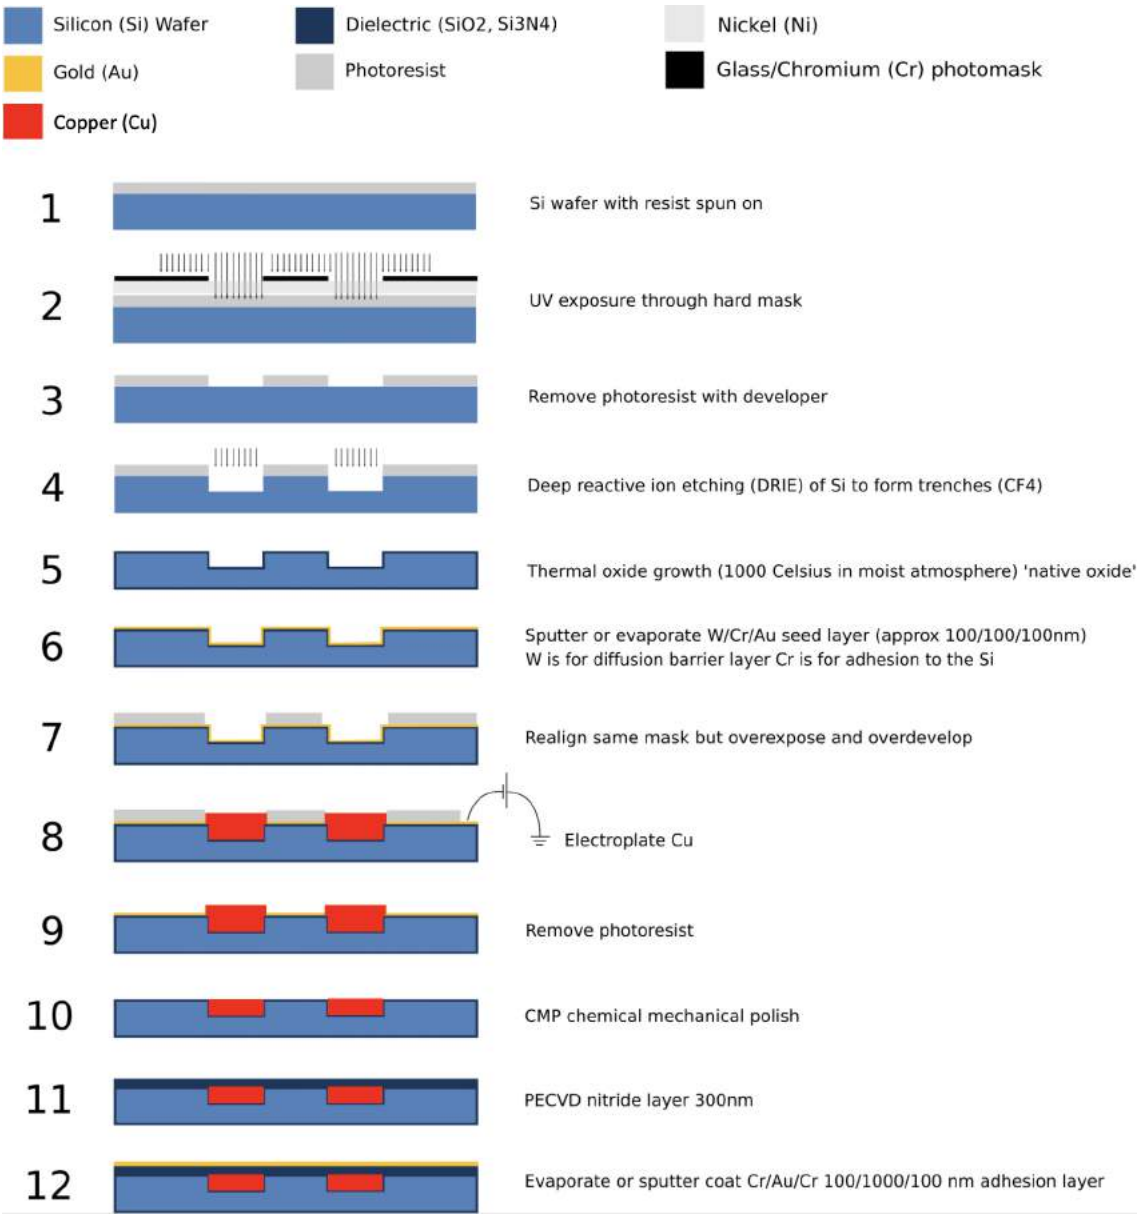

Figure 3.8: Microfabrication steps for constructing copper wires in the silicon substrate of a magnetic gradient microchip.

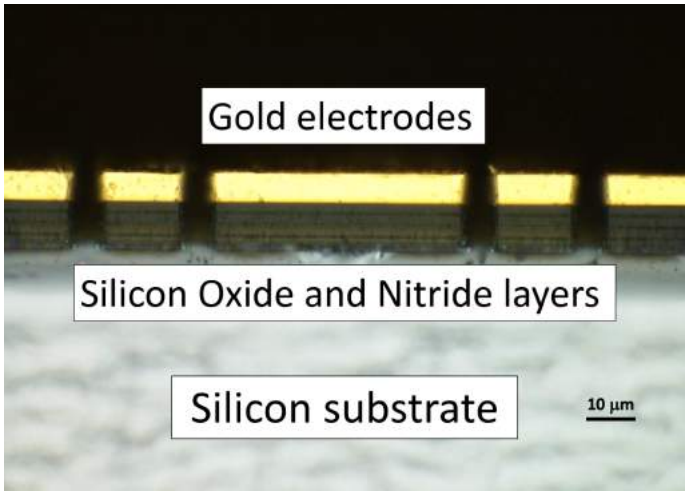

Figure 3.9: Microscope image at 100× magnification showing the edge of a SET150-02 surface-trap microchip. The surface layer of gold electrodes is supported by a stack formed of alternating layers of silicon oxide and silicon nitride, deposited onto a silicon substrate.

sparks due to electrical arcing. The first detection of arcing was at Vpeak $\sim 360\,\mathrm{V}$, which was attributed to breakdown of small fragments or filaments of material trapped within the trenches between the electrodes. Intermittent arcing of this type was observed as the voltage was increased further, until the electrodes shorted to the underlying ground plane due to complete breakdown through the substrate at a voltage of Vpeak $= 1.1\,\mathrm{kV}$.

3.4 Analysing the effect of ion trap separation on trapping potential

Numerical simulations of the two-module processor were conducted to investigate how a misalignment between the ion-trap modules affects the trapping potential in the gap region between chips, and the consequences on ion dynamics. These simulations were performed by Zak Romaszko [63] using COMSOL[4], a commercial software package for Finite Element Method (FEM) simulation. In this method the electric field is calculated for a three-dimensional mesh of points within a volume of space, where the potential at neighbouring points is related by low-order (linear or quadratic) functions evaluated using an iterative solver [113].

The ion traps were each modelled as being identical, having a five-wire geometry consisting of a central ground electrode surrounded by two RF electrodes and two DC electrodes within a single layer. For each trap a ground plane layer was modelled at a depth of $8\,\mu\mathrm{m}$ below the electrodes. Figure 3.10 shows the simulation geometry and defines the reference

[4]COMSOL release 5.4 (https://www.comsol.com/release/5.4)

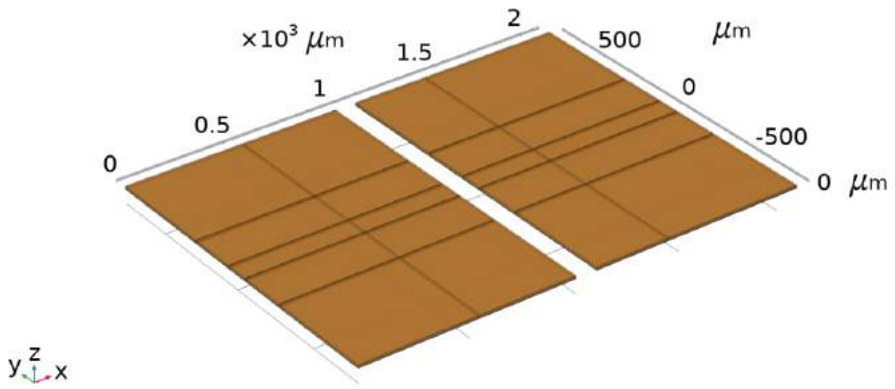

(a) Traps displaced in x direction only.

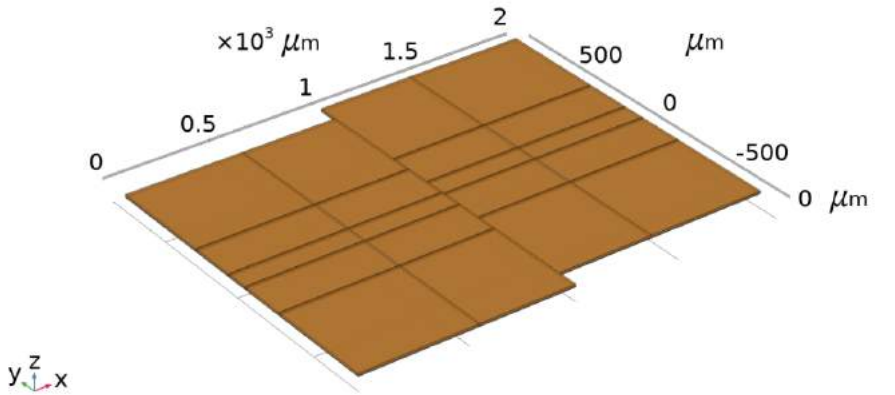

(b) Traps displaced in y direction only.

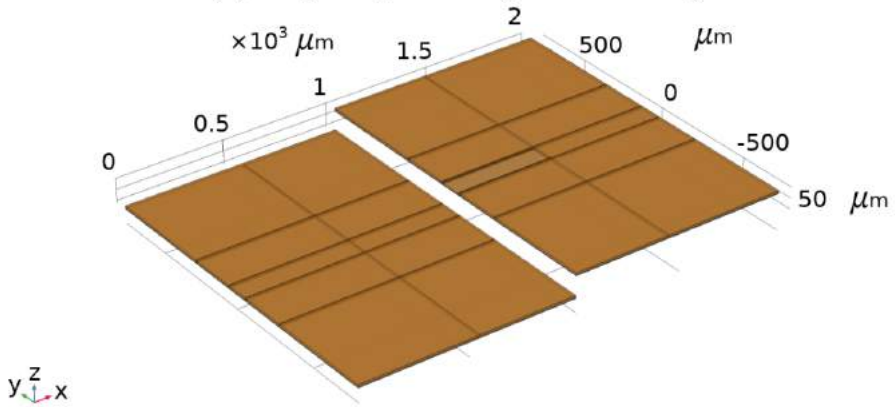

(c) Traps displaced in z direction only.

Figure 3.10: Definition of the reference axes used in simulation to specify trap displacement. A displacement of $x = y = z = 0\,\mu m$ represents perfect alignment of the trap electrodes with no gap between chips [114].

axes used to specify trap displacements. The axial direction of the traps is along the x-axis.

3.4.1 Analysis parameters

The simulations were used to calculate the electric potential above the surface of the traps when driven at equal frequencies of 15 MHz at a voltage amplitude of $V_{RF} = 200\,\mathrm{V}$. The RF signals are assumed to be exactly in-phase for simulation purposes. In practice it is possible to apply dissimilar frequencies to the traps, or to have a non-zero phase difference between the RF signals. This is undesirable since it results in an RF barrier forming between the chips, with non-zero pseudopotential at the ion position. This is qualitatively the same effect as described in this section due to a separation between traps.

The pseudopotential was calculated by first modelling a static voltage of 1 V applied to the RF electrodes, with DC electrodes set to 0 V. This gives a base potential which can be scaled by the peak voltage, V_{RF}, using the principle that the electric field scales linearly with voltage. The pseudopotential was then calculated using the time-independent approximation

$$\Phi_p(\boldsymbol{r}) = \frac{e^2}{4m\Omega_{RF}^2}(V_{RF}\boldsymbol{\nabla}\widetilde{\Phi}_{RF}(\boldsymbol{r}))^2, \tag{3.1}$$

where Ω_{RF} is the RF frequency, $\widetilde{\Phi}_{RF}(\boldsymbol{r})$ is the potential due to a static voltage of 1 V applied to the RF electrodes, and for m the mass of the $^{171}\mathrm{Yb}^+$ ion was used. By solving for the minimum of the potential at a series of points along the axial direction, the ion position was determined and the axial pseudopotential barrier was calculated. In this section all potentials represent energy in meV (and so carry an additional factor of e, the charge on the ion, compared to electric potentials measured in V).

Ideally within a trap the pseudopotential at the ion position will be zero but where a barrier occurs it is non-zero and the gradient of the barrier contributes to motional heating of the ion. The heating rate is given by equation 2.18, which includes a term $\left[\frac{\partial}{\partial z}E_0^2(z)\right]^2$ to represent heating due to the pseudopotential gradient. Although logic gates are, to some extent, insensitive to the initial motional excitation of trapped ions [44], the time cost associated with applying Doppler cooling or sympathetic cooling of ions contributes to decoherence of the qubit state and reduces gate fidelities [115].

3.4.2 Analysis of the pseudopotential barrier

To investigate the pseudopotential barrier, I have analysed the output from additional simulations performed by Martin Siegele to reflect the electrode geometry of the second

generation ion traps (as detailed in section 3.3.2).

Figure 3.11 shows the pseudopotential barrier along the trap axial direction (x direction) for separations of 10 µm, 20 µm and 30 µm in the x direction. For each case the separation in the y and z directions was set to 0 µm. The barrier is symmetric about the centre position between the traps, and at the centre position there is a nill in the pseudopotential.

The axial barrier due to trap separation in the y direction follows a qualitatively similar shape to Figure 3.11 and remains symmetric about the centre position between the chips. In contrast, a separation in the z direction leads to an asymmetric barrier as shown in Figure 3.12. For this modelling the separation in x was fixed at 0 µm, with no separation in y, and the separation in z was simulated at 0 µm, 5 µm and 10 µm. The larger barrier peak occurs above the surface of the lower trap, while the upper trap has a lower barrier peak. The location of the pseudopotential nill is displaced in the direction of the higher trap as separation increases in the z direction.

3.4.3 Analysis of trap depth, barrier and heating rate

Next, the simulations were used to investigate the effect of trap separation in the x, y and z axes on the trap depth, pseudopotential barrier and ion heating rate. Separation was

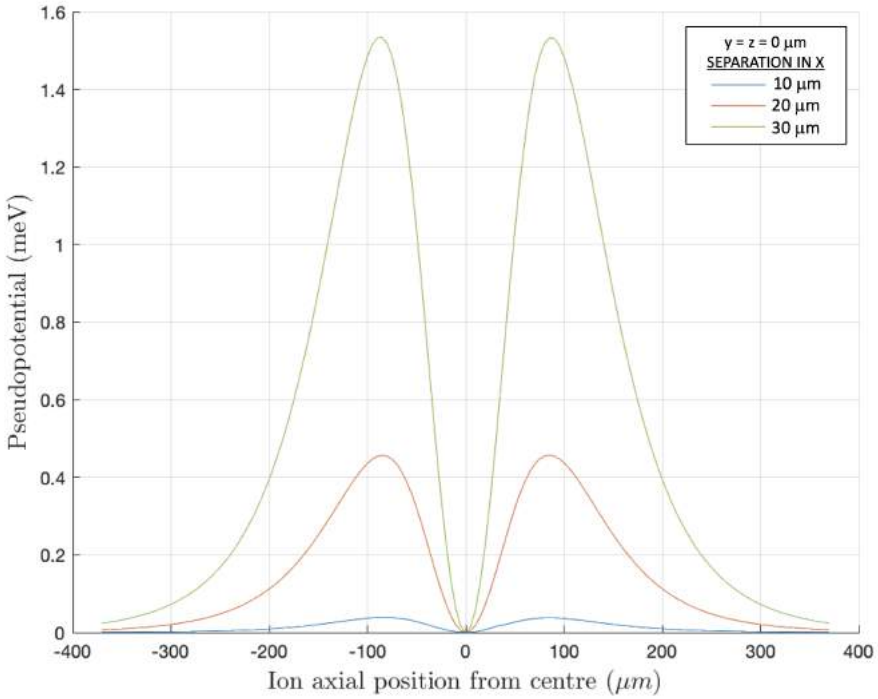

Figure 3.11: Simulation of the pseudopotential barrier along the x-axis (trap axial direction) when the trap separation in the x direction was set at 10 µm, 20 µm and 30 µm. In each case there was no separation in the y and z directions.

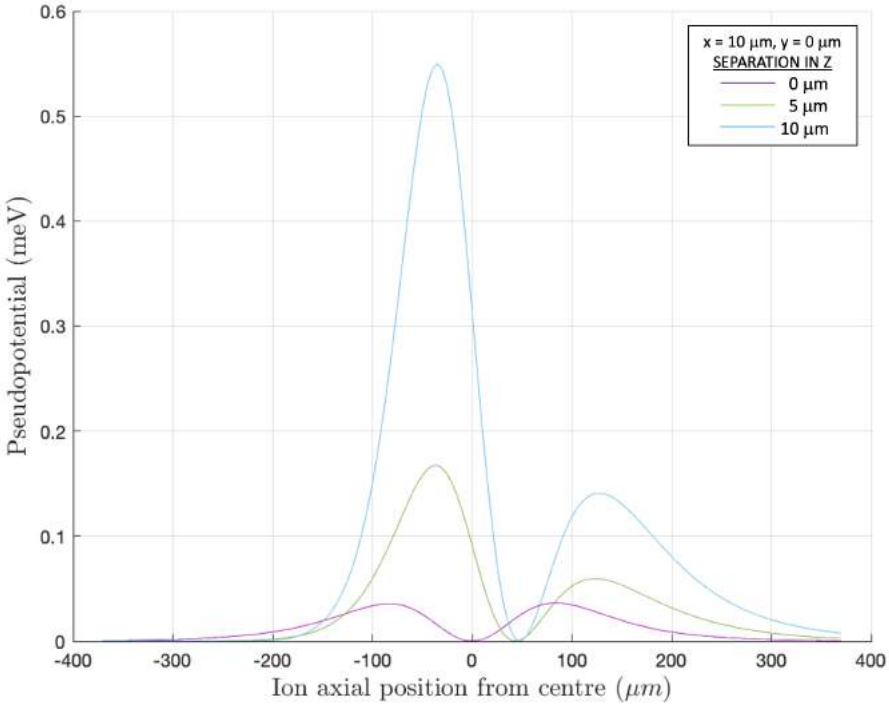

Figure 3.12: Simulations of the axial pseudopotential barrier for trap separations of 0 µm, 5 µm and 10 µm in the z direction. In each case the separation in x was fixed at 10 µm, with no separation in y. The asymmetry of the barrier becomes more pronounced as the separation in z increases, and the location of the central nill is displaced in the direction of the higher of the two traps.

increased for each axis individually, whilst maintaining a zero offset for the perpendicular directions. The trap depth is defined as the energy difference between the pseudopotential minimum and the escape energy of the trap. For this analysis trap depth was calculated at the location of maximum barrier in the axial pseudopotential. When the traps are misaligned, the local axial direction (and ion trajectory) within the gap region is defined by positions along the minimum of the radial potential, where the ion potential energy is minimised.

Figure 3.13 shows the effect of the separation on trap depth. The initial value of the trap depth with perfect alignment between chips is 226 meV. The results show that increasing separation in the x and z axes causes a decrease in trap depth, as expected. For small separations of less than 5 µm in the x and z axes the trap depth remains unchanged. However, separation in the y axis leads to a small apparent increase in trap depth. This result is consistent with previous findings that the trap depth is sensitive to the geometrical ratio of the RF width to the ground width [91]. Therefore, as the misalignment in the y direction increases, the ion is effectively exposed to a narrower central ground electrode at the interface between chips, and the dimension ratio of $\frac{RF}{GND}$ increases. The trap depth is

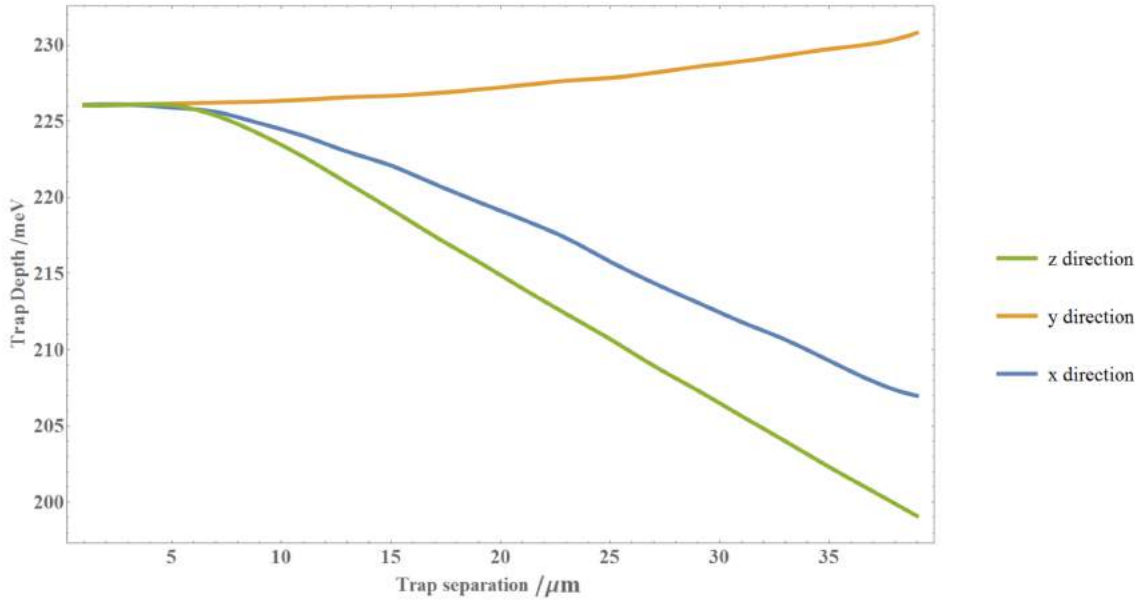

Figure 3.13: Simulation results showing trap depth as a function of trap separation in the x, y and z directions [114].

most sensitive to separation in the z direction, for which the steepest gradient occurs. This is because a rise in the z height of one chip relative to the other exposes the edge of the underlying ground plane located 8 µm below the surface electrodes, forming a RF-ground-RF structure in the axial direction. Adjustments to the microfabricated edge structure may be effective in reducing the barrier height due to this grounding effect. Potential solutions include increasing the undercut below the electrodes at the edge of the trap to limit the fringing field from the RF rails to ground, or alternatively forming a gold surface over the side of the trap edge to effectively extend the electrodes over the edge and shield both the exposed ground and dielectric layers.

There is a maximum separation at which the trap depth falls to zero and so the ion can no longer be confined by the potential. For displacement purely in the x direction, and assuming perfect alignment in the other axes, zero trap depth was found to occur at a trap separation of 230 µm. From a practical perspective this is far from a reliable measure of capability since trapping lifetimes would reduce to become impractical before this limit was reached. Nevertheless it is an interesting metric for the alignment constraints of the two-module processor, which are not overly stringent; commercial piezo actuators are available with at least $\sim$ few μm precision or better (such as the stages described in section 3.5.2).

Figure 3.14 shows how increasing trap separation affects the axial RF pseudopotential barrier. For x and y displacements of less than 5 µm the simulation results have artefacts due to the mesh size of 1 µm which affects the resolution of the output. However, the results

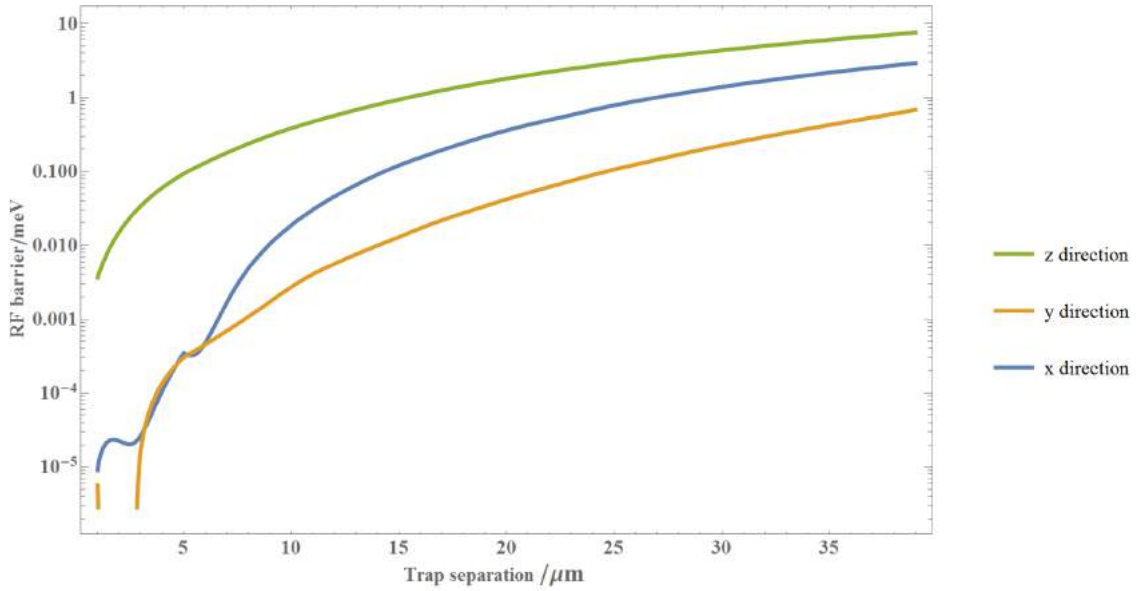

Figure 3.14: Simulation results showing axial pseudopotential barrier as a function of trap separation in the x, y and z directions [114].

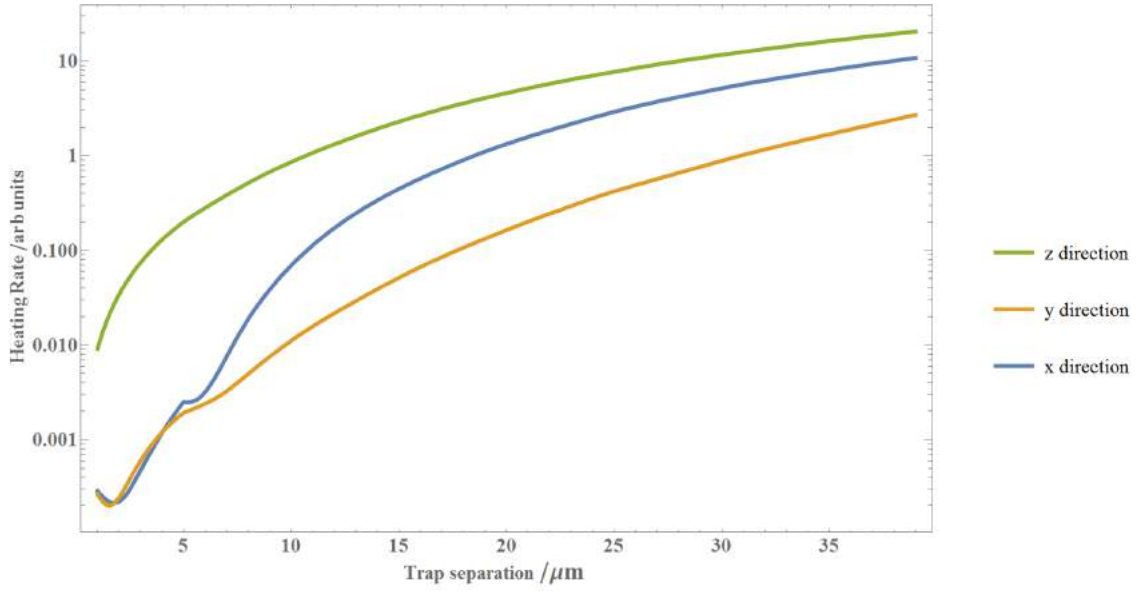

Figure 3.15: Simulation results showing the ion heating rate due to a pseudopotential gradient along the axial direction, as a function of trap separation in the x, y and z directions [114].

show that for separations up to a few tens of μm in any one axis, the barrier remains less than a few percent of the trap depth. The barrier is most sensitive to displacement in the z direction because of the effect of exposing the underlying ground plane. At 5 µm the barrier is two orders of magnitude greater than for the same separation in the x or y axes.

Figure 3.15 shows the ion heating rate resulting from the pseudopotential barrier as the trap separation is increased. The heating rate is presented here in arbitrary units as it represents the gradient term $\left[\frac{\partial}{\partial x}E_0^2(x)\right]^2$ in Equation 2.18. This is directly proportional to the heating rate, $\dot{\bar{n}}$ due to spectral voltage noise on the RF potential. The shape of the curve is similar to Figure 3.14 for the pseudopotential barrier because a larger barrier exhibits a steeper gradient which causes a larger heating rate. Both the barrier height and ion heating rate rise quickly with increasing trap separation, with the largest sensitivity in the z direction. The gradient of the barrier is more important than the height for reducing motional heating. However, a significant pseudopotential gradient can be tolerated without causing significant heating if effective RF filtering is implemented to limit voltage noise near the ion secular frequency [62].

In summary, simulations have shown that for ion trap separations of the order of 10 µm in each axis the decrease in trap depth is expected to be relatively small, and there is a confining potential within the gap region between chips for significantly larger separations. However, the RF barrier and associated ion heating rate resulting from the pseudopotential gradient both increase rapidly with trap separation, particularly in the z direction. Therefore, experimental control over positioning in this direction should be given careful consideration. To maximise the clock speed of quantum computations using the two-module processor it is desirable to achieve high-rate, low-loss shuttling of ions between modules and maximise the number of shuttling operations that can be performed before Doppler cooling or ion reloading is needed. This requires that the ion heating rate due to the barrier be kept as low as possible, which can be achieved with close alignment of the traps and proper filtering of the RF voltage signal to reduce noise. Therefore, a practical target for the positioning system could reasonably be set at 10 µm in each axis. The positioning system that was developed for the experiment is discussed in the next section.

3.5 Positioning system for alignment of ion trap modules

The positioning system was designed to meet the primary experimental requirement of aligning the two ion trap modules to within 10 µm in each of the three spatial axes. As

described in the previous section, this separation distance was chosen to limit the pseudopotential barrier that arises in the gap region so as to achieve a low rate of ion loss during shuttling between modules. To achieve this, it was not sufficient to set up the modules with an initial 10 µm separation during installation into the vacuum chamber of the experiment because thermal expansion and contraction of the mounting structure (shown in Figure 3.1) was expected to occur during UHV baking treatment and cryogenic cooling that would cause the module positions to drift. The results of thermo-mechanical modelling of the mounting structure performed by Raphaël Lebrun-Gallagher show that during bakeout, when the temperature increases from room temperature at 20 °C (293 K) to 140 °C (413 K), the relative displacements of the modules are expected to be 90 µm vertically in z (with the heat sink side higher than the piezo side) and 100 µm away from

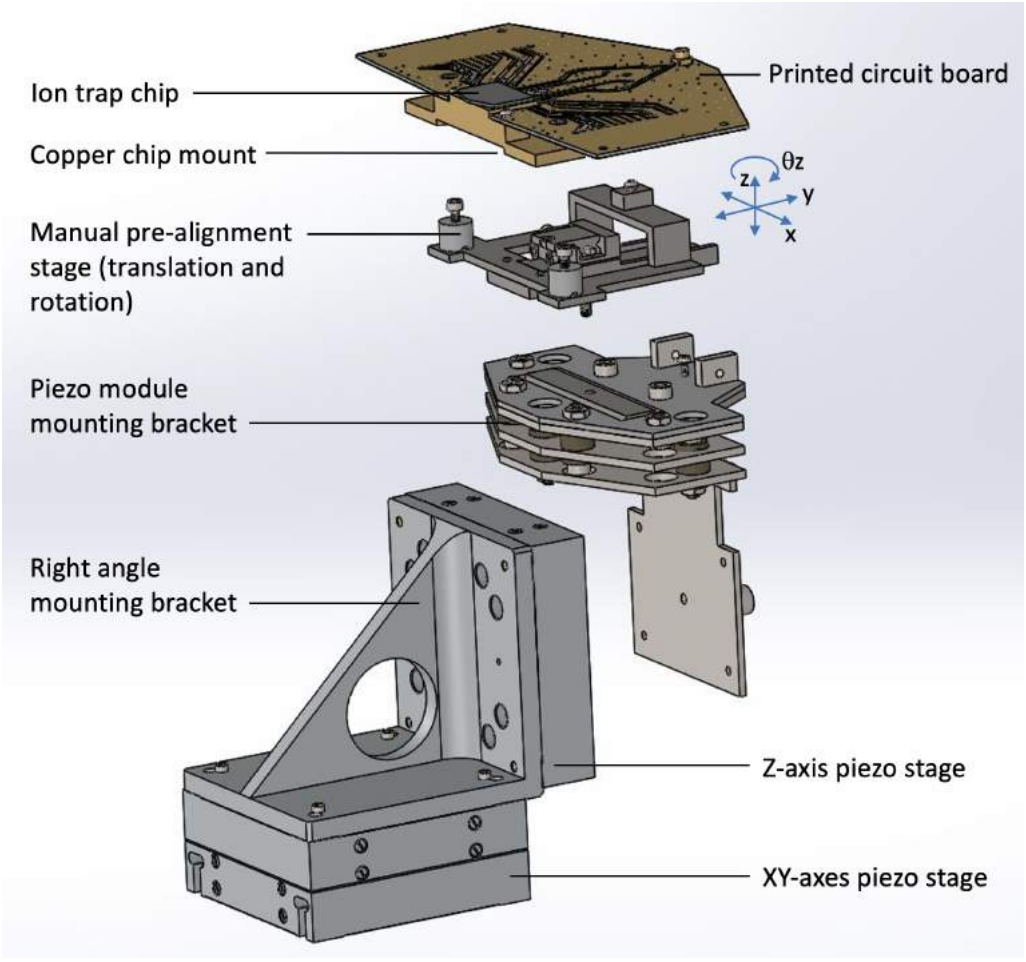

Figure 3.16: Positioning system for the movable ion trap module. The trap and printed circuit board at the top of the structure are mounted to a manually-adjusted pre-alignment stage, which enables small corrections to be made to the installed position with respect to the adjacent heat-sink module. The pre-alignment stage is attached to a high thermal-resistance mounting bracket that is connected to the piezo translation stages at the base of the positioning system.

each other in x, with no movement laterally in y. In the idealised model the modules return to their original positions after cooling to room temperature, with no overall shift in position. During operation of the cryogenic cooling system the temperature decreases from 293 K to 50 K, resulting in an expected relative displacement of 240 µm vertically in z (with the heat sink side below the piezo side) and 60 µm towards each other in x, with no change laterally in y. In practice, the resultant change in the relative position of the modules is also dependent upon any additional movement between component parts of the mounting structure itself, which contributes to positional drift but is difficult to account for directly. The magnitude of this drift was expected to be comparable to the results above. Therefore, to ensure a safe initial distance between modules and avoid risk of collision, there was a requirement to be able to make small adjustments to the module positions after mounting them within the vacuum chamber. Additionally, the requirement to achieve final alignment within 10 µm places a stringent constraint on the maximum permissible rotation angle between traps, since too large a deviation from parallel would cause the edges of the adjacent traps to come into contact. This contact may prevent successful alignment at the RF electrodes on each chip, and may also cause electrical shorts to occur between the traps. The angle for the second generation chips (which are larger than the first generation so the constraint is more stringent) can be calculated using the edge length, $L = 17.4\,\text{mm}$, and the 10 µm separation between chips at the RF electrodes, to give a contact angle of $\sin^{-1}\left(\frac{10\,\mu\text{m}}{L/2}\right) \approx 0.07°$. Therefore a method was required to control the rotation angle between the modules. During installation onto the mounting structure the angle of misalignment was measured by imaging the gap region using a high resolution digital camera.

The above requirements were met using three separate parts of the positioning system: a manual pre-alignment stage for setting the initial positions and angles of the modules during installation into the vacuum chamber; a remotely-operated, piezo-based translation stage for in-vacuum positioning; and finally an imaging system to measure the position and displacement of the modules. For simplicity, one module was mounted to the positioning system and moved relative to the adjacent module which remained stationary. Figure 3.16 shows the components of the positioning system, with the ion trap chip at the top of the structure. The asymmetric printed circuit board surrounds the chip on one side and provides routing and filtering of trapping voltages applied to the chip as described later in this chapter. The copper chip mount provides a thermal interface with the cooling system via a flexible OFHC copper braid (not shown) to allow free movement of the module.

3.5.1 Manual pre-alignment stage

A manually adjustable pre-alignment stage was built by Raphaël Lebrun-Gallagher to provide three-axis translation and rotation about the z axis for the ion trap, copper chip mount and printed circuit board of the movable module. This permits the initial positions of the modules to be setup when installing the modules within the vacuum chamber. The pre-alignment stage was constructed from 316L stainless steel for its relatively low thermal conductivity and coefficient of thermal expansion compared to other candidate materials such as aluminium, while being easier to machine than titanium. The stage incorporates a sliding clamp apparatus with socket screws that can be adjusted set the position of the module. From this initial position, thermo-mechanical modelling indicated no overall rotation was expected as a result of thermal expansion or contraction of the mounting structure. This result eliminated the requirement for a powered rotation stage to be incorporated into the remotely-operated part of the positioning system, which is described next.

3.5.2 Piezo-based XYZ translation stage

A UHV-compatible, high-displacement, three-axis piezo translation stage was constructed by combining a two-axis[5] (XY) stage and a single-axis[6] (Z) stage to align the ion trap modules within the vacuum chamber of the demonstrator experiment. The stages are shown in Figure 3.17.

The stages were selected following research into remotely operated actuator candidates for the positioning system. The requirements were specified as follows:

1. **UHV compatibility.** No materials prohibited by the Vacuum Compatible Materials List[7] of the Laser Interferometer Gravitational-Wave Observatory (LIGO). Capable of operating at 10^{-11} mbar.

2. **Linear displacement.** Since thermal drift was expected to be in the order of a few hundred µm, the travel range was specified to be at least 500 µm in each axis.

3. **Step size.** Given the target for module alignment was 10 µm in each axis, the minimum step size was required to be in the order a few µm or less.

4. **Maximum applied force.** The force applied in the direction of motion. The Z stage was required to operate against the direction of the weight force of the components

[5] Physik Instrumente P625.2 Hera stage
[6] Physik Instrumente P625.1 Hera stage
[7] Current version of the list as at Sep 2021: LIGO Document E960050-v13

attached to it: the chip and copper mount, printed circuit board, pre-alignment stage and module mounting bracket. These were estimated to have a combined weight of 0.30 kg.

5. **Loading capacity.** For the XY stage at the base of the positioning system, the additional weight of the attached Z stage (0.24 kg) must be combined with the weight of the attached module components (0.30 kg). Thus the total loading force is 0.54 kg or 5.4 N applied at 90° to the direction of travel.

6. **Temperature range.** The maximum temperature must be compatible with UHV bakeout temperatures (at least 120 °C). The minimum operating temperature determines the level of thermal insulation that must be added between the stages and the cooling system heat sink.

The piezo stages are constructed from an aluminium frame in which a piezo crystal stack applies force to a movable platform mounted onto zero-play flexure guides. A leveraged design increases the travel range of the device, since the piezo crystal itself typically changes length by only 0.1% due to the piezoelectric effect when a voltage is applied. This flexure technology requires no lubricants or greases and is an inherently vacuum-compatible mode of operation. The stages used in the experiment were certified by the manufacturer for operation at UHV pressures of 10^{-11} mbar. This was achieved by removing all non-compatible materials from the product at the assembly stage during production. Prohibited materials included the polymer outer coating applied to the standard product, the cable

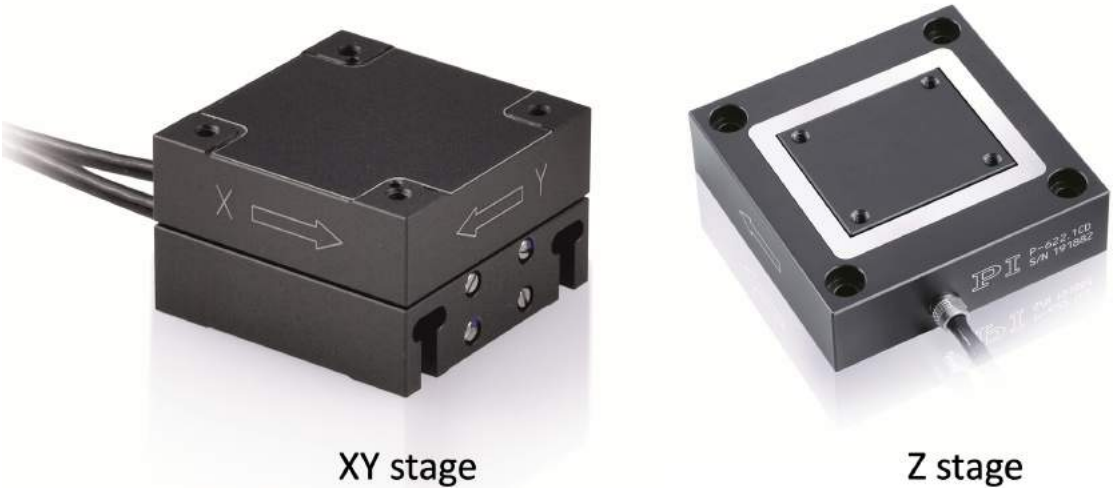

Figure 3.17: Piezo flexure stages manufactured by Physik Instrumente were used to construct the three-axis ion trap module positioning system. The stages deliver displacements of up to 600 µm in each axis with a minimum step size less than 1 µm and were certified compatible with ultra-high vacuum pressures of 10^{-11} mbar.

insulation which was replaced with Kapton polyimide insulation, and a proprietary epoxy used to attach an internal capacitive sensor for positional feedback, which also had to be removed. As a result, operating the stages was only possible in an open-loop mode without the built-in displacement calibration provided by the sensor. Displacement of the stages was measured using the imaging system described in the next section.

The piezo stages provide a maximum linear displacement of 600 µm in each axis, with an achievable precision of <1 µm. Motion of the stages was controlled by an E-727 multi-channel controller module supplied by the manufacturer that applies voltages in the range -20 V to +120 V. This voltage can be finely adjusted in mV increments to permit motion control in step sizes smaller than 1 µm in principle. One of the benefits of this type of piezo stage is the high push/pull force of 10 N in the direction of travel and 10 N load capacity perpendicular to the direction of travel. The operating temperature range is −20℃ to 150 °C, and this meant that multiple layers of thermal insulation were required to be built into the in-vacuum support structure to insulate the stages from the cooling system heat sink. As a further measure to ensure that the piezo stages remained at a temperature above −20℃, the base of the XY stage was thermally anchored via a OFHC copper braid to the vacuum chamber wall which has sufficient thermal mass to remain at room temperature.

To confirm the performance of the stages prior to their installation into the demonstrator experiment, each stage was characterised individually in terms of displacement versus

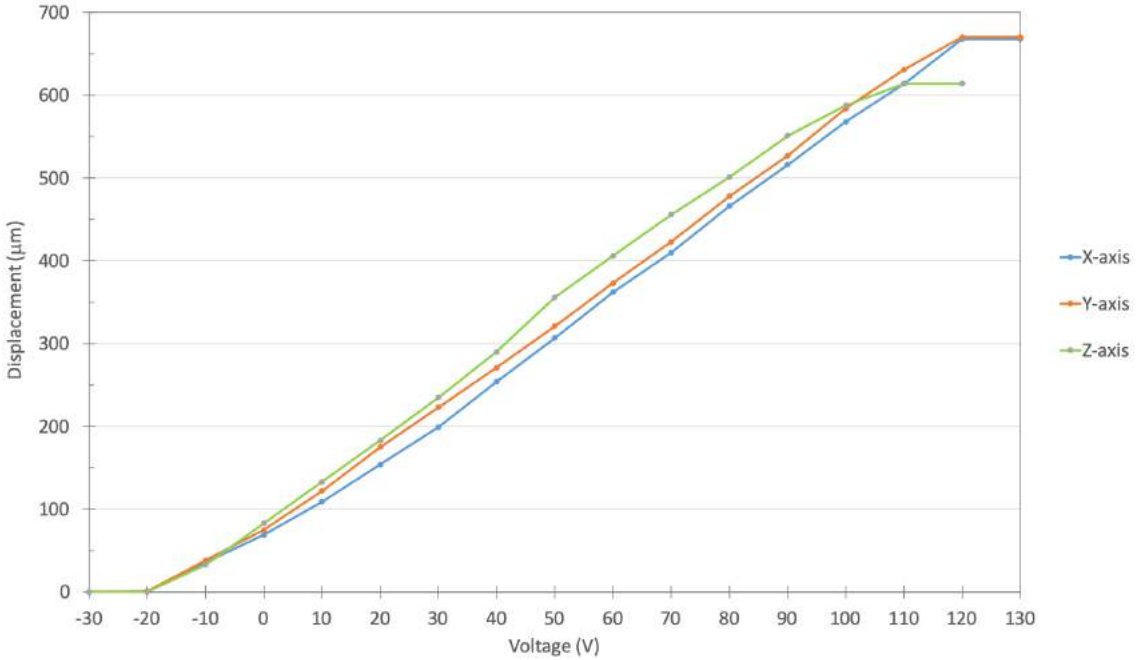

Figure 3.18: Characterisation of the individual piezo stages prior to installation into the demonstrator experiment.

applied voltage. This was achieved by placing the stages under an optical microscope and observing their motion directly, while measuring displacement using calibrated software. The results of the tests are given in Figure 3.18 which demonstrate a linear response with in excess of 600 µm displacement in each axis.

In Chapter 5 the fully assembled XYZ piezo stage is characterised after installation into the experiment and operated both at room temperature and whilst the cooling system was running. During those tests the travel range was reduced to approximately 300 µm.

3.5.3 Imaging system for measurement of module alignment and qubit state detection

The optical imaging system was designed by Raphaël Lebrun-Gallagher to meet the following requirements of the experiment:

1. **High resolution distance measurement for three-axis module alignment.** This required a high magnification for imaging of the trap surfaces in the gap region between modules.

2. **Ion imaging during module-to-module ion transport.** This required low-magnification for a wide field of view across the surfaces of both traps when the modules were aligned.

3. **Single ion and ion chain imaging.** Sufficient resolving power was required to distinguish individual ions and ions within a chain with separation of ∼3 µm.

4. **Fast and accurate qubit state measurement.** This required a high numerical aperture for high photon collection efficiency to achieve high-fidelity readout of the quantum state.

The imaging system consists of a custom-made, high numerical aperture (NA) objective lens[8] to maximise photon collection efficiency, coupled to a reconfigurable secondary lens system to control the overall magnification of the imaging setup. Detection is performed by either a high-resolution, scientific complementary metal-oxide-semiconductor (sCMOS) camera[9] for imaging or a photomultiplier tube (PMT) for photon counting. The objective lens has NA = 0.6 and a fixed magnification, M_o, of -7.5. This results in a theoretical resolution in the xy plane (the plane of the chip surface) of 0.4 µm. However, this resolution is reduced to 1.5 µm owing to undesirable refractive effects introduced by the imaging

[8] Photongear 15920-S large atom imager system
[9] EHD SCM2020-UV sCMOS camera

window of the vacuum chamber. The effective resolution of the imaging system in the z direction is determined by the depth of field (DOF) of the objective lens, which is itself dependent upon NA, and this can be controlled by an iris placed into the optical path and used as an aperture stop. With the iris fully open the objective lens has an unrestricted NA = 0.6 and DOF of 0.8 µm which is impractically small and carries the risk of missing a trapped ion if not perfectly in focus. By reducing the aperture of the iris, the NA can be decreased to 0.2 and the DOF increased to 9.0 µm. After the objective lens, a broadband UV dielectric mirror is fitted into the optical path at an angle of 45°to redirect the optical axis from vertical to horizontal. A 369 nm band-pass filter is placed at the intermediate image plane to optimise the output for fluorescence photons from the $^2S_{1/2}$ to $^2P_{1/2}$ cycling transition used in the state detection procedure.

Before the light reaches the camera or PMT, a reconfigurable secondary lens system is inserted as a means to control the overall magnification of the setup. The secondary lens system consists of an air-spaced doublet assembly within an imaging tube that may be inserted into the optical path to provide magnification, M_s = -3.3. By reversing the direction of the lens system the magnification becomes $1/M_s$. The total magnification of the imaging system is then either $M_o.M_s = 24.7$ or $M_o/M_s = 2.2$. This satisfies the competing demands of requirements (1) and (2) above, for which either a high magnification or lower magnification is needed depending on the experimental application. After the secondary lens system a removable UV mirror is used to control light propagation either to the sCMOS camera or the photon counting head of a PMT.

The objective lens is mounted onto a 3-axis translation stage that provides better than 1 µm precision for focusing onto the ion during trapping or the surfaces of the chips when performing module alignment. The remainder of the imaging system assembly is adjustable using two 2-axis translation stages to align the image focus onto the camera. Finally, a 2-axis translation stage is used to position the PMT.

3.6 In-vacuum printed circuit boards with low-pass filtering

The ion trap modules require in-vacuum routing and filtering of the trapping potentials applied to the microchips. Additionally, high current connections are required for the on-chip magnetic gradient coils. To achieve this, each of the modules incorporates an asymmetric printed circuit board (PCB) that is positioned around the chip on only three sides to leave the alignable edge unobstructed. This constrains the PCB dimensions, which must also be compact enough for mounting within the confined space of the vacuum chamber.

The maximum size of the PCBs is limited specifically by the 160 mm bore diameter of the CF160 flange of the chamber (described in Chapter 5) which the boards need to pass through during installation into the experiment. All of the components used to construct the PCB must be compatible with UHV pressures (in the 10^{-11} mbar regime), and this was ensured by checking materials against the LIGO vacuum compatible materials list, or by certifying inside a test chamber before use. The boards are constructed from Rogers RO4350B material which is UHV compatible, has a low thermal conductivity, and good electrical stability. The board is covered with 1 oz/sq.ft (35 µm thickness) electrodeposited copper foil, and the surface is finished with an electroless nickel immersion gold (ENIG) layer on both sides, which consists of 3-6 µm of nickel covered by a thin layer of gold (0.04 to 0.1 µm). The ENIG finish provides good adhesion for gold or aluminium wire bonds that are used to form electrical connections from the PCB surface to the chip electrodes. The positions of wire bonds must not obstruct laser access to the trapping region near the edge of the chips and so electrical connections are made far from the edge. The boards are produced 'bare', without the solder mask or silkscreen layers, to ensure UHV compatibility. The PCBs incorporate first-order low-pass RC filters using 1 kΩ thick film resistors[10] and 620 pF multi-layer ceramic capacitors[11]. This combination of R and C values gives a filter cut-off frequency of $f_c = 1/2\pi RC$ of 257 kHz to suppress voltage noise on the DC signals. Similar 1 kΩ thick film resistors have been tested at 70 K and shown to have smaller than 1% change in resistance compared to the value at room temperature [116]. The capacitors are made of class C0G ceramic material that has been proven UHV safe in previous experiments conducted within the Ion Quantum Technology research group at Sussex. Although the capacitors are rated to a minimum temperature of -55 °C, the ceramic material has been shown to exhibit less than 1% drift in capacitance when cooled from room temperature down to 4 K [117]. Lead free solder was used for all board components, and the completed PCBs were cleaned to remove solder flux using Chemtronics Flux-Off detergent and isopropyl alcohol following the procedure outlined by the LIGO experiment[12].

In the following sections I describe the individual designs of the specific boards made for the ion trap chips of each generation.

[10] KOA Speer RK73G2ATTD1001C
[11] Kemet C0603C621J5GACTU
[12] LIGO Document T060280-x0, Printed Circuit Boards for Ultra-high Vacuum

3.6.1 PCBs for first generation ion trap chips

The PCB for use with the SET150-02 linear ion trap chip is shown in Figure 3.19. The PCB schematic is shown in Figure 3.20. The PCB provides 2 RF traces and 30 filtered DC traces on a two-layer board of core thickness 0.76 mm and overall board thickness 0.85 mm including surface copper and ENIG layers.

The RF trapping voltages are carried to the PCB from an SMA feedthrough flange on the vacuum chamber by UHV compatible 50 Ω SMA-to-SMP cables of length 150 mm. On the top surface of the PCB, the RF cable connects to a 50 Ω gold-plated, beryllium copper SMP socket[13] with PTFE insulator that is soldered to impedance-matched RF traces to minimise reflected power. The DC voltages are delivered to PCB from a 50-pin DSUB feedthrough on the vacuum chamber by a Kapton insulated copper ribbon cables[14] with 50 individual wires. The ribbon cable is fitted with a PEEK DSUB female connector with 50-pins for attachment to the vacuum feedthrough at one end and splits into two 15-pin PEEK DSUB connectors at the PCB end. The cable connects on the bottom surface of the PCB to two 15-pin male DSUB connectors[15] soldered to the DC traces which pass through vias from the top to the bottom of the board.

The PCB designed for the EX150 X-junction ion trap chip is shown in Figure 3.21. The PCB schematic is shown in Figure 3.22. The PCB provides 2 RF traces and 66 filtered DC traces on a two-layer board with 0.85 mm total thickness. This PCB uses the UHV-compatible resistors and capacitors described at the beginning of this section, and the same SMP connector is used to connect RF signals to the bpard as previously described. However, owing to the increased number of DC connections for the X-junction chip, the PCB uses two custom-made 36-pin connectors to interface with the ribbon cables carrying DC potentials from the vacuum feedthroughs to the PCB. This was necessary because it was not possible to find a suitable UHV-compatible connector off the shelf to fit the compact dimensions of the board.

Custom DC connector

The custom connector is shown in Figure 3.23 in the closed position. It has three pieces: a PCB-side housing, a cable-side housing, and a cable-side back plate, all made of PEEK. The parts were manufactured by the university engineering workshop at Sussex. Holes in

[13] Rosenberger 19S101-40ML5.
[14] Allectra 50 wire ribbon cable, part number 380-D50FXPR
[15] 15 pin female vertical SMT socket, product code 194278

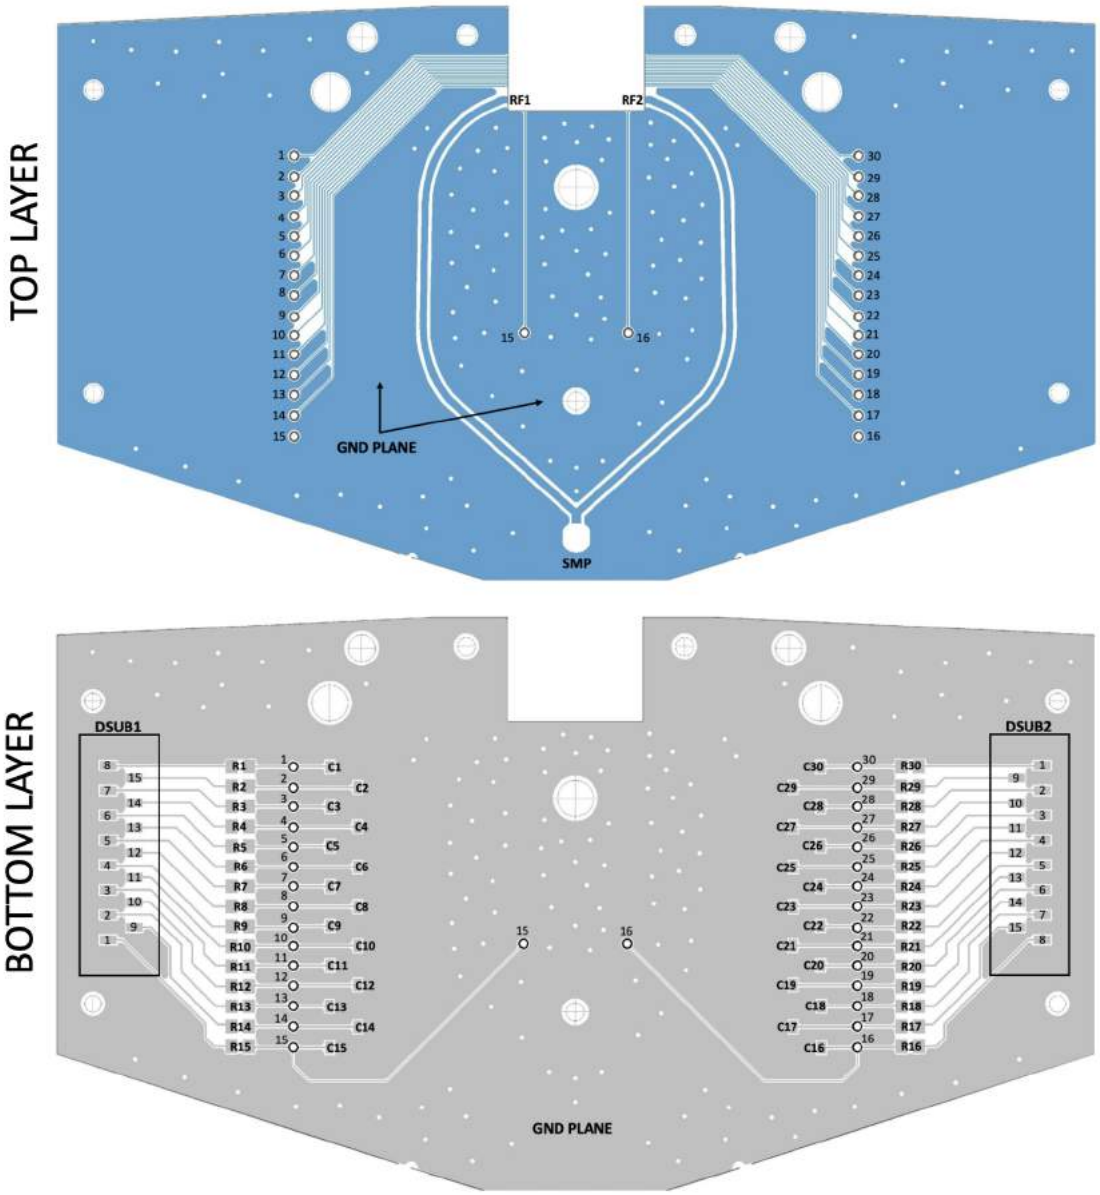

Figure 3.19: Ultra-high vacuum compatible printed circuit board for routing and filtering of voltages sent to SET150-02 linear surface trap microchip. The board dielectric material is Rogers RO4350B with copper conductor and an ENIG surface finish for good wire bond adhesion. The top layer (*top image*) carries two RF traces from a single SMP connector near the edge of the board, and 30 DC traces pass from the top layer to the bottom layer of the board through via connections. On the bottom layer (*bottom image*) each DC trace has first-order low pass RC filtering with R=1 kΩ and C=620 pF resulting in a cut-off frequency of 257 kHz. The bottom layer also provides solder pads for two 15-pin subminiature D-type connectors (labelled DSUB1 and DSUB2) to interface with ribbon cables that deliver the DC signals to the board. Signal traces on both layers are surrounded by ground planes that are electrically connected by vias through the board.

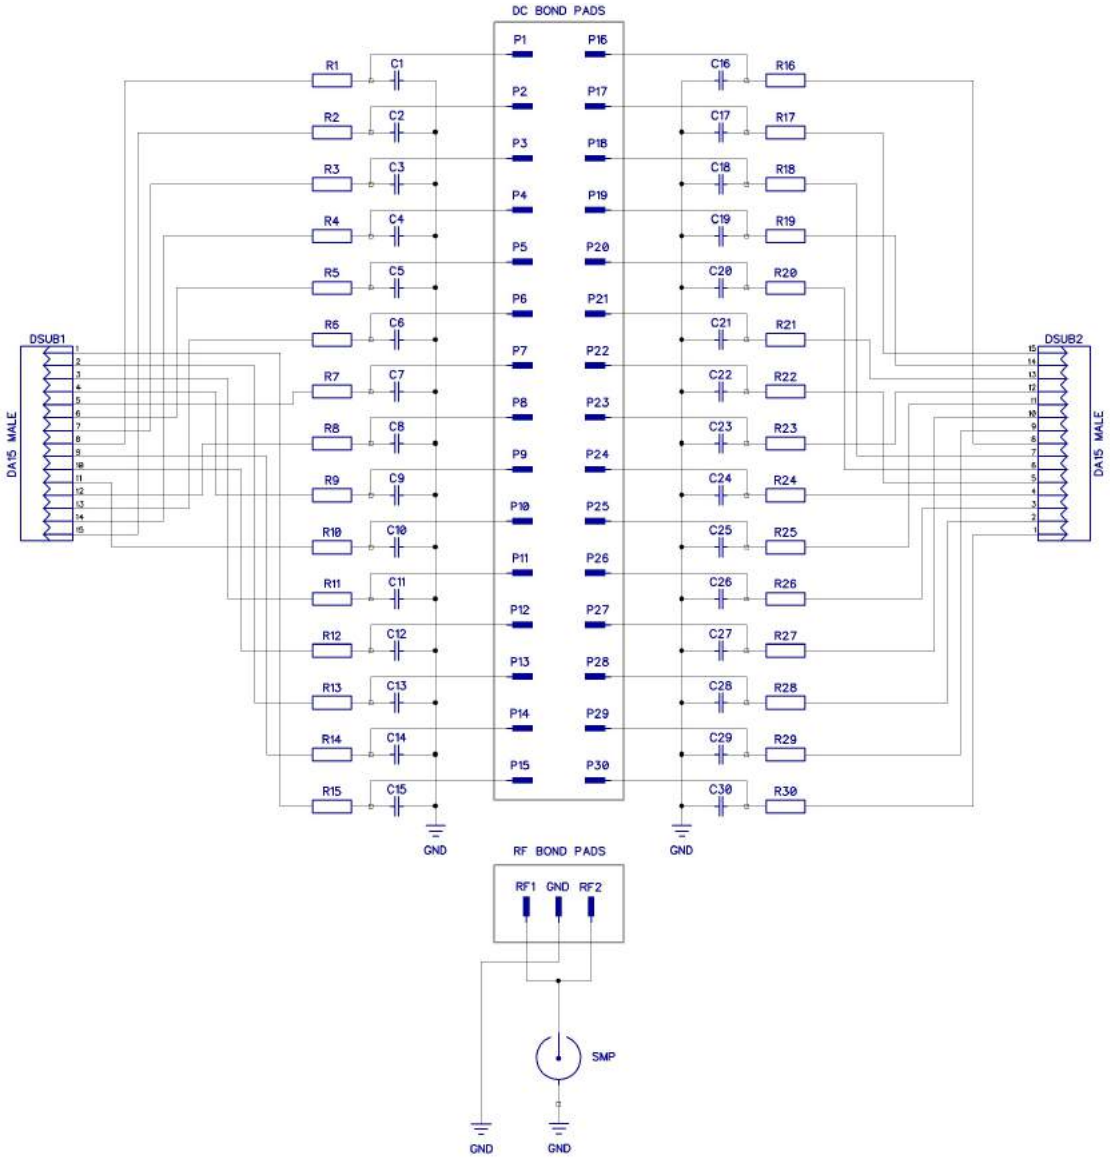

Figure 3.20: Printed circuit board schematic showing signal routing and filtering for voltages applied to the SET150-02 linear surface trap chip. The board provides two RF traces (labelled RF1 and RF2) for the oscillating trapping potential applied to the RF rails at a frequency near 19.3 MHz, and a ground connection (GND) for the central electrode between the RF rails. There are 30 DC traces (with wire bond pads labelled P1 to P30) for quasi-static control potentials which are used as endcaps and for applying shuttling waveforms. First-order RC low pass filters are added to each DC trace, with values of R=1 kΩ and C=620 pF giving a cut-off frequency of 257 kHz. Filtering resistors and capacitors are labelled R1 to R30 and C1 to C30, respectively. Two 15-pin subminiature D-type connectors (labelled DSUB1 and DSUB2) are soldered to the board and attach to ribbon cables carrying the DC potentials.

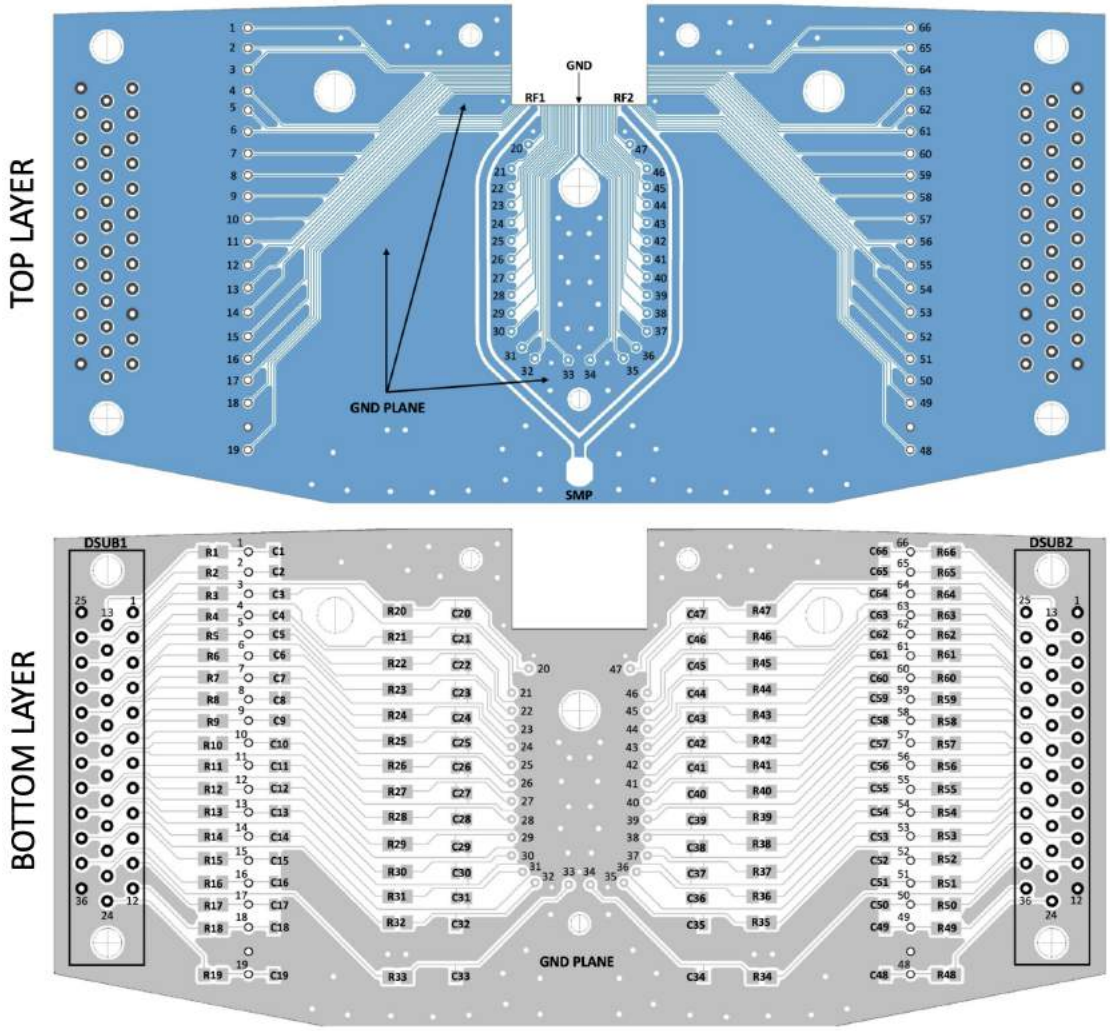

Figure 3.21: Ultra-high vacuum compatible printed circuit board for use with EX150 X-junction surface trap microchip. The board dielectric material is Rogers RO4350B with copper conductor and an ENIG surface finish to support strong wire bond welds. The top layer of the board (*top image*) carries two RF traces from a single SMP connector near the board edge. The 66 DC traces pass through vias from the top to the bottom layer. The bottom layer (*bottom image*) provides first order low pass RC filtering on the DC traces, with R=1 kΩ and C=620 pF resulting in a cut-off frequency of 257 kHz. The bottom layer also provides solder pads for two 36-pin custom D-type connectors (labelled DSUB1 and DSUB2) to interface with ribbon cables that deliver the DC signals to the board. On both layers all signal traces are surrounded by ground planes that are electrically connected by vias through the board.

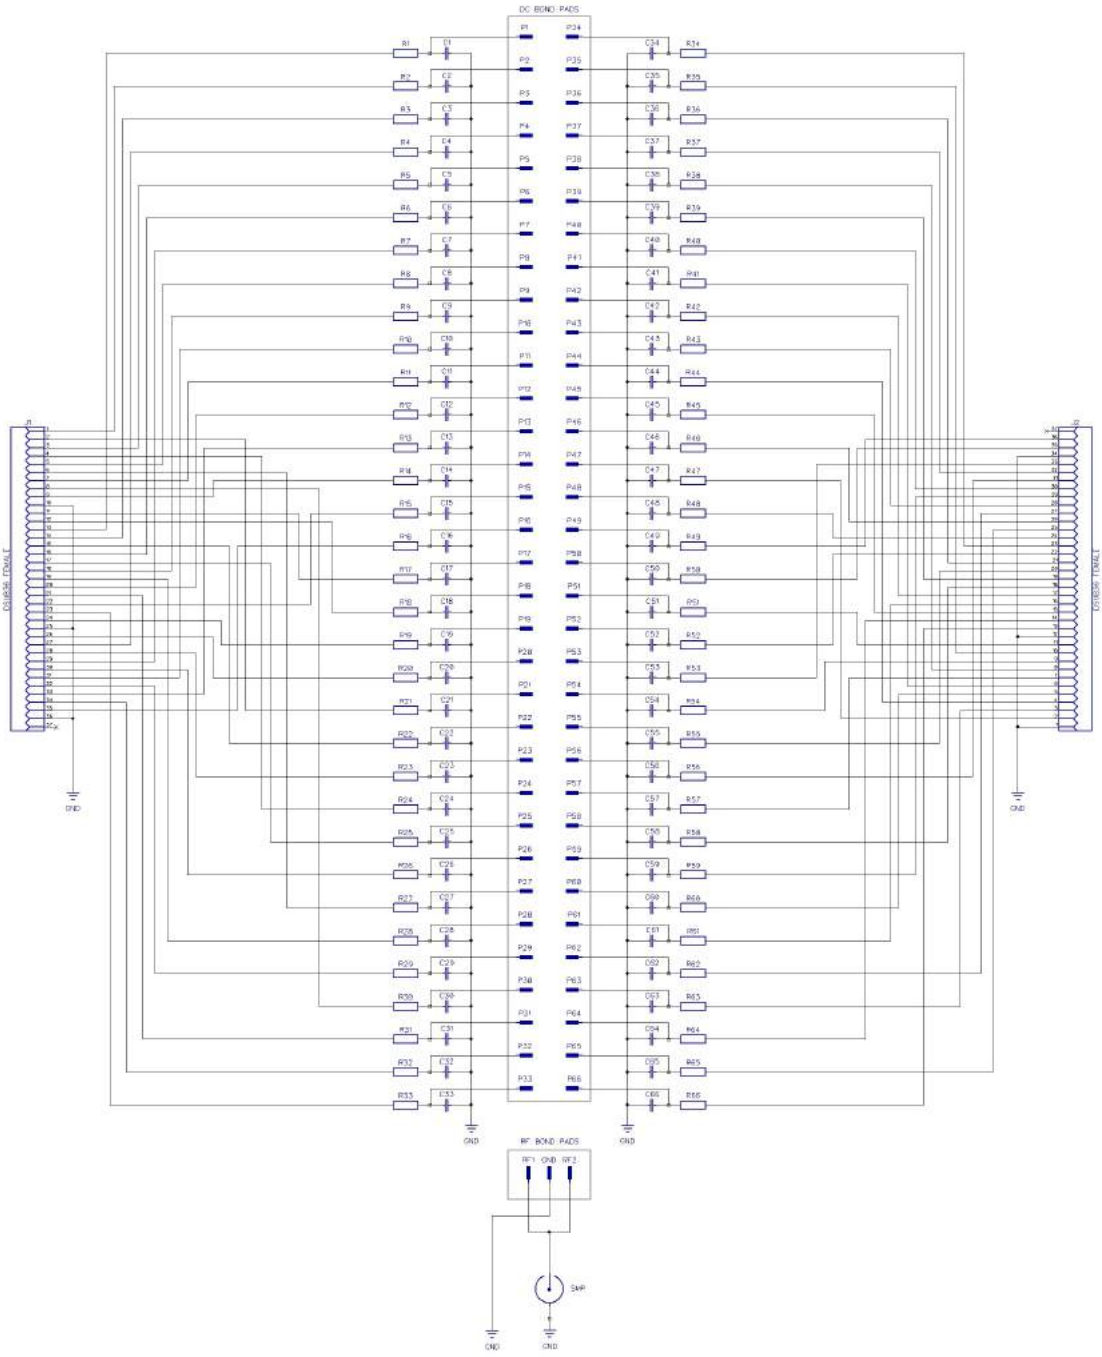

Figure 3.22: Printed circuit board schematic for electrical connections to the EX150 X-junction surface trap chip. The board provides two RF traces (labelled RF1 and RF2) for the oscillating potential applied to the RF rails, and a ground connection (GND) for the central electrode between the rails. There are 66 DC traces (with wire bond pads labelled P1 to P66) for quasi-static control potentials which are used as endcaps and for applying shuttling waveforms. First-order RC low-pass filters are applied to each DC trace, with values of R=1 kΩ and C=620 pF giving a cut-off frequency of 257 kHz. Filtering resistors and capacitors are labelled R1 to R66 and C1 to C66, respectively. Two 36-pin custom D-type connectors (labelled DSUB1 and DSUB2) are soldered to the board, which attach to ribbon cables carrying the DC potentials. Two of the 36 pins are dedicated for the external ground signal connection to the PCB ground planes, and one pin is not connected.

the PEEK housings are designed to accept 36 gold-plated beryllium copper crimp pins[16], which are used for the DSub contacts and push-fitted into place. The peek back plate is required to securely hold the pins of the cable-side housing, after the pins have been crimped onto individual wires of the ribbon cable which connects to the feedthrough of the vacuum chamber. On the PCB there are two sets of 36 through-hole pads into which the pins shown at the top of the connector are soldered. The assembly is held together by two M3x20 mm screws, which pass through the PCB from the top side to secure the connector onto the board.

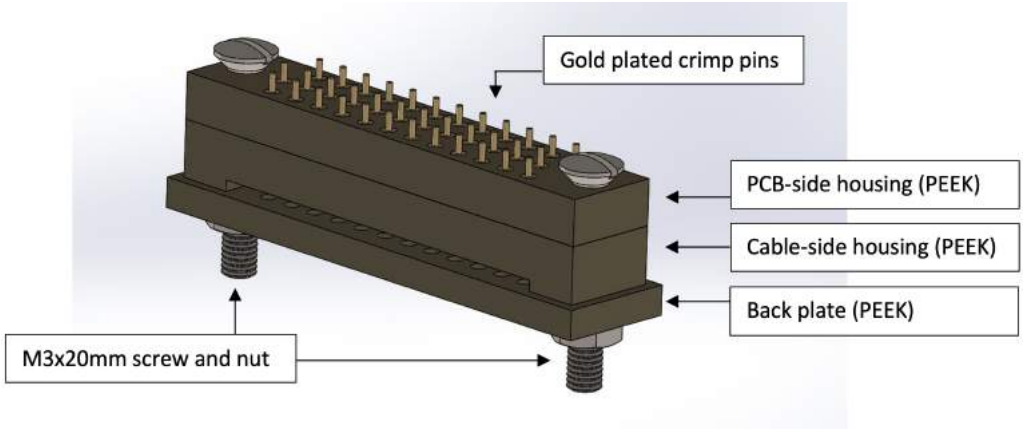

Figure 3.23: A custom-made 36-pin DSUB-style connector for use with compact UHV-compatible circuit boards. Gold plated pins are used as contacts fitted within a PEEK housing. Two of these connectors were attached onto the underside of the in-vacuum PCB designed for the first generation chip with X-junction geometry, and similarly for the PCB used with the second generation linear chip with integrated magnetic gradient coils. Individual wires of a ribbon cable are crimped to the contact pins within the cable side housing. Contact is made between the connector and the DC traces on the bottom layer of the PCB by soldering the contact pins of the PCB-side housing into through-hole pads on the board surface.

Calculation of RF trace impedance

To maximise the efficiency with which RF power is delivered to the ion trap chip it is important to minimise reflections of the RF signal that occur due to impedance mismatches along the RF path. To achieve this, all standard connections within the RF circuit have 50 Ω impedance, including the output from the signal generator, inputs and outputs at the RF amplifier, the SMA feedthroughs on the vacuum chamber and the cables and connectors used to transmit the signal to the in-vacuum PCB. The impedance of the RF trace on the PCB was also approximately matched to 50 Ω. This was achieved by modelling the trace as a coplanar waveguide as shown in Figure 3.24 and calculating the values of trace width, a,

[16] Mac 8 straight through-hole female socket pin, part no. PD-10

and gap size, w, to provide a characteristic impedance, Z_0, that is close to 50 Ω. However, a trade-off was made by imposing a minimum gap, w, of 0.8 mm to avoid the risk of electrical breakdown given that RF voltage amplitudes of ∼ 200 V were to be applied to the board. Although the RO4350B board material has a reported breakdown voltage of 31 kV mm^{-1}, this can be reduced by roughness at the trace edge or the presence of unwanted conducting filaments between the trace and ground plane.

The characteristic impedance of a coplanar waveguide [118] can be calculated using the equation:

$$Z_0 = \frac{60\pi}{\sqrt{\varepsilon_{\text{eff}}}} \frac{1}{\frac{K(k)}{K(k')} + \frac{K(k_l)}{K(k_l')}}, \tag{3.2}$$

where we define the following terms based on the dimensions shown in Figure 3.24:

$$k = \frac{a}{b}, \tag{3.3}$$

$$k' = \sqrt{1 - k^2}, \tag{3.4}$$

$$k_l' = \sqrt{1 - k_l^2}, \tag{3.5}$$

$$k_l = \frac{\tanh \frac{\pi a}{4h}}{\tanh \frac{\pi b}{4h}}, \tag{3.6}$$

and the effective dielectric constant is given by

$$\varepsilon_{\text{eff}} = \frac{1 + \varepsilon_{\text{r}} \frac{K(k')}{K(k)} \frac{K(k_l)}{K(k_l')}}{1 + \frac{K(k')}{K(k)} \frac{K(k_l)}{K(k_l')}}, \tag{3.7}$$

where $K(k)$ represents an elliptical integral of the first kind (defined in Ref [119]) evaluated numerically using the approximation given in Ref [120].

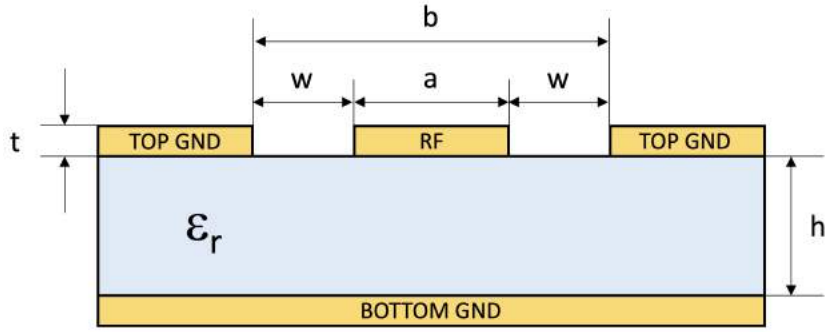

Figure 3.24: The RF signal trace on the PCB was modelled as a coplanar waveguide for calculation of its characteristic impedance. The dimensions a and b were chosen to approximately match the 50 Ω impedance of the on-board SMP connector, while imposing a sufficient gap, w, to avoid the risk of electrical breakdown at voltage amplitudes of 200 V.

Dimension	SET150-02 PCB	EX150 PCB
Width of RF signal trace (a)	0.8 mm	0.8 mm
Length of RF signal trace	57 mm	45 mm
Gap from signal trace to top ground plane (w)	0.4 mm	0.4 mm
Thickness of RF trace (t)	38 µm	38 µm
Height above bottom ground plane (h)	0.76 mm	0.76 mm
Relative dielectric constant (ε_r)	3.76	3.76
Estimate of characteristic impedance, Z_0:	65.7 Ω	65.7 Ω

Table 3.1: Dimensions of RF signal traces on the PCBs made for first generation chips (PZ refers to the piezo side of the experiment, while HS refers to the heat sink side). Terms shown in brackets refer to dimensions labelled in Figure 3.24 that were used to calculate the characteristic impedance, Z_0, of the RF traces which were made close to 50 Ω.

The values in Table 3.1 show the dimensions chosen for the RF traces on the PCBs designed for first generation chips and the resulting characteristic impedance, Z_0 for each trace calculated using the coplanar waveguide model. Since the only difference is the length of the traces, which does not factor into the model, the characteristic impedances have the same value of 65.7 Ω.

3.6.2 PCBs for second generation ion trap chips

The PCBs for the magnetic gradient chips are required to route RF and DC potentials for ion trapping and transport, and provide high current connections to on-chip magnetic gradient coils. The gradient coils were connected only on the heat sink side of the experiment due to the direct, high thermal conductance connection to the cooling system to dissipate the heat load effectively at this location. For the piezo module, the gradient coils were not connected so as to reduce the total power dissipated by the experiment. This decision was taken due to the copper braid that connects the piezo module to the heat sink and the resulting temperature gradient that arises along the length of the braid, which reduces cooling efficiency for this module. Therefore the PCB for the piezo side of the experiment does not provide high current traces. The capability to perform quantum logic operations using the two-module processor is maintained using the magnetic gradient coils on the heat sink module, which provide a gate zone accessible to all ions via transport between modules.

The PCB designed for the EL125 chip on the piezo side of the experiment is shown in Figure 3.25 and the schematic is shown in Figure 3.26. The PCB designed for the EL125 chip with magnetic gradient connections on the heat sink side of the experiment is shown in Figure 3.27, and the schematic is shown in Figure 3.28.

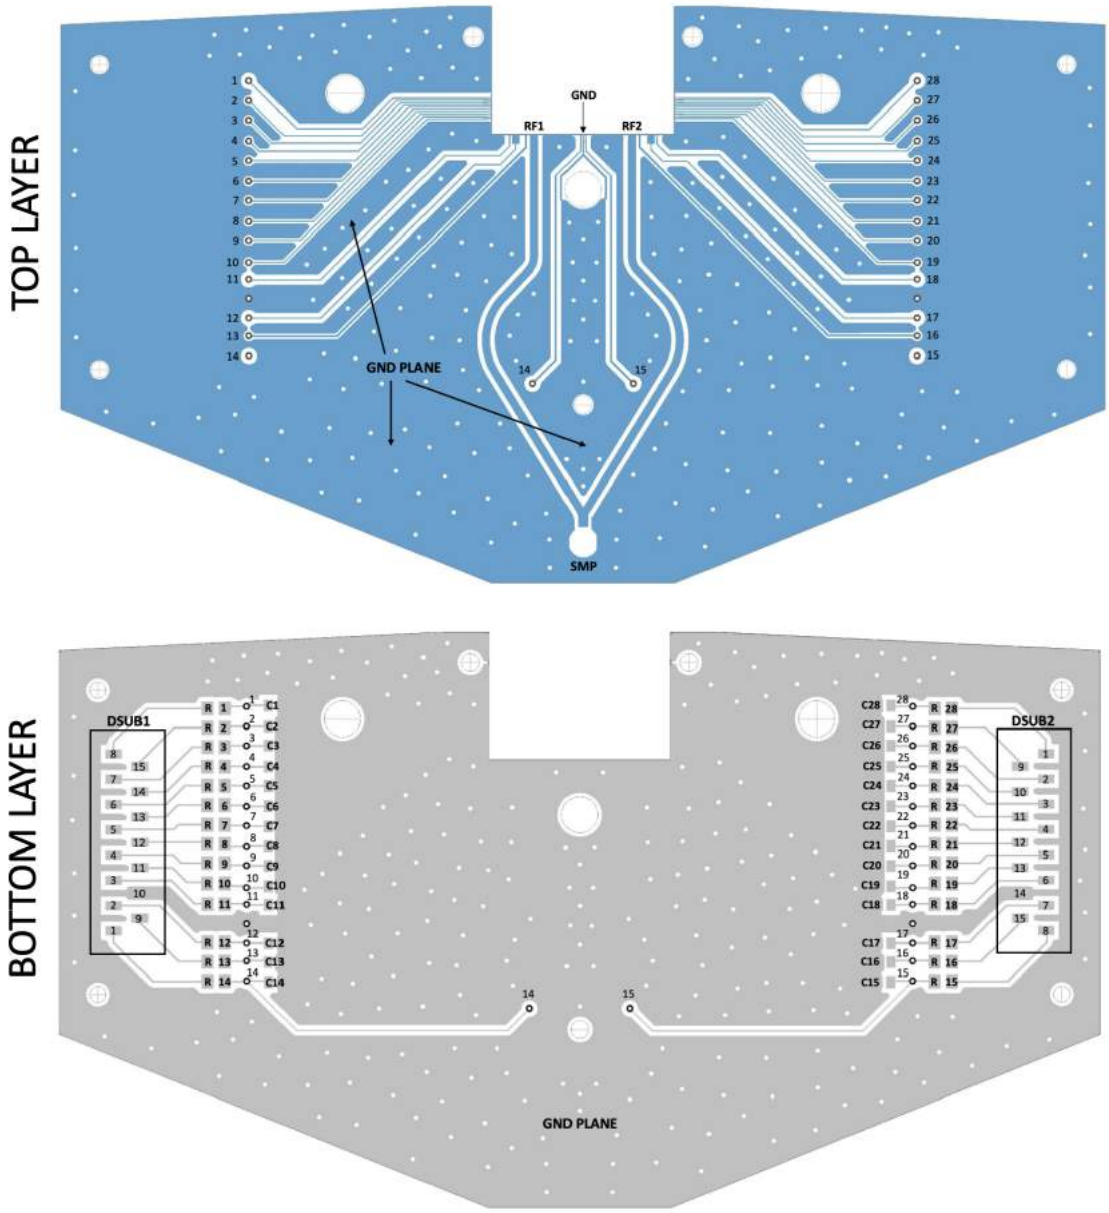

Figure 3.25: Printed circuit board for routing and filtering of voltages sent to EL125 linear surface trap microchip. The top layer of the board (*top image*) carries two RF traces from a single SMP connector near the edge of the board, and 28 DC traces that connect through vias to the bottom layer of the board. The bottom layer (*bottom image*) provides first-order low-pass RC filtering on the DC traces, with R=1 kΩ and C=620 pF resulting in a cut-off frequency of 257 kHz. The bottom layer also provides solder pads for two 15-pin subminiature D-type connectors (labelled DSUB1 and DSUB2) to interface with ribbon cables carrying the DC voltages to the board. Signal traces on both layers are surrounded by ground planes electrically connected by vias through the board. One pin on each D-type connector is dedicated to an external ground connection for the PCB.

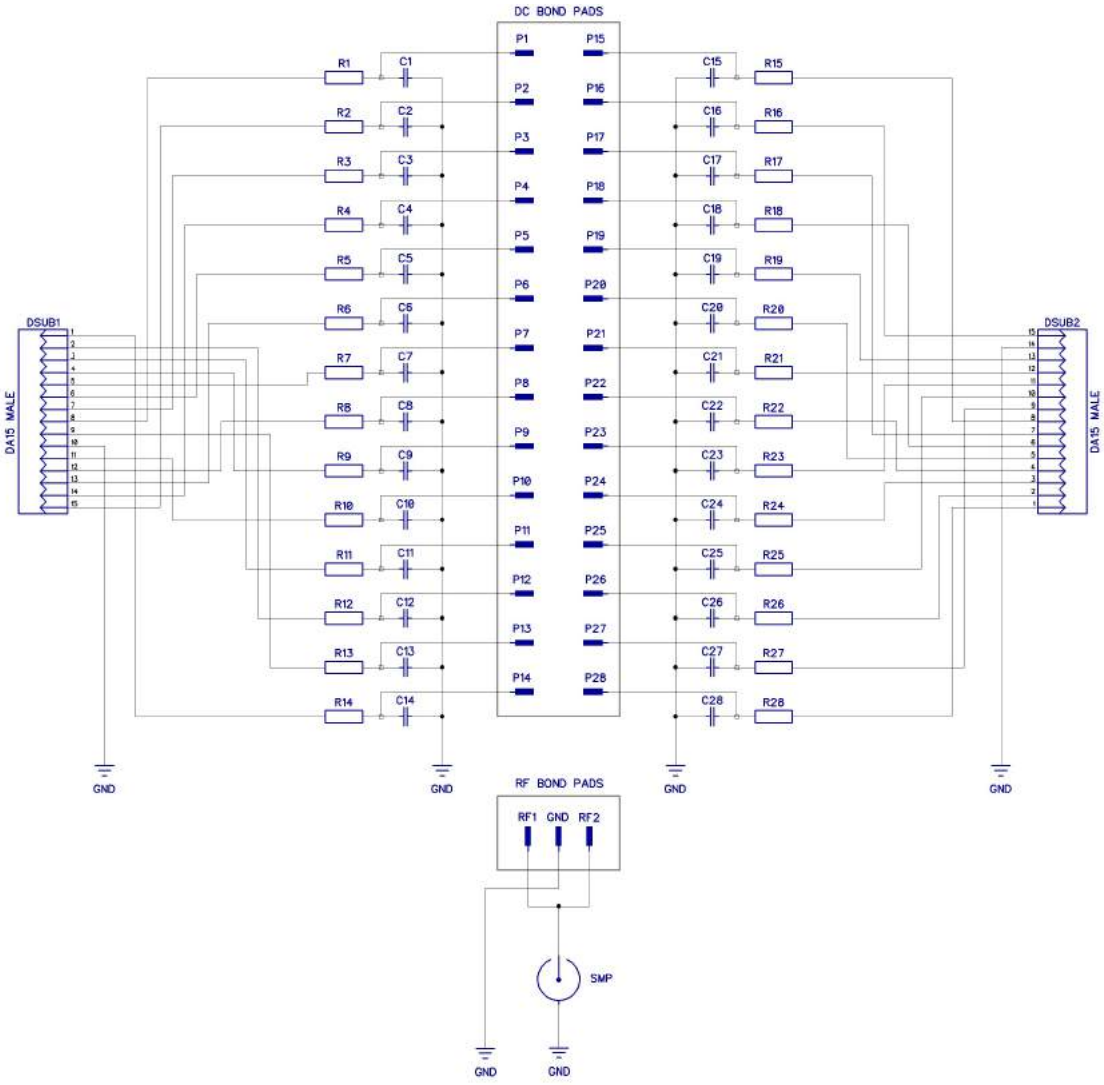

Figure 3.26: Printed circuit board schematic routing and filtering of voltage signals applied to the EL125 linear surface trap chip. The board provides two RF traces (labelled RF1 and RF2) for the oscillating trapping potential applied to the RF rails at a frequency near 19.3 MHz, and a ground connection (GND) for the central electrode between the RF rails. There are 28 DC traces (with wire bond pads labelled P1 to P28) for control potentials which are used as endcaps and for applying shuttling waveforms. First-order RC low-pass filters are added to each DC trace, with values of R=1 kΩ and C=620 pF giving a cut-off frequency of 257 kHz. Filtering resistors and capacitors are labelled R1 to R28 and C1 to C28, respectively. Two 15-pin subminiature D-type connectors (labelled DSUB1 and DSUB2) are soldered to the board and attach to ribbon cables carrying the DC potentials. One pin on each D-type connector is dedicated to an external ground connection.

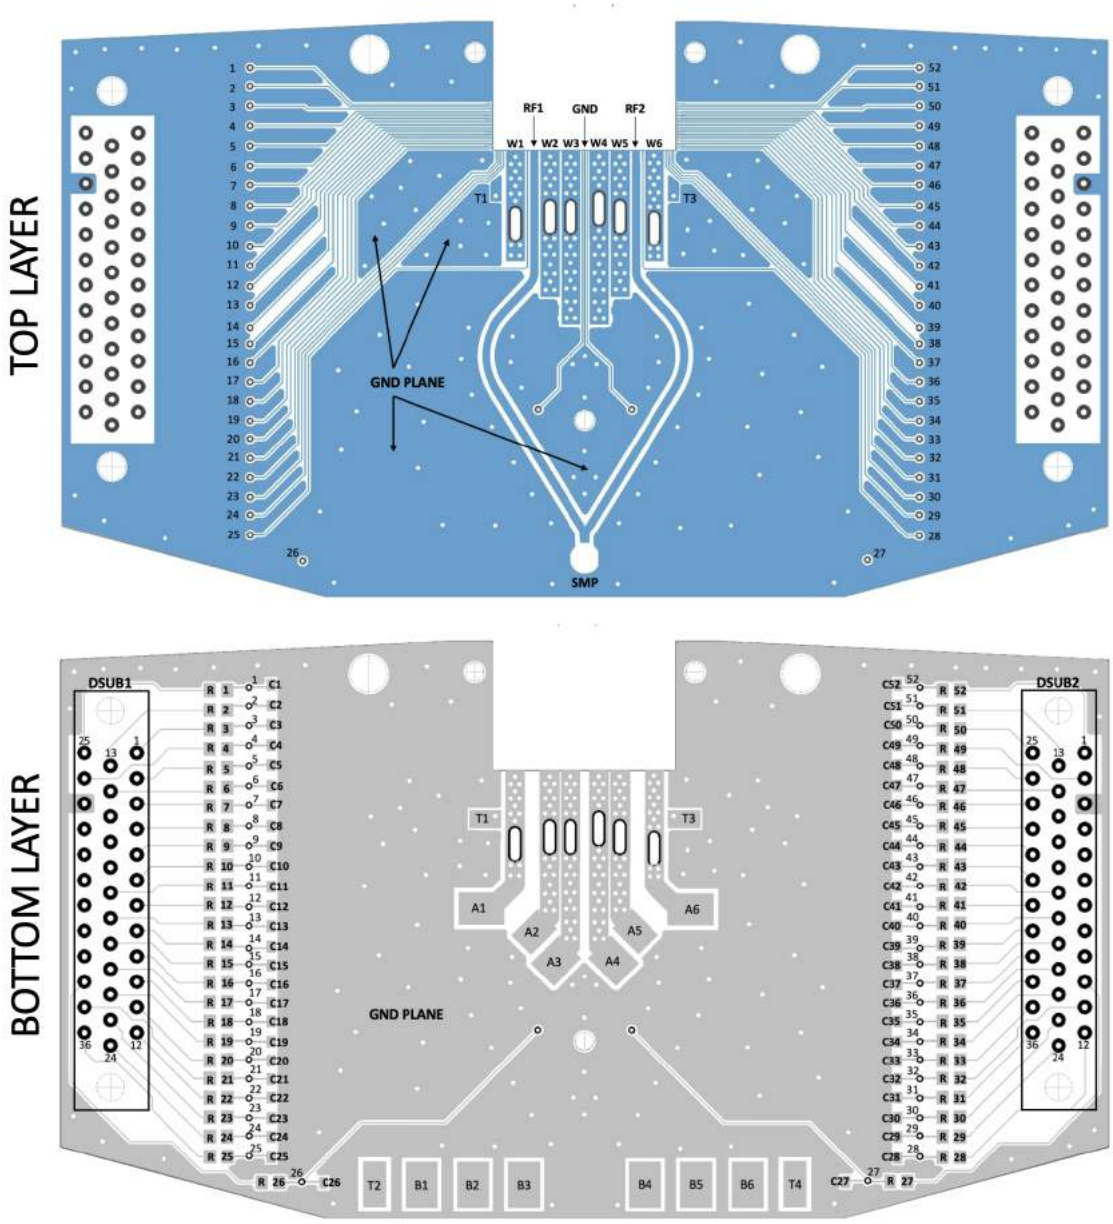

Figure 3.27: Printed circuit board for routing and filtering of voltages sent to EL125 linear surface trap microchip with integrated magnetic gradient. The top layer of the board (*top image*) carries two RF traces from a single SMP connector near the edge of the board, partial DC traces that connect through vias to the bottom layer of the board, and high-current traces that also pass between the top and bottom layers. The bottom layer (*bottom image*) provides first order low pass RC filtering on the DC traces, with R=1 kΩ and C=620 pF resulting in a cut-off frequency of 257 kHz. The bottom layer also provides solder pads for two 36-pin custom D-type connectors (labelled DSUB1 and DSUB2) to interface with ribbon cables that deliver the DC signals to the board. Signal traces on both layers are surrounded by ground planes that are electrically connected by vias through the board. The pads labelled A1 to A6 and B1 to B6 are connected pairwise by 1 mm copper wires soldered to the board to carry high currents for the on-chip magnetic gradient coils that are electrically connected using beryllium copper clamps screwed onto W1 to W6.

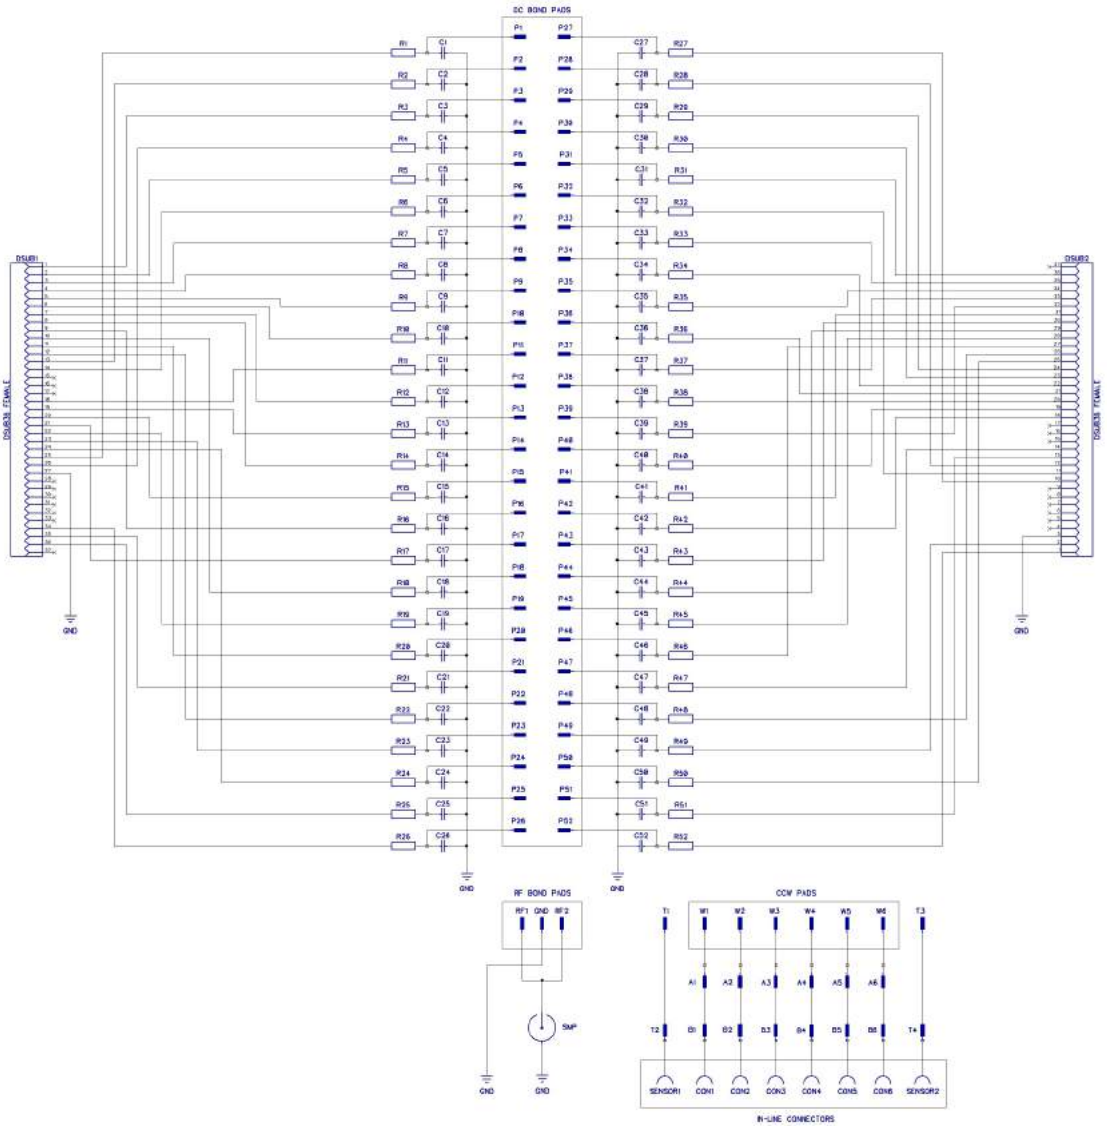

Figure 3.28: Printed circuit board schematic showing signal routing and filtering for voltages applied to the EL125 linear surface trap chip. The board provides two RF traces (labelled RF1 and RF2) for the oscillating trapping potential applied to the RF rails at a frequency near 19.3 MHz, and a ground connection (GND) for the central electrode between the RF rails. There are 52 DC traces (with wire bond pads labelled P1 to P52) for quasi-static control potentials which are used as endcaps and for applying shuttling waveforms. First-order RC low-pass filters are added to each DC trace, with values of R=1 kΩ and C=620 pF giving a cut-off frequency of 257 kHz. Filtering resistors and capacitors are labelled R1 to R52 and C1 to C52, respectively. Two custom 36-pin subminiature D-type connectors (labelled DSUB1 and DSUB2) are soldered to the board and attach to ribbon cables carrying the DC potentials.

PCB for the piezo side of the experiment

The PCB provides two RF traces and 28 filtered DC traces on a two-layer board of core thickness 0.76 mm and overall board thickness 0.85 mm including surface copper pour and ENIG layers. All of the surface mount components used to construct this PCB are the same as described in the previous sections. The RF traces connect to a 50 Ω SMP connector on the top surface of the board. First-order low-pass RC filters with a cut-off frequency of 257 kHz are applied to the DC traces. The bottom layer provides solder pads for two 15-pin subminiature D-type connectors that interface with ribbon cables carrying DC control potentials to the board. Ground planes on the top and bottom layers are electrically connected by vias through the board.

PCB for the heat sink side of the experiment

This PCB provides two RF traces and 52 filtered DC traces on a four-layer board constructed from Rogers RO4350B material and RO4450F prepreg material. The board has two internal ground planes to provide shielding between the RF traces on the top layer and the high-current traces on the bottom layer. The overall PCB thickness is 0.94 mm. The same surface mount components are used for this PCB as for the others described previously. The DC potentials are supplied to the board using the custom manufactured 36-pin connector described in section 3.6.1.

Current-carrying traces on the PCB were designed to maximise conductor area within the space available in order to reduce resistive power dissipation. The high current electrical connections between the PCB and on-chip gradient coils were made using gold-plated beryllium copper clamps. These attach to the PCB on the pads labelled W1 to W6 in Figure 3.27, and are fastened in place using M1 screws that pass through the elongated through-holes. The shape of the through-holes allows the tip of the clamp to be positioned onto the CCW pads on the chip before the screws of the clamps are tightened to make good electrical contact. On the backside of the PCB, connections to the external high-current source are provided by in-line connectors soldered to the pads labelled B1 to B6. Kapton insulated copper wires with 1 mm diameter were soldered between pads B1 to B6 and pads A1 to A6 to carry current to the clamps on the top side of the board (since pads W1 to W6 are connected to pads A1 to A6 by vias). The pads labelled T1 to T4 were added to the PCB for temperature sensors.

Calculation of RF trace impedance

The RF traces on the PCBs were impedance matched as closely as possible to the 50 Ω SMP connector to maximise the power efficiency of RF trapping signals applied to the board, as described in section 3.6.1. The coplanar waveguide model shown in Figure 3.24 was used to calculate the appropriate dimensions of the RF traces, while imposing a sufficient gap of 0.9 mm to the ground planes to prevent the risk of electrical breakdown. Table 3.2 shows the dimensions of the traces and the resulting characteristic impedance, Z_0. The values were 63.5 Ω and 63.5 Ω for the piezo side and the heat sink side of the experiment, respectively. The difference in these values was due to the different height, h, of the trace above the underlying ground plane which was considerably smaller for the PCB on the heat sink side.

Dimension	EL125 PCB (PZ)	EL125 PCB (HS)
Width of RF signal trace (a)	0.9 mm	0.9 mm
Length of RF signal trace	46 mm	46 mm
Gap from signal trace to top ground plane (w)	0.5 mm	0.5 mm
Thickness of RF trace (t)	38 µm	38 µm
Height above bottom ground plane (h)	0.76 mm	0.45 mm
Relative dielectric constant (ε_r)	3.76	3.76
Estimate of characteristic impedance, Z_0:	63.5 Ω	50.8 Ω

Table 3.2: Dimensions of RF signal traces on the PCBs designed for second generation chips with magnetic gradient coils (PZ refers to the piezo side of the experiment, while HS refers to the heat sink side). Terms shown in brackets refer to dimensions labelled in Figure 3.24 that were used to calculate the characteristic impedance Z_0 of the RF traces which were adjusted to match the 50 Ω SMP connector.

3.7 Summary

In this chapter I have introduced the two-module ion-trap processor. For each of the component systems that make up the alignable modules, the experimental requirements and design considerations were discussed. I have described the microchip geometry and how the electrode dimensions were optimised for maximum trap depth at a given applied voltage, and to enable ion transport operations. Simulations of the trapping potential above the surface of the aligned microchips informed the requirements of the positioning system, where the goal was to achieve reliable, low-loss shuttling of ions between the traps. The alignment requirement was set at 10 µm for trap separation in each axis to limit the axial pseudopotential barrier that arises within the gap region, which contributes to the ion motional heating rate due to voltage noise on the electrodes. A positioning system

was constructed that was capable of: 600 µm translation in three axes with a precision of better than 1 µm; manual adjustment during installation of the modules into the vacuum chamber to correct for rotational misalignment and account for expected thermal drifts of order 100 µm during baking and cooling of the experiment; and measurement of the displacement and trap separation during module positioning. The next chapter discusses the broader experimental setup within the laboratory.

Chapter 4

Experimental setup

4.1 Introduction

This chapter describes the equipment and control systems that were developed to trap and manipulate Yb ions to perform quantum information processing experiments using the scalable demonstrator device. As an overview, Figure 4.1 shows the individual component parts of the experiment (with an image of the two-module processor at the centre). Each sub-system was designed to fulfil specific experimental requirements and I explain these requirements and how they were met in each section.

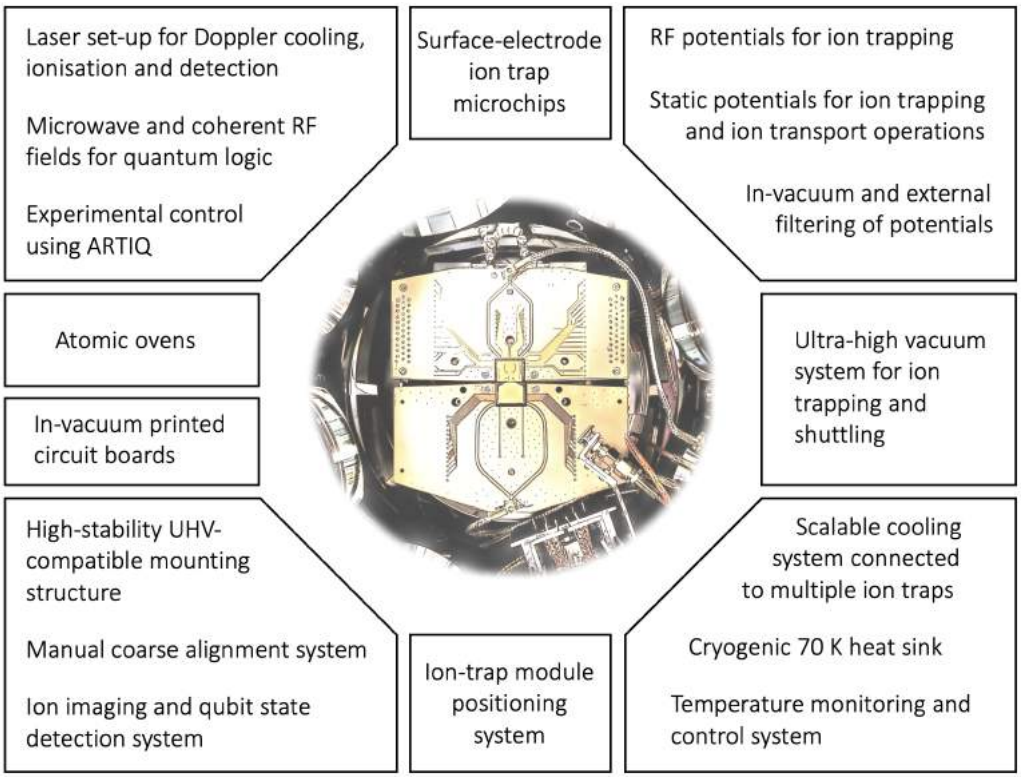

Figure 4.1: An overview of the component parts of the demonstrator experiment.

I begin with an overview of the demonstrator experimental set-up and then explain each component of the experiment in detail. The two-module ion trap processor was described in Chapter 3, including the surface-electrode ion trap chips and the positioning system for module alignment. These are mounted within a vacuum system which provides the necessary isolation from the environment by achieving background pressures in the UHV range of order 10^{-11} mbar. In the following sections of this chapter I explain the electronics for delivering RF and static voltages to the chips including external filtering of signals for noise reduction, the atomic ovens that provide neutral atom beams, the laser-setup for ionisation, Doppler cooling, qubit state preparation and measurement, application of microwave and RF fields for driving logical operations, and a scalable cooling system designed to interface with the demonstrator experiment and three other ion-trap experiments within our laboratory.

4.2 The scalable demonstrator experimental setup

An overview of the layout of the experiment within the laboratory is shown in the schematic diagram of Figure 4.2. The experiment is built onto an optical table[17] with a 400 kg non-magnetic 304L stainless steel breadboard supported by active isolation legs. This provides vibration isolation with extremely low vertical and horizontal resonant frequencies of 1.3 Hz and 1.0 Hz respectively. The non-magnetic breadboard was chosen to minimise ambient magnetic field noise that contributes to dephasing of the magnetically sensitive qubit states. At the centre of the set-up, the vacuum system houses the two-module ion trap processor. External cavity diode lasers are used for ionisation of neutral ytterbium, Doppler cooling of ions and qubit state detection. The laser diodes are located separately from the experiment and the beams are transmitted to the experiment table via optical fibre. On the table, sidebands are added to the lasers for Doppler cooling and qubit initialisation, and spatial filtering provides a Gaussian beam profile prior to the beams being combined and steered into the vacuum chamber and directed at the ion traps. Figure 4.2 also shows the various electronic systems for generating RF and DC trapping potentials, supplying high currents for the on-chip current carrying wires (CCWs), and generating microwave fields for quantum logic. The imaging system is shown mounted above the vacuum chamber and fitted with an sCMOS camera (and PMT) as previously described in Chapter 3. A laminar airflow canopy was installed above the optical table, which consists of two air

[17] Thorlabs Nexus T1225CN non-magnetic 304L stainless steel optical table, PTS603 700 mm tall legs with active isolation.

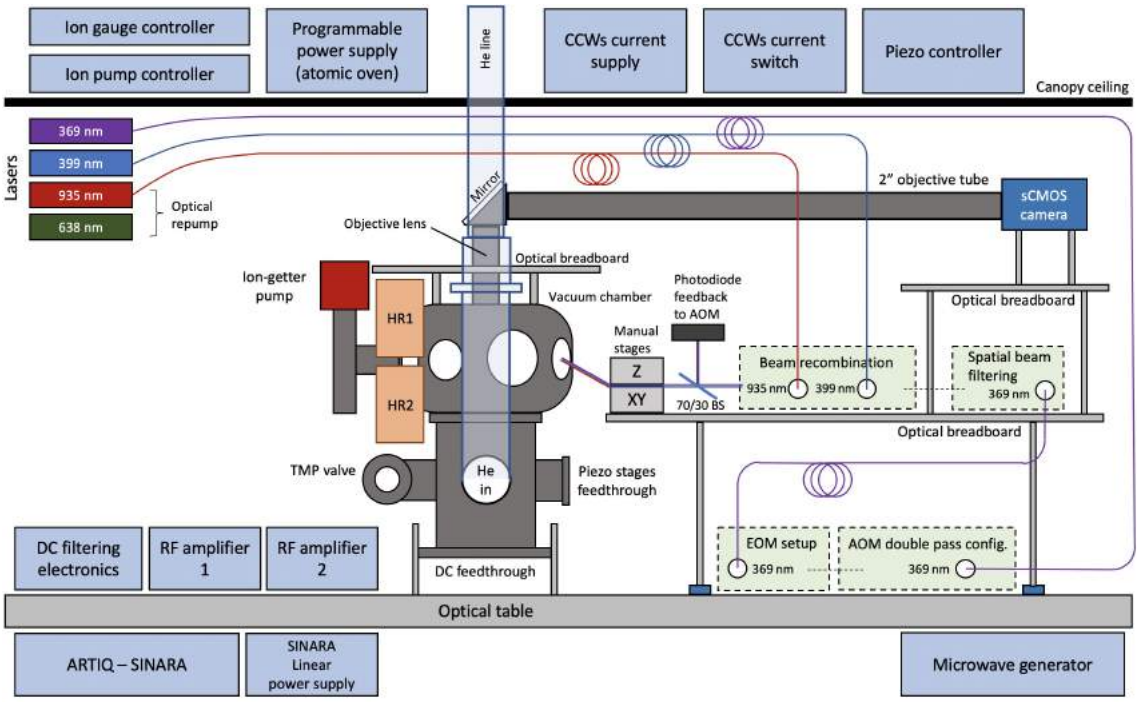

Figure 4.2: Overview of the experimental setup for the demonstrator, showing the vacuum chamber, lasers and optical subsystems, and electronics used for experimental control.

filtration fans fitted with HEPA particle filters in combination with polyvinyl chloride (PVC) anti-static curtains surrounding the table. This provides a continuous flow of dust-free air to the surface of the table to maintain the cleanliness of optical components fitted to the breadboard. The cooling system supplies cold helium gas to the in-vacuum heat sink via transfer lines which pass down through the canopy above the table and connect to cryogenic feedthroughs fitted to the vacuum chamber. In the following sections I describe each of the various parts of the demonstrator experiment in detail.

4.3 Ultra-high vacuum system

In this section I describe the development of a versatile ultra-high vacuum (UHV) system for the two-module ion trap processor. For ion trapping experiments, the most basic requirement of a vacuum system is to provide environmental isolation to minimise the probability of collisions with background gases. The stainless steel chamber also acts as an effective Faraday cage to shield the experiment from external electric field noise. This isolation and protection from external influence serves to increase both the trapping lifetime and coherence times of ions by minimising disturbance to the ion dynamics within the trap and preserving the fragile quantum state. Beyond these requirements, the vacuum system also includes many of the essential subsystems needed to operate the ion trap.

These components must be precisely mounted into position, as well as being powered and controlled by connecting to external electronics by passing signals (or in the case of the cooling system, helium gas) through the sealed chamber.

The vacuum system for the demonstrator experiment was designed to fulfil the following specific requirements:

- Achieve and maintain a chamber pressure of 10^{-11} mbar
- Provide accurate pressure measurement and monitoring
- Allow for mounting of the two-module ion trap processor and associated experimental components within the chamber
- Provide laser access to the ion traps for ionisation of the neutral ytterbium beam emitted by the atomic ovens, Doppler cooling of ions, qubit state preparation and detection
- Provide feedthrough connections for RF and DC potentials applied to the chip electrodes for ion trapping and shuttling, high currents for on-chip magnetic gradient coils, control signals for the piezo-based positioning system, atomic ovens, microwave antenna, and helium gas circulated by the cooling system

The vacuum system was also designed to have sufficient space and connections to accommodate future planned experiments, such as more complex microchip ion traps having multiple junctions and increased numbers of transport electrodes.

Many of the components of the vacuum system were custom developed to meet the requirements of our microchip alignment and quantum information processing (QIP) experiments. Of primary importance is the compatibility of a material with UHV pressures, and in the next section I describe the procedure for selecting suitable materials for use within the vacuum chamber. In the following sections I describe each of the individual elements of the vacuum system in detail.

4.3.1 Qualifying materials for use in UHV

Before any material could be selected for use within the vacuum chamber, it was required to be checked for suitability against the Laser Interferometer Gravitational Wave Observatory (LIGO) Vacuum Compatible Materials List, which is a publicly available document that is continuously updated and maintained by the LIGO collaboration. If a material is not included on the list, then a sample must be tested within a small qualification chamber

to establish what background pressure can be achieved. To be suitable for use within our experiment, a pressure in the 10^{-11} mbar regime must be achieved. Some materials cannot achieve such low pressures because of high vapour pressures owing to sublimation of matter from their surfaces, or because gases are slowly released from inside the material, in a process known as outgassing. These materials must be avoided because a small quantity of outgassing can prevent UHV pressures from being reached.

4.3.2 Vacuum chamber

At the planning stage of the project, consideration was given to materials for the construction of the vacuum chamber, specifically whether titanium could provide an advantage over stainless steel, given the getter properties of titanium alloys[18]. Research discovered that for projects in which titanium had been used as a vacuum chamber material, the reasons behind the choice were concerned primarily with the weight of the chamber, such as for space applications where stringent payload constraints apply. Nevertheless, activation of the getter properties of a titanium chamber would be possible by heating to ∼450 °C. However, that temperature is substantially higher than is achievable with the existing baking oven within the group and so the decision was made to proceed using stainless steel which is a well proven chamber material.

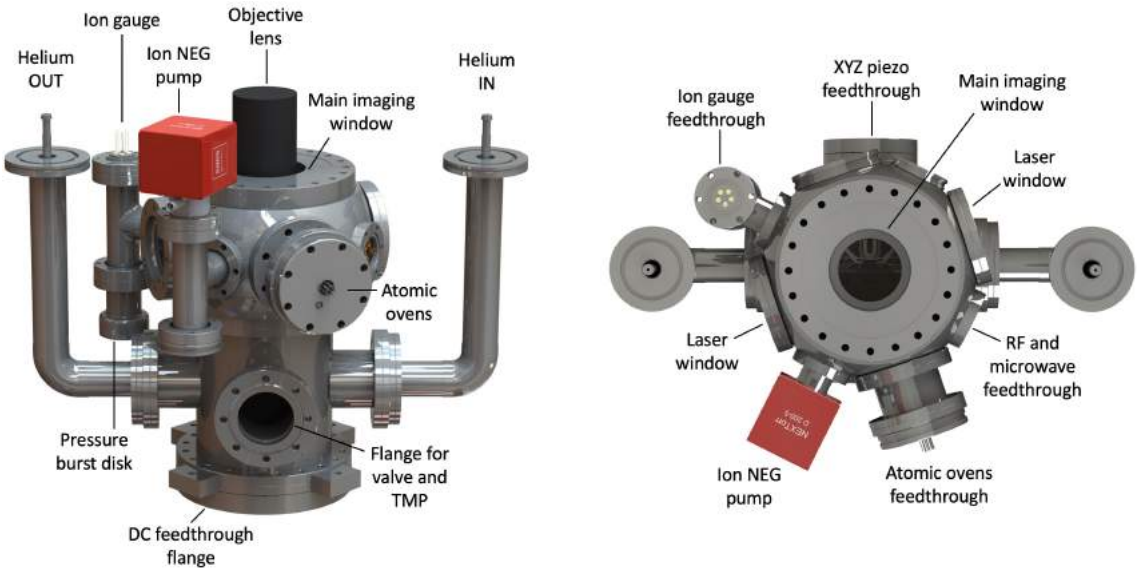

Figure 4.3: Vacuum system set-up for the scalable demonstrator experiment.

A diagram of the fully assembled vacuum system is shown in Figure 4.3. The chamber is centred around an 8.0" spherical square[19] that provides a total of ten Conflat (CF) flanged

[18] Such as are used in a titanium sublimation pump for achieving UHV pressures.
[19] Kimball Physics MCF800-SphSq-G2E4C4, which provides 2x CF160, 4x CF63 and 4x CF40 flanges.

ports of various sizes in an octagonal configuration to deliver the flexibility needed to connect additional vacuum hardware with many orientation options. The spherical square was chosen based on its capacity to accommodate the expected 'footprint' area of the two ion trap modules described later in this chapter. It is machined from a single block of 316L stainless steel[20] using Computer Numerical Control (CNC) milling and is therefore highly leakproof since there are no welds or joins, and the surface is electropolished to create a clean, mirror finish. The spherical square chamber incorporates mounting grooves built into the inside circumference of every flanged window, and these permit components to be fixed into position within the chamber by using proprietary Groove Grabber clamps. To add further volume to the chamber, a six-way cross[21] made from 304L stainless steel[22] is connected below the spherical square, and together these two elements form the core of the vacuum chamber.

As can be seen in Figure 4.3, a wide variety of equipment and individual subsystems attach onto the main chamber, and these can be grouped by function into three main categories: (1) achieving UHV pressures, (2) cryogenic cooling and (3) operating the two-module ion trap processor that resides within the chamber. On the chamber itself, a CF63 flange on the six-way cross connects via an all-metal angle valve[23] to a pumping station consisting of a turbomolecular pump (TMP)[24] and roughing pump[25]. To achieve very low pressures, an ion-NEG pump[26] is used and the pressure within the chamber is monitored with an ion gauge[27]. Two L-shaped cryogenic cooling feedthroughs connect on either side of the six-way cross to supply cold helium gas to the in-vacuum heat sink (as described in Chapter 3). A pressure burst disk[28] is required to protect the vacuum system in case of a high pressure helium leak into the chamber. For ion trapping, laser access is provided by viewport windows made from high transparency fused silica, with an anti-reflective coating to optimise transmission at wavelengths of 369 nm, 399 nm, 935 nm, 638 nm, and 493 nm. The latter two wavelengths relate to the trapping of barium, $^{138}Ba^{+}$, that we plan to use in future experiments for sympathetic cooling of $^{171}Yb^{+}$. On the bottom of the chamber a custom CF160 feedthrough flange provides 400 DC connections for quasi-static voltages that drive the transport electrodes on the ion-trap microchips. Inside the chamber a high-

[20] A low carbon stainless steel with improved corrosion resistance and strength at high temperatures.
[21] Six way cross manufactured by Creative Vacuum. Provides 2x CF160 flanges and 4x CF63 flanges.
[22] Both 316L and 304L grades of stainless steel are UHV compatible alloys.
[23] KJ Lesker CF63 VZCR60R
[24] Leybold Turbovac SL80
[25] Pfeiffer Scrollvac SC5D
[26] SAES NexTorr D200
[27] Epimax PCVX with EMIGTX twin filament Thoria-coated Iridium UHV ion gauge
[28] Allectra 461-PBD-C40-LP DN40CF Pressure Burst Disc, Low Pressure Type, Stainless Steel

stability structure supports the two ion-trap modules. The modules include the ion-trap microchips, the circuitry for in-vacuum routing and filtering of trapping voltages, cryogenic cooling via the heat sink, and thermometry for temperature monitoring. The module positioning system is also housed within the main mounting structure. It is based around a UHV compatible, three-axis piezo translation stage, as described in Section 3.5.2. An array of atomic ovens are internally mounted within the chamber to provide neutral beams of ^{171}Yb and ^{174}Yb atoms for the experiment. A high-frequency feedthrough[29] provides the RF voltage signals required for trapping, and the microwave frequency voltage signals required to drive an internal patch antenna mounted close to the ion trap surfaces. All of these systems are described in the following sections.

4.3.3 Shielded main imaging window

Readout of the qubit state (as described in Chapter 2) requires detection of fluorescence photons from a trapped ion using either a PMT or CCD camera. To achieve high detection fidelity, the optics need a large numerical aperture to collect as many photons as possible from the ion, and this requires a wide diameter objective lens placed close to the ion trap (the detection optics were described in Chapter 3). To support these requirements, the chamber features a large imaging window with a diameter of 72.5 mm, which is recessed down into the chamber by 54 mm to provide close access to the ion trap for the objective lens. The distance between the trap surface and the window is 10 mm,. One issue with having the detection window so close to the ion trap chip is that the fused silica of the window can become a charged dielectric surface, which at such close range can affect the trapping potential and prevent trapping from being achieved, or interfere with ion dynamics within the trap. Charge buildup in the chamber can occur by the photoelectric effect when laser light (particularly at 369 nm and 399 nm) scatters from metal surfaces, since the UV photons provide sufficient energy for electrons to overcome the work function of the metal. To overcome this, the window must be grounded and this is achieved using a high transparency conducting wire mesh[30] as shown in Figure 4.4. The mesh has wire diameter 30 μm and hole diameter 224 μm, giving a 79 % open area.

One drawback of the mesh is that it significantly reduces the transmission of microwave and RF radiation into the chamber that is required for Doppler cooling and quantum state manipulation of $^{171}Yb^{+}$. Therefore an aperture of diameter 16.7 mm is cut into the mesh in the centre of the window, above the position of the ion trap chips. This diameter of

[29] Allectra 242-SMAD18G-C40-4 DN40CF 4x SMA, 50 OHM, 18 GHz
[30] TWP Inc. 100X100T0012W48T 316SS 100 mesh

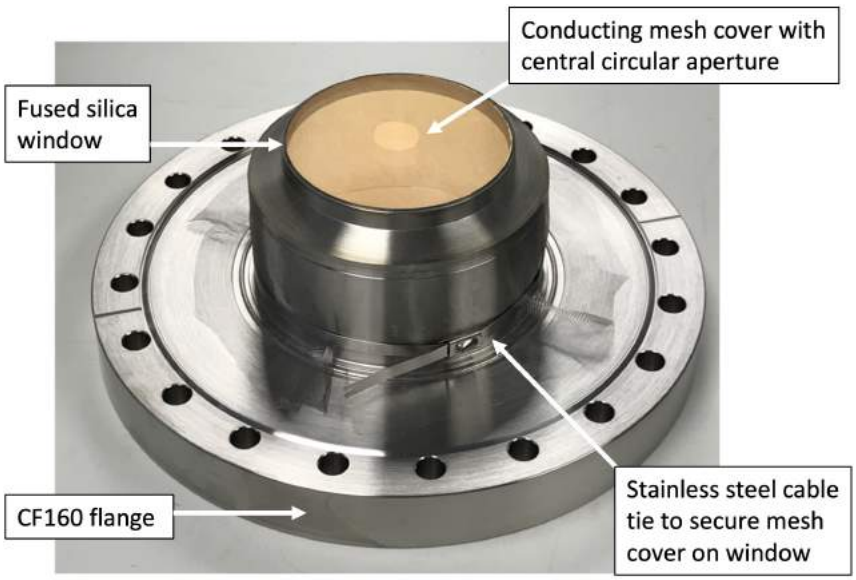

Figure 4.4: The main imaging window is covered by a stainless steel grounding mesh to prevent charge build-up on the dielectric surface close to the ion trap. An aperture of diameter 16 mm was cut into the centre of the mesh to permit RF and microwave fields into the chamber, and to improve the efficiency of fluorescence detection.

aperture is known to be safe from problematic charge buildup, and it was calculated in [90]. The calculation derives the solid angle between the ion and the exposed area of window for an existing ion-trap chamber in the laboratory, that has already been used to trap ions successfully. This is then used as the maximum permitted solid angle for our setup, which allows the maximum diameter of 16.7 mm to be calculated.

4.4 Experimental control using ARTIQ

ARTIQ (Advanced real-time infrastructure for quantum physics) is a modular, scalable platform for experimental control that is used by atomic physics groups, particularly for quantum information processing. We chose to adopt Artiq since it provides a flexible programming environment based on Python that permits automated control of instrumentation and experiment scheduling.

The core ARTIQ system is based around the Sinara hardware, housed within a standard Eurocard rack, which consists of an FPGA master that provides real-time control over a number of daughter devices. The FPGA uses a reference clock signal at 25 MHz to synchronise time-critical outputs for operations such as ion transport and quantum gates. The Urukul signal generator provides four direct digital synthesis (DDS) RF outputs for ion trap drive voltages and broadcasting from the MW and RF antennas. The Zotino 16 bit DACs supply DC voltages for ion confinement and transport, which are passed through

to a Yenesei filter card with second-order low pass filtering, prior to entering the vacuum system. Digital input/output voltages based on fast transistor-to-transistor logic (TTL) pulses are used for controlling laser and microwave pulses for qubit state manipulation.

4.5 RF setup for ion trapping

For experiments using surface trap microchips, high voltage RF signals need to be applied to the chips to generate the ion-trap pseudopotential. RF amplitudes in excess of 200 V, at frequencies between 10-100 MHz, are required to create a sufficient trap depth for ion confinement, given typical surface trap geometries [90, 112]. While it is possible to directly apply voltages from an RF source to a chip, this method suffers from inefficiencies due to the impedance mismatch between the source output and the chip, which results in reflections and the need to supply high RF power. In addition, noise on the RF signal contributes to heating of the trapped ion, and so we wish to filter the signal to attenuate unwanted frequencies. To address these issues, we apply the RF signal to each chip via a helical resonator, as shown in the circuit diagram of Figure 4.5. The figure shows one circuit, but a second circuit is also constructed since there are two ion trap chips.

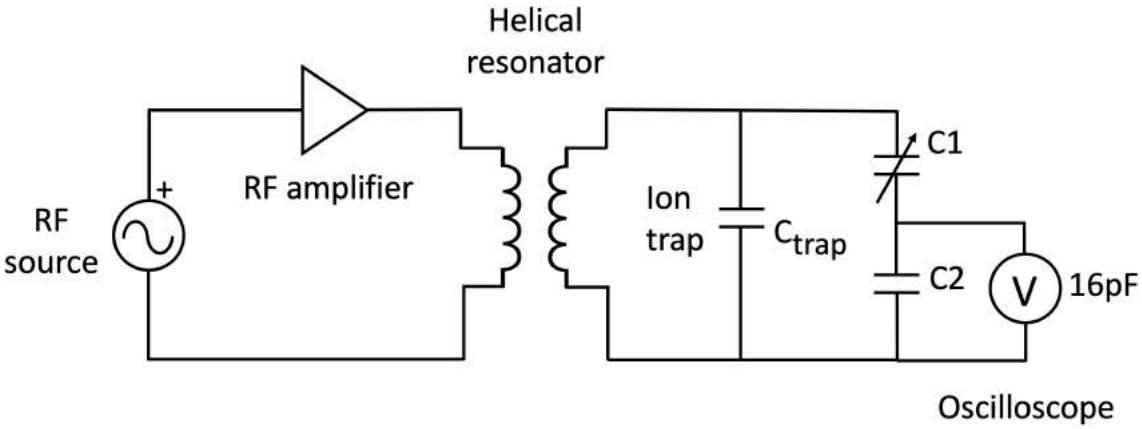

Figure 4.5: RF trapping circuit with helical resonator for impedance matching to the ion trap. The variable capacitor allows for tuning of the resonant frequency of the circuit.

For the RF source we use the Artiq Sinara system which provides RF signal generation using the Urukul AD9910 DDS card with 4-channel SMA outputs. The Urukul card is used to supply RF potentials at a frequency of ∼16 MHz for use as the trap drive frequency. As outlined in Figure 4.5, the signal is passed through a high-power 43 dB amplifier[31] which increases the power to a few W, before passing through a directional power meter[32] to the helical resonator. The signal travels through the resonator which acts as a bandpass

[31] Mini Circuits LZY-22+
[32] Rohde and Schwarz Naus 3

filter centred on the resonant frequency, and is applied to an SMA feedthrough on the vacuum chamber. Inside the chamber, vacuum-compatible SMA cables connect to UHV PCBs, where the signal is routed to the ion trap chips via wire bonds. It is important that both of the in-vacuum cables connecting to the two PCBs are of near identical length to ensure that the phases of the RF signals applied to the chips are equal. To measure the RF voltage applied to each chip, a capacitive divider is used to send a fixed ratio of the voltage signal to an oscilloscope, and this method is described further in section 4.5.2, along with the use of a variable capacitor for tuning the resonant frequency of the circuit.

4.5.1 Helical resonators

A helical resonator can be considered as an LCR circuit, and the ion trap chip adds a capacitance to this circuit that can be impedance matched by the inductance of the resonator. A helical resonator consists of two coils mounted within a grounded shield. Two almost identical resonators were built for the demonstrator experiment, as shown in Figure 4.6. The smaller coil is referred to as the antenna coil, and the larger coil is the pickup coil. Both coils are secured to the resonator shield using BNC connectors. When the resonator lid is attached, the proximity of the coils causes an inductive coupling. An RF voltage signal applied to the antenna coil induces a voltage in the pickup coil, in the same way as for a transformer. The pickup coil is grounded to the shield near the antenna by a low resistance BNC terminator. The other end of the pickup coil is not grounded, and connects to the ion trap chip via the vacuum chamber feedthrough. The grounded shield protects the coils from ambient sources of noise in the laboratory.

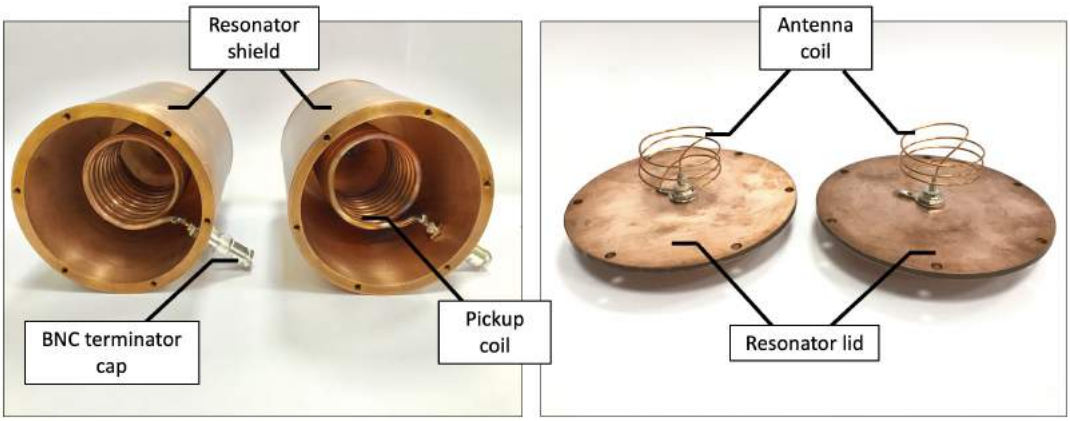

Figure 4.6: The helical resonators used for filtering and amplification of the RF voltage signals applied to the ion trap microchips.

To build the resonators, the framework set out in [121] was followed, which is a the-

oretical analysis of the resonator behaviour and a guide to construction. We can model the resonator as a low resistance, series LCR circuit in which the dominant inductance, L, is provided by the large pickup coil and the capacitance, C, is the sum of the individual parallel capacitances in the circuit (these mainly come from the chip and SMA cables, but also the capacitance of the large pickup coil to the outer shield is significant). The resonant frequency of the circuit is given by

$$\begin{aligned} \omega_0 &= \frac{1}{\sqrt{LC}} \\ &= \frac{1}{\sqrt{L_c(C_t + C_w + C_c + C_s)}}, \end{aligned} \tag{4.1}$$

where L_c is the pickup coil self inductance, C_t is the capacitance of the ion trap, C_w is the capacitance of the SMA cables, C_c is the pickup coil self capacitance, and C_s is the capacitance of the pickup coil to the surrounding shield. To build a resonator for a particular resonant frequency, C_t and C_w can be measured and the other values can be calculated using the framework in [121].

The quality factor, Q, is given by

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}}, \tag{4.2}$$

and so to achieve a high Q, the resistance and capacitance in the circuit must be as small as possible. The main source of resistance comes from the solder joints where the coils attach to the BNC connector, so these were made as large and robust as possible. The SMA cables add significant capacitance to the circuit and so these must be kept as short as possible. On resonance, the Q factor is maximised and the reflected power is minimised, which minimises the heating of the chip and provides strong attenuation of noise signals at frequencies other than the resonant frequency.

The root mean square (RMS) voltage at the chip is given by

$$V_{RMS} = (L/C)^{1/4}\sqrt{PQ}, \tag{4.3}$$

and so to achieve a high voltage with a low power we wish to maximise the Q value of the resonator, and have a large inductance, L, with a low capacitance, C.

To achieve the optimised parameters above, the resonators were constructed from high purity copper, with the following dimensions:

- Shield inner/outer diameter: 89/103 mm

- Shield height: 130 mm
- Pickup coil diameter: 50 mm
- Pickup coil wire diameter: 4 mm
- Number of turns on pickup coil: 9.5
- Pickup coil winding pitch: 8 mm
- Antenna coil diameter: 40 mm
- Number of turns on antenna coil: 3.5

The minimise the resistance of the resonators, the grounded shield has a large wall thickness of 14 mm and the pickup coils are constructed from large 4 mm diameter wire. To solder the wire to the BNC connectors, a blowtorch was used to heat the end of the coil, which was then pre-soldered before being held in position and reheated, while more solder was added to make a strong connection.

Testing of the resonators was performed using a vector network analyser (VNA). A chip analogue with a capacitance of 24 pF was used for the tests, to represent a typical linear chip. The chip analogue was made from a capacitor connected to an SMA plug so that it could be attached to an SMA cable for measurement. The chip analogue was connected to the output of the resonators by a 10 pF SMA cable and the VNA was used to input a -10 dBm signal to each resonator, that was scanned to find the resonant frequency. The resonant frequencies were measured to be 17.58 MHz and 17.81 MHz. The Q factors were measured to be 81 and 85.

4.5.2 Voltage measurement and frequency tuning

To measure the high voltages applied to the chip, a capacitive divider is set up to send a fixed ratio of the voltage to an oscilloscope, as shown in Figure 4.5. The capacitive divider is also required to shield the resonant circuit from the large capacitance of the oscilloscope, which would otherwise add 16 pF to the capacitance of the circuit if connected in parallel with the ion trap. Instead, we use the divider with C1 $\approx$ 5 pF and C2 = 400 pF. For these values, the contribution to the total capacitance is

$$\begin{aligned} C_{total} &= C_{trap} + \left(\frac{1}{C_1} + \frac{1}{(C_2 + C_{scope})}\right)^{-1} \\ &= C_{trap} + \left(\frac{1}{5\,pF} + \frac{1}{(400\,pF + 16\,pF)}\right)^{-1} \\ &= C_{trap} + 4.9\,pF. \end{aligned} \tag{4.4}$$

The smaller value of C1 = 5 pF limits the value of the entire network formed by C1, C2 and C_{scope}. If this capacitor is variable, then it allows the capacitance of the circuit to be tuned, which by Equation 4.1 will tune the resonant frequency of the circuit. This is useful since the two chips can have significantly different capacitances, with the linear chip expected to be ~ 20 pF and the X-junction chip ~ 40 pF capacitance. Tuning the circuits into resonance therefore permits the resonators to be used to efficiently drive the chips at the same frequency.

The experiment therefore uses a variable capacitor[33] for C1 which has a capacitance measured to be between 5.9 pF and 34 pF. The large value of C2 = 400 pF ensures that even at the higher end of this range, the voltage ratio sent to the oscilloscope remains less than 10 %. The resonators were tuned by a frequency range of more than 3 MHz using the method described above, but large changes in frequency are achieved at the expense of resonator Q value, and a reduction in the voltage amplitude applied to the trap at a given power.

4.6 DC setup for generating ion trap control potentials

The ARTIQ Zotino DACs are used to provide DC potentials for static endcap electrodes and for shuttling of ions using time-dependent voltage waveforms applied to trap electrodes. An overview of the setup used for the experiment is shown in Figure 4.7.

4.7 Atomic ovens

An atomic oven is a device filled with an atomic sample that when heated produces a beam of neutral atoms to supply the ion-trap. Atoms within the beam must first be ionised using the 399 nm and 369 nm lasers in a two step process (that was described in Chapter 2) to be confined within the ion-trap potential. To achieve this, the two lasers are made to overlap and are directed to pass through the potential minimum, where the atomic beam is also targeted. Upon passing through the lasers, the neutral atoms are ionised which

[33] Knowles Voltronics JZ300HV

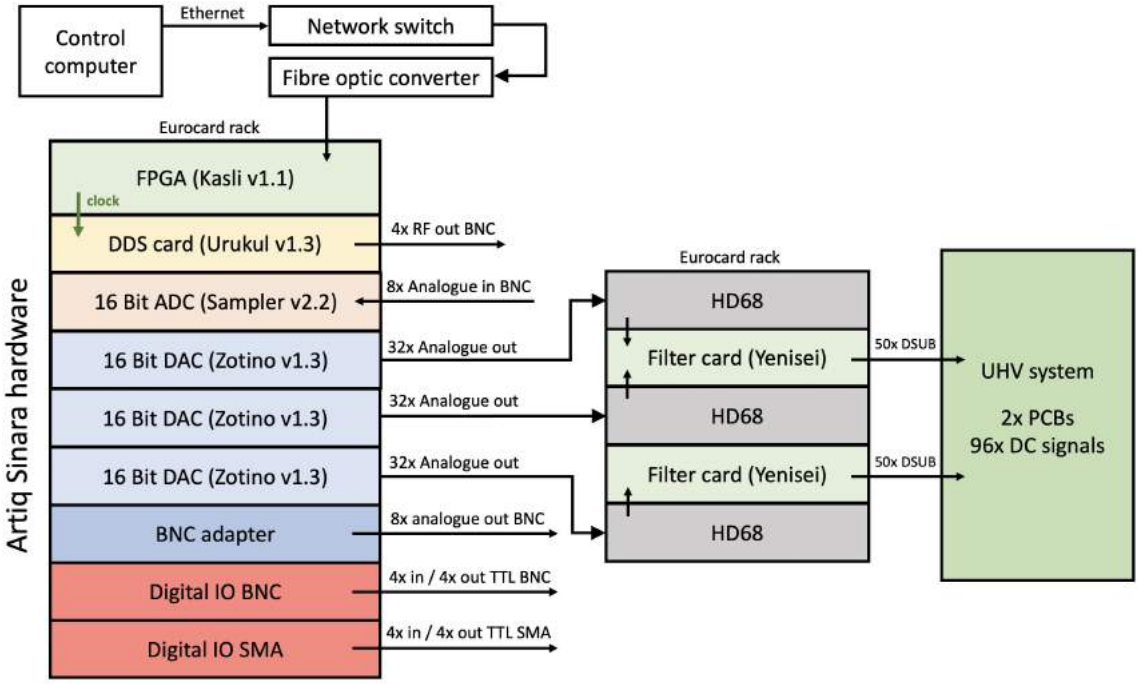

Figure 4.7: Artiq Zotino DACs are used to generate ion trap control potentials for endcap electrodes and ion transport operations. The Zotino cards have built-in low pass filtering with cut off frequency of 72 kHz. Additional second-order low-pass filtering is applied to the DC signals prior to entering the vacuum chamber using the Yenesei filter card also with cut-off frequency of 72 kHz. Shielded cables are used to minimise noise pickup between the filters and the vacuum chamber.

immediately causes them to become trapped by the potential generated by the chip.

Three ovens were constructed for the experiment: two ovens for ytterbium in natural abundance (that can be used as a source of ^{174}Yb by tuning the wavelength of the applied ionisation laser for isotope selectivity as described in Chapter 2), one for enriched ^{171}Yb and one for barium in natural abundance. The ovens feature temperature monitoring for fast fluorescence control using a programmable power supply controlled by a Labview script. The ovens were mounted into a custom designed support structure that orientates the tubes to target the ion traps, collimates the atomic flux, and provides thermal management to limit the temperature of the structure during operation of the ovens to reduce radiative heating of the experiment. The complete oven assembly is shown in Figure 4.8.

In the following sections I describe how the ovens were constructed and tested.

4.7.1 Construction of the atomic ovens

Each ytterbium oven is made from a 316L stainless steel tube with 1.5 mm outer diameter (OD) and 1.3 mm inner diameter (ID). The barium oven was built by Mitchell Peaks, using a larger 316 stainless steel tube with 3.0 mm OD and 2.4 mm ID. The reason for the larger

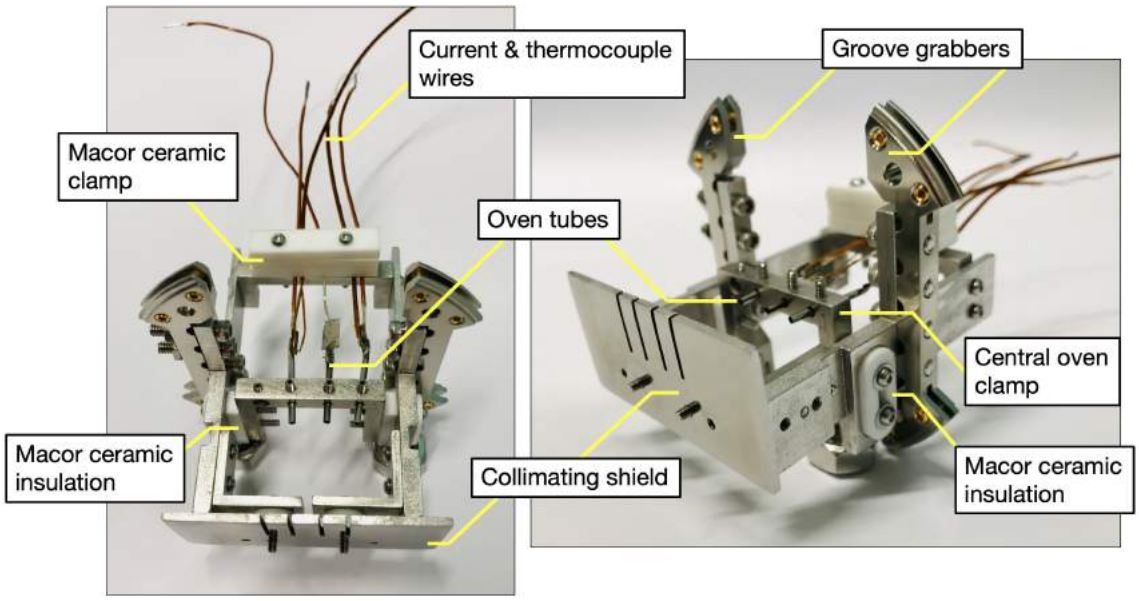

Figure 4.8: Top and side views of the atomic ovens and mounting structure. The barium oven tube is not shown. The structure is designed to attach to the inside of a CF63 circular window on the vacuum chamber, where it is secured in position using groove grabbers. A shielding plate with 0.8 mm wide slits is used to collimate the atomic beams emitted from the oven tubes.

diameter was to support fast loading of the sample material into the oven (and then into the vacuum chamber) since barium reacts relatively quickly in air to produce a surface layer of oxide. Other than the sizes, the ytterbium and barium ovens were constructed using the same method, as described below.

A length of stainless steel tube is cut to size (20 mm and 30 mm for ytterbium and barium ovens respectively) and one end is crimped closed to create a flat tab. Onto this tab a rectangular piece of high resistivity constantan foil is spot welded, which assists in heating the oven when current is applied. A length of silver plated copper wire[34] is then spot welded to the constantan foil, with the position of the spot weld made as close as possible to the closed end of the tube, which is where the atomic sample was located after loading the oven. Finally, a type K thermocouple[35] is spot-welded onto the oven, to measure the temperature. After conducting fluorescence tests of the ovens (described below), the minimum recorded fluorescence temperature was then used as a target for a Labview script designed to hunt the required temperature as quickly as possible, by controlling the current from a programmable power supply.

[34] LewVac KAPW1X0813 kapton insulated single core 600 VAC 2 kVDC
[35] Allectra 311-KAP-TCK-5M type K, kapton insulated

4.7.2 Loading of the ovens

Before loading with an atomic sample, the ovens are cleaned by immersing them into a beaker of isopropyl alcohol and sonicating in an ultrasonic bath for 10 minutes. After leaving to dry, the ovens are loaded with small fragments of ytterbium metal, prepared using a scalpel to cut pieces from an ytterbium sample of natural abundance (31.8 % ^{174}Yb) or enriched ^{171}Yb. The ovens are loaded just prior to being mounted in the vacuum chamber, since this minimises the possibility of losing ytterbium from inside the oven and ensures that the oven is clean and free from contamination.

4.7.3 Oven mounting structure

The oven mounting structure fulfils four main purposes, (1) to hold the ovens in position within the vacuum chamber, in the correct orientation so as to target each of the atomic beams at the ion traps, (2) to collimate the atomic flux from the ovens to reduce the likelihood of coating the ion trap surfaces with ytterbium, (3) to thermally insulate the ovens from the rest of the structure so as to prevent heating of the collimation shield and other in-vacuum components close to the ion traps, and (4) to provide a common electrical ground connection for each oven.

The mounting structure is constructed from 316 stainless steel, and it consists of the following pieces: two 'n' shaped cross-beam oven clamps, one to hold the oven tubes and one to clamp the trailing wires; two side arms that connect onto groove grabbers[36] for attaching to the vacuum chamber; and a slotted shield for collimation of the atomic beam, which connects to the front of the structure by two angle brackets. The central 'n' shaped clamp has four holes into which the open ends of the oven tubes are passed, and then secured by tightening a grub screw above each hole. The collimation slits in the shield are 0.8 mm wide, and the positioning of the shield between the ovens and the ion traps results in atomic beams of ∼3 mm width at the location of the chips. This width is sufficiently broad to cover all arms of the X-junction and to permit trapping on any region of the chips.

The ovens are insulated from the rest of the mounting structure by macor ceramic spacers placed either side of the central oven clamp. The sides struts of the mounting structure are thermally anchored to the large thermal mass of the vacuum chamber, which remains close to room temperature, and this helps to limit the temperature rise due to heating of the ovens. It is undesirable to raise the temperature of the collimation shield,

[36] Kimball Physics 4.5" MCF450-GrvGrb-C02

since it will contribute to radiative heating within the chamber, and so the shield is also thermally anchored to the chamber by an OFHC copper braid.

4.7.4 Oven skimming

An effective technique to prevent the atomic beam from coating the surfaces of the chips is referred to as 'skimming', and it has previously been implemented by Dr David Murgia [90]. It is also used here, in addition to the collimation shield. Although the collimation shield is designed to reduce the left-right spread of the atomic flux from the ovens, the skimming method is intended to shadow the chip surfaces from the atomic beam, up to a given height above the surface. This is achieved by positioning the apertures of the oven tubes slightly below the height of the chips, as shown in figure 4.9.

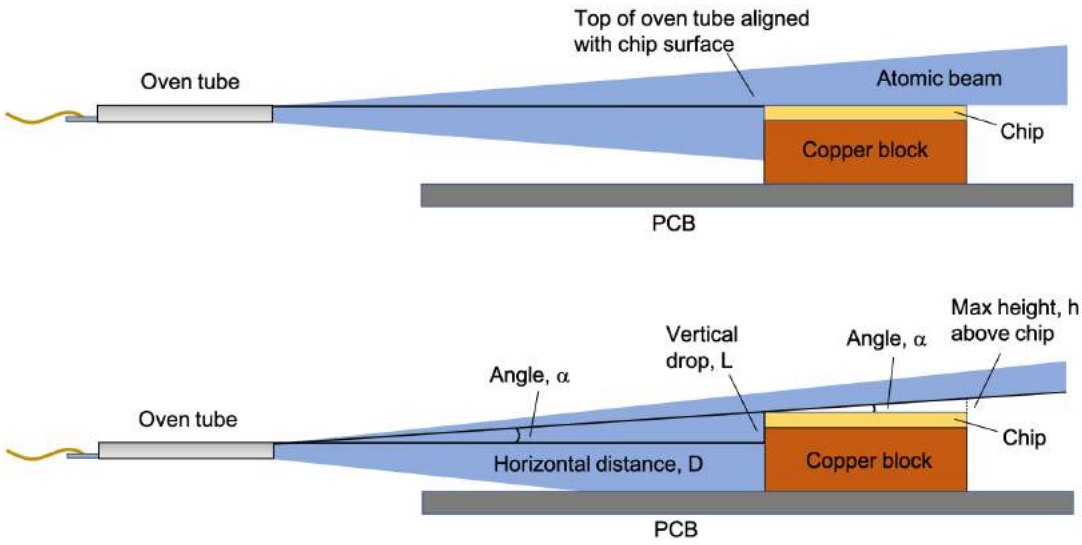

Figure 4.9: Skimming method used to prevent the beam of neutral Yb atoms from the atomic ovens from coating the microchip surface.

The mounting structure is designed to allow for small final adjustments to the height of the oven clamp, after the assembly has been fitted into the vacuum chamber, in order to accurately target the ovens at the chip mount to implement skimming. Given the ion height of 150 µm above the chip surface, and using the known measurements for the width of the chip which is close to 10 mm across, and the horizontal distance, D, of 80 mm between the oven and the copper block, we can calculate the maximum permitted vertical drop, L. The expression for L is:

$$L = D \times \tan\alpha = 80\,\text{mm} \times \frac{0.15\,\text{mm}}{10\,\text{mm}}. \tag{4.5}$$

This gives a maximum drop in the position of the ovens of 1.2 mm in order to ensure atomic flux passes above the chip surface at the required ion height. To achieve this fine adjustment, one can observe the position of the oven aperture by eye, looking into the vacuum chamber across the chip surface, with the chip surface at eye level. Once the top of the oven tube is observed to disappear just below the chip surface, then the oven clamp is fastened into place.

4.7.5 Fluorescence testing

Fluorescence tests were conducted to verify that the atomic ovens could be operated successfully, and to compare thermocouple measurements from different junction positions on the ovens. The reason for making this comparison was to explore the finding from [106] that a significant error in reading the thermocouple output can occur if electric current applied to the oven tube travels along a small section of the thermocouple wire where the junction makes contact with the tube. This is based on the reasoning that since the small section of thermocouple wire has finite resistance, then any current flow implies a small

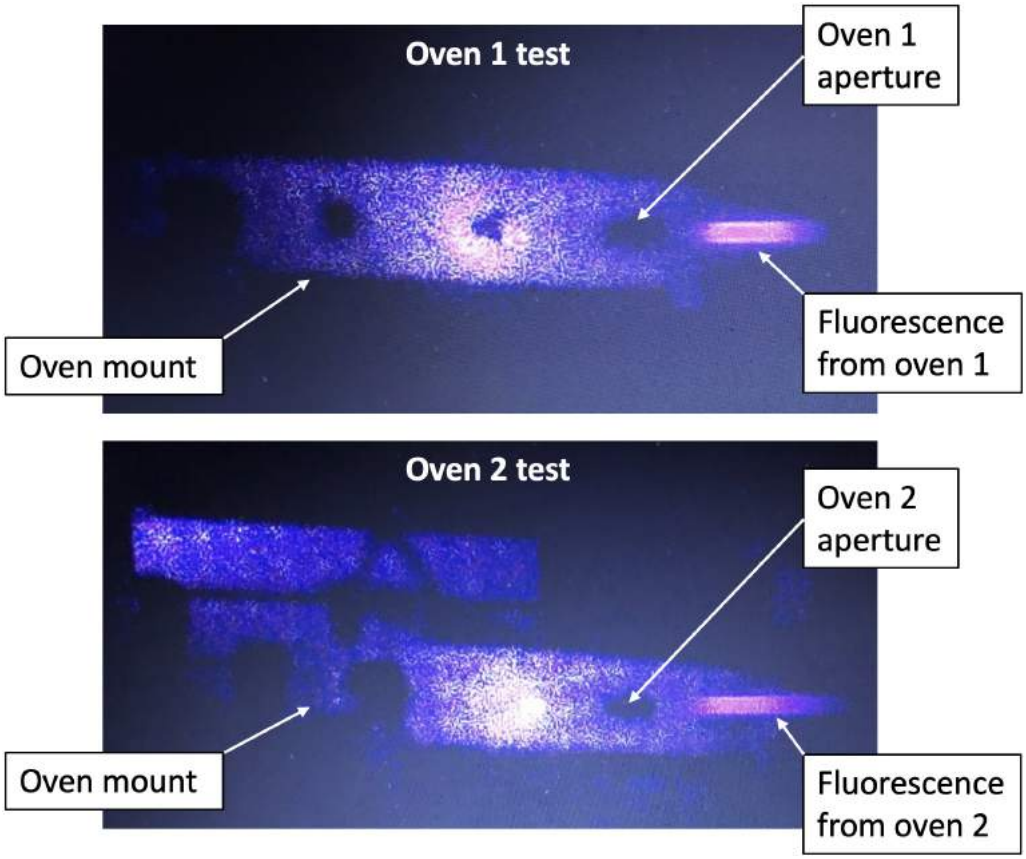

Figure 4.10: Laser light with wavelength near 399 nm was used to excite fluorescence in beams of neutral ^{174}Yb emitted from atomic ovens mounted adjacent to one another during testing.

voltage drop across the wire and this causes a misreading of the thermocouple. At the time of the tests, two identical ovens (here referred to as oven 1 and oven 2) had been constructed and loaded with ytterbium in natural abundance. A thermocouple was attached to each oven by spot welding, so that for oven 1 the junction was welded directly onto the surface, while for oven 2 the junction was displaced by approximately 2 mm from the same position so that it didn't contact the surface, and the wire was spot welded just past the junction and across the tube at an angle of 90 degrees. The ovens were then secured into the mounting structure and fitted into the chamber for testing. A webcam was set-up to observe into the chamber to record fluorescence.

After pumping down to a pressure below 10^{-5} mbar, the 399 nm laser was first directed across the aperture of oven 1 and the current supply to the oven was slowly ramped up while scanning the frequency of the laser. Fluorescence was observed at a temperature reading of 402 °C and a minimum current of 7.62 A. Fluorescence for oven 2 was observed at 158 °C and a minimum current of 7.79 A. Figure 4.10 shows the atomic fluorescence from the ovens. The unexpected low temperature of oven 1 was well explained by the 2 mm offset of the thermocouple junction from the surface, and was attributed to the temperature gradient along the section of wire. Since this offset method was also a more fragile way to connect the thermocouple to the oven, we decided to take forward the more robust on-surface method for spot-welding of the junction.

Although each oven was run continuously at high temperature for durations of greater than 30 minutes, and the two ovens were separated by only 5 mm, the rise in temperature of the stainless steel mounting clamp was not observed to cause one oven to trigger the other during the tests.

4.8 Laser set-up

As described in Chapter 2, three principal lasers with wavelengths of 369 nm, 399 nm and 935 nm are needed to ionise, Doppler cool, manipulate and detect the quantum state of $^{171}Yb^{+}$ ions. A fourth laser can be used during Doppler cooling to address the $^{2}F_{7/2}$ state of the ion after collisions with background gas, but this was observed to occur so infrequently in other experiments within the group that its use was discontinued with no adverse effect.

369 nm laser

This is an M Squared laser assembly with output wavelength of 369.5 nm. It consists of a Lighthouse Photonics Sprout-G pump laser, an M Squared Ti:sapphire SolsTiS module that outputs 739 nm, and an M Squared ECD-X frequency doubler that is used to achieve 369.5 nm light at an output power of 1.3 W. The experiment requires 100 µW of laser power at this wavelength for ionisation, Doppler cooling and fluorescence for state detection.

399 nm and 935 nm lasers

These are Toptica Photonics external cavity diode lasers with output wavelengths of 398.9 nm and 935.2 nm. For the 398.9 nm laser[37], the experiment requires around 500 µW of power for the first stage of the two-step ionisation process described in Chapter 2. For the 935.2 nm laser[38], the experiment requires around 8 mW for the Doppler cooling repump in order to be power-broadened over the relevant hyperfine states in $^{171}Yb^{+}$.

The lasers are located on a separate optical table since they are shared with an adjacent experiment, and so the beams are distributed to our demonstrator experiment by coupling into polarisation maintaining optical fibres. The laser set-up for the experiment was constructed by Mariam Akhtar and Raphaël Lebrun-Gallagher. The fibres and optical components were distributed over two height levels, owing to the mounting position and height of the vacuum chamber. The lower level was on the optical table surface and the upper level was supported on a raised breadboard at a height that brings the beams in line with the viewport windows of the chamber. Figure 4.11 shows the laser set-up at the lower level, while Figure 4.12 shows the upper level set-up on the raised breadboard.

4.8.1 Optics lower level - AOM and EOM setup

On the lower level of the optical setup, the 369 nm laser beam emerges from an optical fibre and passes through an iris to control the beam width before entering into the acousto-optic modulator[39] (AOM) in a double pass configuration. The AOM allows the laser frequency to be detuned from resonance with the $^{2}S_{1/2} \leftrightarrow {}^{2}P_{1/2}$ transition by ∼MHz for Doppler cooling. After this, the laser beam passes through an electro-optic modulator[40] (EOM) to add sidebands of ∼1 GHz to address the hyperfine levels in $^{171}Yb^{+}$ to form a closed

[37] TOPTICA DL pro HP 029050
[38] TOPTICA DL pro 020275
[39] A/R coated Isomet 1206C-833
[40] QUBIG EO-T1055M3

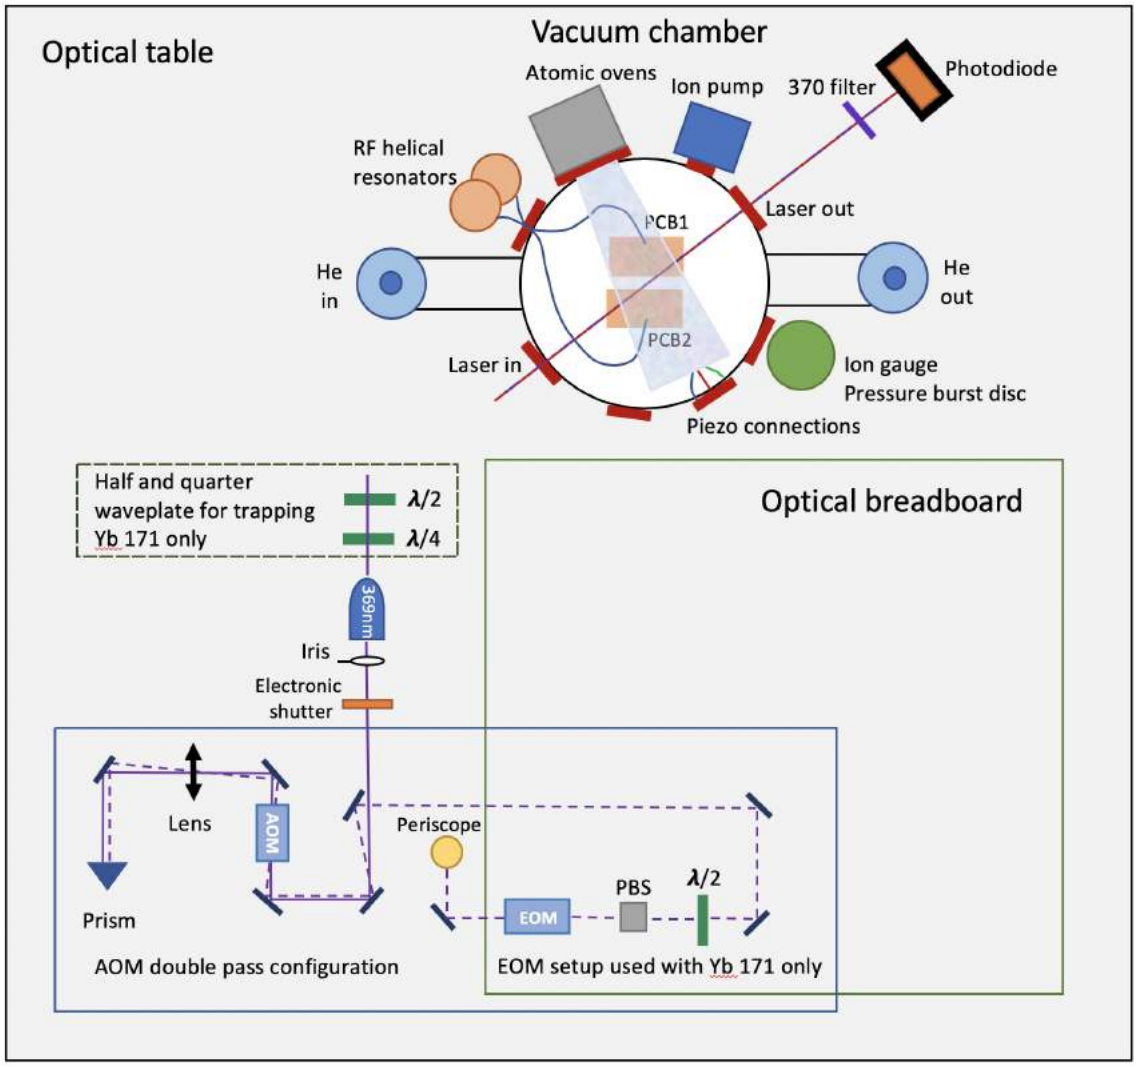

Figure 4.11: Laser setup at optical table level (lower level).

Doppler cooling cycle, as described in Chapter 2. The electronic shutters allow for fast on-off control of the beam for heating rate measurements and quantum logic experiments.

4.8.2 Optics upper level - filtering and beam combination

After the AOM and EOM, the 369 nm beam is periscoped up to the optical breadboard level. Here the beam is filtered using two lenses and a pinhole to produce a Gaussian beam profile. Also at this level, the 399 nm and 935 nm laser beams emerge with Gaussian beam profiles from single mode fibres and no spatial filtering is required. To combine the 369 nm and 399 nm laser beams, a band pass filter is used that transmits 369 nm light and acts as a mirror to reflect the 399 nm light. Next, the 935 nm beam is overlapped with the other beams using a dichroic mirror, before the lasers are directed into the vacuum chamber and targeted to pass above the ion trap chips. This is achieved using a three-axis translation stage to position the beams through a viewport window, and an achromatic lens is also

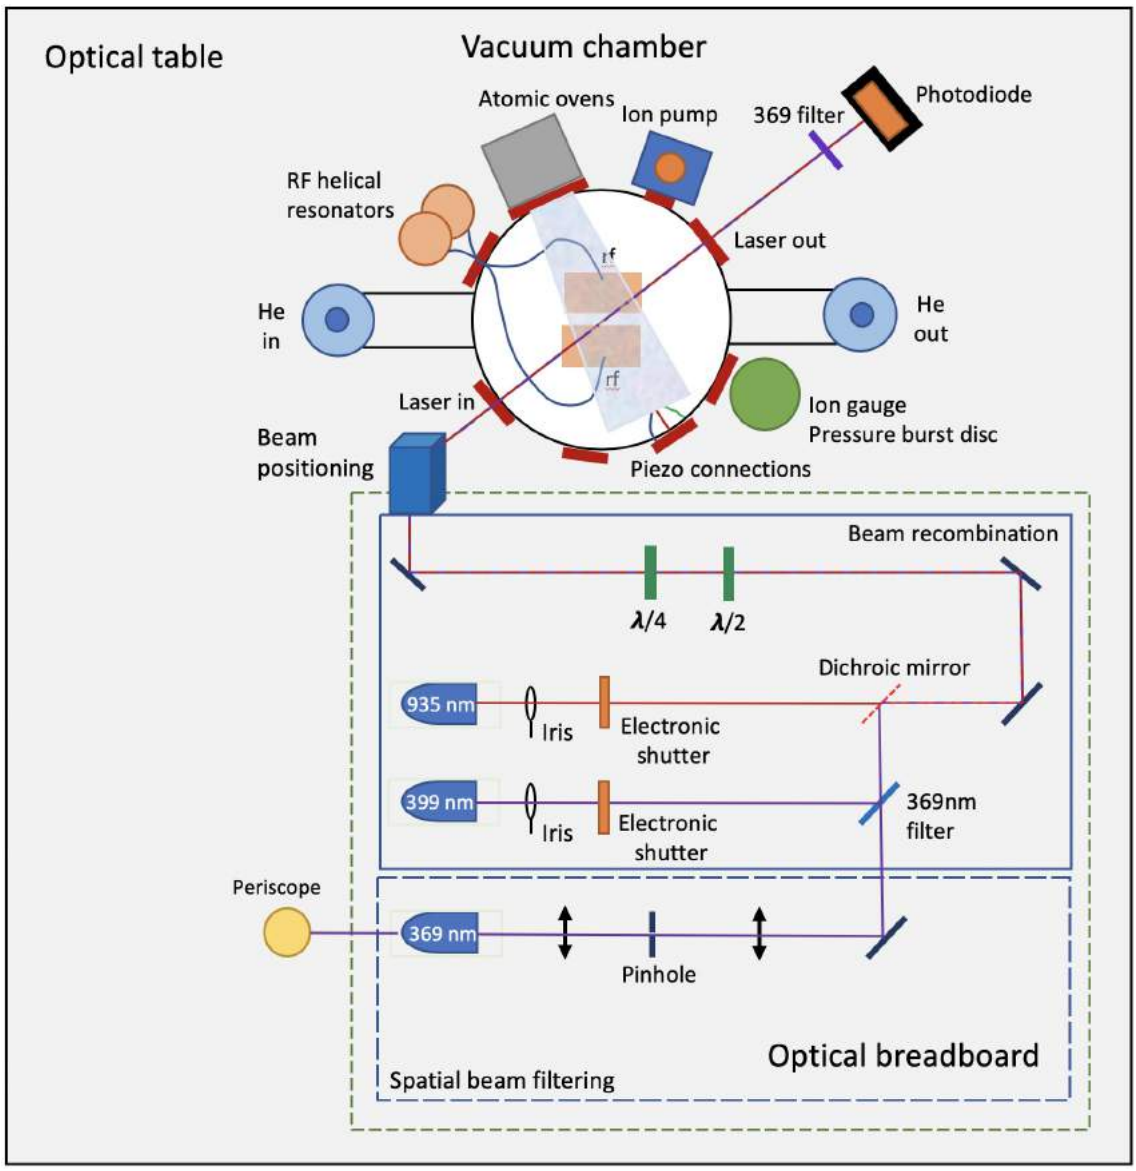

Figure 4.12: Laser setup on the raised optical breadboard (upper level).

used (not shown in Figure 4.12) to focus the beams before they enter the chamber. This results in a narrow beam diameter (∼20 µm) which is necessary to prevent scattering of laser light from the surface of the ion traps. The beams emerge from the opposite viewport window and pass through a 369 nm bandpass filter to a photodiode. The signal from the photodiode is used to stabilise the output from the AOM using a Labview program to control the RF input power.

4.9 Microwave and RF setup for qubit logic

Microwave and RF radiation fields with frequencies close to 12.64 GHz are required to address the atomic transitions between states within the $^2S_{1/2}$ hyperfine manifold of $^{171}Yb^+$,

for Doppler cooling and coherent manipulation of qubits as described in Chapter 2. This section describes the setup used to generate and broadcast the fields to ions in the demonstrator experiment. The microwave and RF set-up is shown in Figure 4.13.

4.9.1 Microwave setup

The microwave setup is shown in the upper part of Figure 4.13. A 1 channel programmable Arbitrary Waveform Generator (AWG) is used to send a 100 MHz RF signal (that is monitored using an oscilloscope) to a pre-amplifier with 18 dB gain and then through a high isolation RF switch to a Vector Signal Generator (VSG). The RF switch can be controlled by the experimental computer using a TTL signal from the Artiq Sinara system. At the VSG the 100 MHz waveform is mixed with a 12.54 GHz microwave carrier to produce sidebands at 12.44 GHz and 12.64 GHz. The relative power of the carrier and sideband signals can be controlled by the VSG to optimise for the 12.64 GHz signal. Experiments prioritising high-fidelity gates require high powers for fast gate operations in order to limit decoherence over the gate duration. In that case a bandpass filter can be added to remove power from the unwanted 12.54 GHz carrier signal and 12.44 GHz sideband before high-gain amplification of the signal. In our case, this is not required since the experiment prioritises

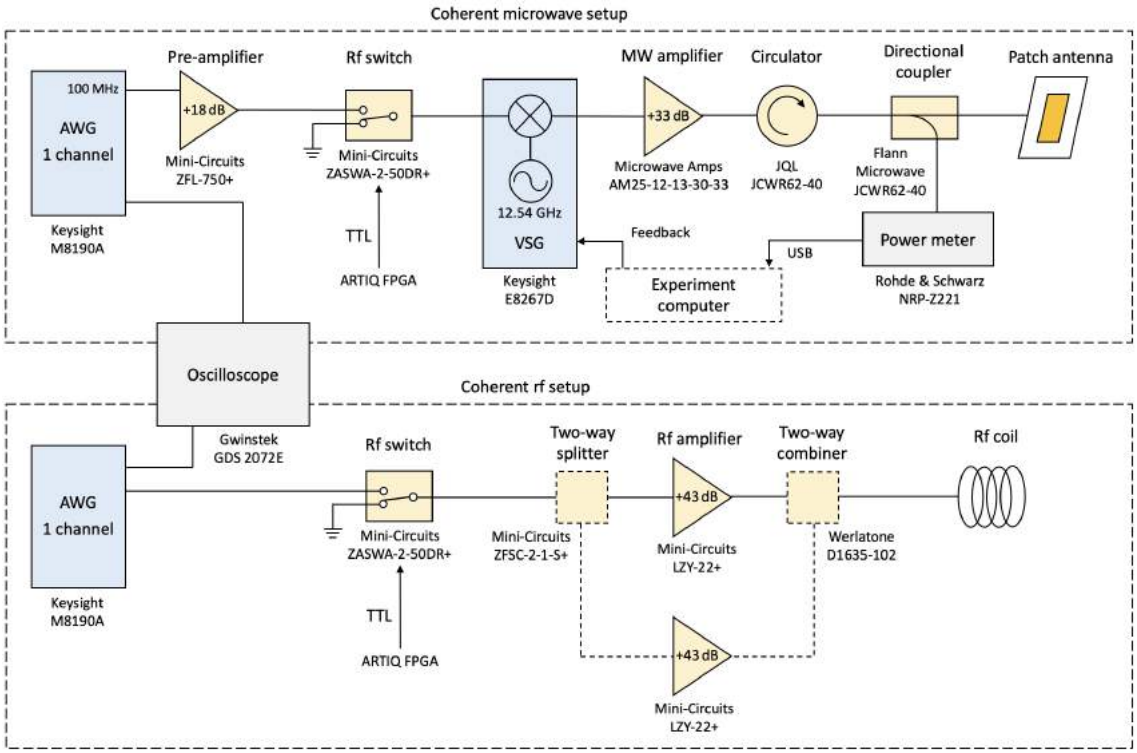

Figure 4.13: The experimental set-up for generating and broadcasting microwave and RF fields for coherent manipulation of $^{171}Yb^{+}$ hyperfine qubits.

ion shuttling between trap modules. Therefore the combined signals are sent directly to a microwave amplifier with +33 dB gain and some power is lost to the unnecessary frequencies. The amplified signal is sent first to a circulator to prevent power being reflected back to the amplifier from the antenna. Next, a waveguide directional coupler is inserted to tap off -40 dB to a power meter that can be used to provide feedback via the experimental computer to the VSG for stabilisation of the output power of the 12.64 GHz signal. Finally the signal is sent to the microwave patch antenna mounted inside the vacuum chamber.

4.9.2 Coherent RF setup

The setup that was used to generate coherent RF radiation for the experiment is shown in the lower portion of Figure 4.13. A programmable, 1 channel AWG is used to generate the RF waveform and, as before, one output of the signal is monitored using an oscilloscope while the other output is sent to a high-isolation RF switch with TTL control. From the switch the RF signal is amplified by a single high-power amplifier with +43 dB gain. Additional power may be realised by connecting two high power amplifiers in parallel which is beneficial for high fidelity gates owing to reduced gate times. This was not utilised for our experiments. The RF signal(s) are then combined (if required) using a high power two way combiner. Finally the coherent RF signal is passed to a 3-turn antenna coil inserted into the recessed imaging window of the vacuum system.

4.9.3 Microwave patch antenna

To address the qubit transitions within the $^2S_{1/2}$ hyperfine manifold of $^{171}Yb^+$, previous experiments in the Ion Quantum Technology research group have used a microwave horn external to the vacuum chamber [90], positioned as close to the main imaging window as possible in an effort to achieve a large field strength at the ion position. This approach is not possible with the demonstrator experiment because the large diameter objective lens of the imaging system almost completely fills the space inside the recessed main imaging window. Therefore, an in-vacuum antenna was required for the experiment. The choice of a patch antenna as an in-vacuum emitter was driven by its small size, its ease of construction using UHV compatible materials and its well-known directional radiation pattern.

A schematic of the antenna constructed for the experiment is shown in Figure 4.14. This type of patch antenna is called a probe-fed antenna because the signal is supplied by the central conductor, or probe, of a coaxial cable. The probe-fed design is compact and well-suited to mounting within the constrained space of the vacuum chamber. By placing

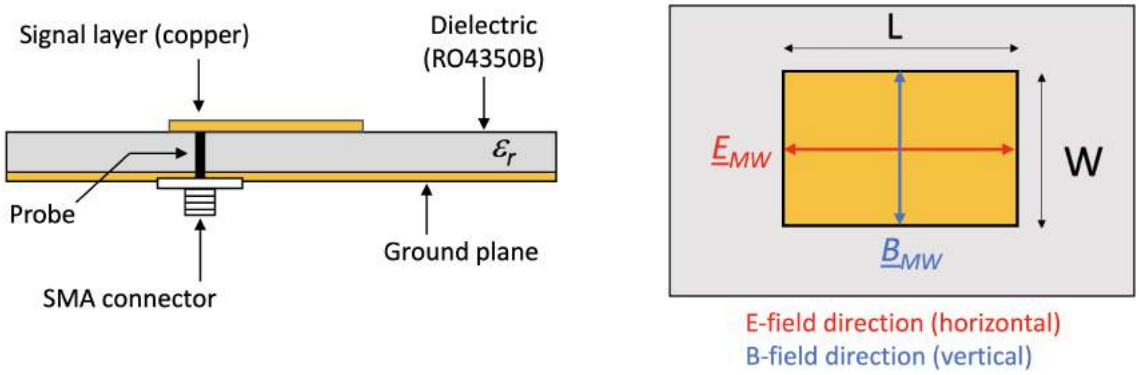

Figure 4.14: Schematic diagram showing the key features of a probe-fed patch antenna. (*left*) Side view of the antenna showing the backside SMA connection with probe passing to the front signal layer. (*right*) Top view showing the resonant length, L, and the impedance width, W. The length determines the wavelength of the radiation emitted from the patch which is a plane wave with linear polarisation where the E-field direction of oscillation is parallel to the length and the B-field direction is parallel to the width.

an SMA connector on the backside ground plane the patch can be positioned close to, and facing, the ion trap chips to maximise the field strength at the ion position.

For the patch described here, the substrate is Rogers RO4350B material with thickness 0.5 mm, and a relative dielectric constant ϵ_r of 3.6. For a rectangular patch antenna as shown in Figure 4.14, the top conducting surface is effectively a half-wave antenna, and the length, L, sets the frequency based on the modified wavelength of the radiation within the dielectric substrate. The width, W, determines the impedance of the patch along the L direction, and it is necessary to calculate the position along the length to attach the probe in order to impedance match to the coaxial cable at 50 Ω.

To calculate the resonant length and impedance width required of the patch, I will use the patch transmission line model outlined in [122]. The width is calculated using,

$$W = \frac{c}{2f_r}\sqrt{\frac{2}{\epsilon_r + 1}}, \tag{4.6}$$

where W is the width, c is the speed of light in free space, f_r is the free space frequency of the radiation, and ϵ_r is the relative dielectric constant. For a frequency of 12.6 GHz and a dielectric constant of 3.6, the width is 7.8 mm.

Fringing fields at the edge of the patch enlarge the effective size, so that we need to calculate the effective dielectric constant, which is

$$\epsilon_{rEFF} = \frac{\epsilon_r + 1}{2} + \frac{\epsilon_r - 1}{2}\left(1 + 12\frac{h}{W}\right)^{-1/2}, \tag{4.7}$$

where h is the thickness of the dielectric, in this case 0.5 mm. This yields an effective

dielectric constant of $\epsilon_{rEFF} = 3.28$.

Now we need to calculate the increase in length due to the fringing field, which is given by

$$\Delta L = 0.412 \frac{(\epsilon_{rEFF} + 0.3)}{(\epsilon_{rEFF} - 0.258)} \frac{(\frac{W}{h} + 0.264)}{(\frac{W}{h} + 0.8)} h, \tag{4.8}$$

and this results in a length extension of 0.573 mm. This permits the calculation of the final length of the patch, using

$$L = \frac{1}{2 f_r \sqrt{\epsilon_{rEFF}} \sqrt{\mu_0 \epsilon_0}} - 2\Delta L, \tag{4.9}$$

where μ_0 is the permeability of free space and ϵ_0 is the permittivity of free space. This length, L, is 5.4 mm.

Therefore the dimensions of the patch are $W = 7.8$ mm and the length is $L = 5.4$ mm. The antenna was constructed with slightly larger dimensions than these (+0.5 mm) and then trimmed down to the correct size using a scalpel. The assembled patch antenna with SMA cable and mounting arm is shown in Figure 4.15.

The patch antenna was tested to identify the resonant frequency using a VNA. The results for the S11 parameter, which shows the reflected power applied to the patch antenna, are given in Figure 4.16. The antenna has a resonant frequency of 12.71 GHz. At the desired frequency of 12.64 GHz the reflected power is -16.1 dB which is less than 2.5 % of the input signal power.

The polarisation of the microwave radiation is important for driving the magnetic dipole transitions between the $^2S_{1/2}$, F = 0 $\leftrightarrow$ $^2S_{1/2}$, F = 1 states of the $^{171}Yb^+$ ion. The

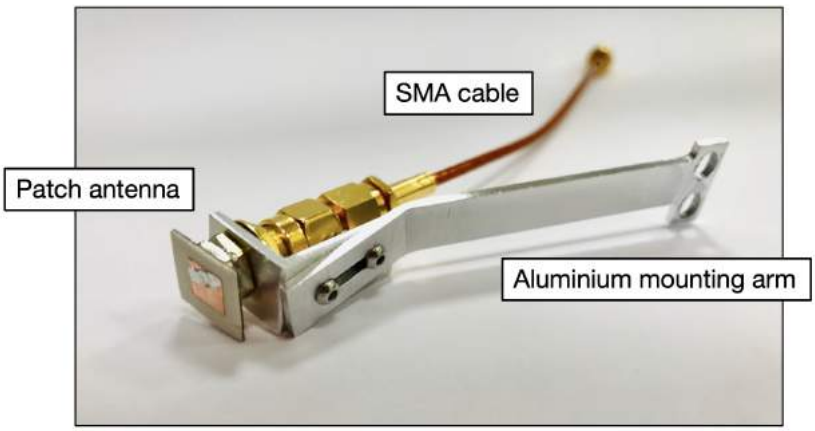

Figure 4.15: The microwave patch antenna with mounting arm and SMA cable assembly prior to installation into the vacuum chamber of the experiment.

transitions and polarisations are shown in Figure 4.17. A π-transition is a linearly polarised transition with no change of angular momentum between states. The $\sigma^{\pm}$-transitions are circular polarised transitions and they have an angular momentum change of ± 1. It is necessary to address all of the states simultaneously during Doppler cooling of $^{171}Yb^{+}$ to continuously excite fluorescence and prevent the ion becoming trapped in a dark state. The orientations of the polarisations required to excite transitions are determined by the direction of the applied magnetic field, $\vec{B_z}$, generated by the on-chip coils which causes Zeeman splitting and sets the quantisation axis. The field, $\vec{B_z}$, points along the axial direction of our traps. To drive the π-transition requires that there is a magnetic component $\vec{B}$ of the linearly polarised microwave field parallel to the quantization axis. To drive the $\sigma^{\pm}$-transitions requires a magnetic component of the microwave field to be perpendicular to the quantization axis. To drive both π- and $\sigma^{\pm}$-transitions equally requires the magnetic component of the microwave field to be oriented at a 45° angle to the quantization axis. The position of the patch antenna within the vacuum chamber was constrained by the large recessed imaging window above the chips, and so it was mounted in the orientation shown in Figure 4.18. The wavevector, $\vec{k}$, of the antenna points towards the trap surfaces at an angle of 10° from horizontal, and makes an angle with the trap axis of close to 60°. This orientation provides magnetic components of the microwave field both parallel and perpendicular to the quantization axis set by $\vec{B_z}$ as required to drive all transitions.

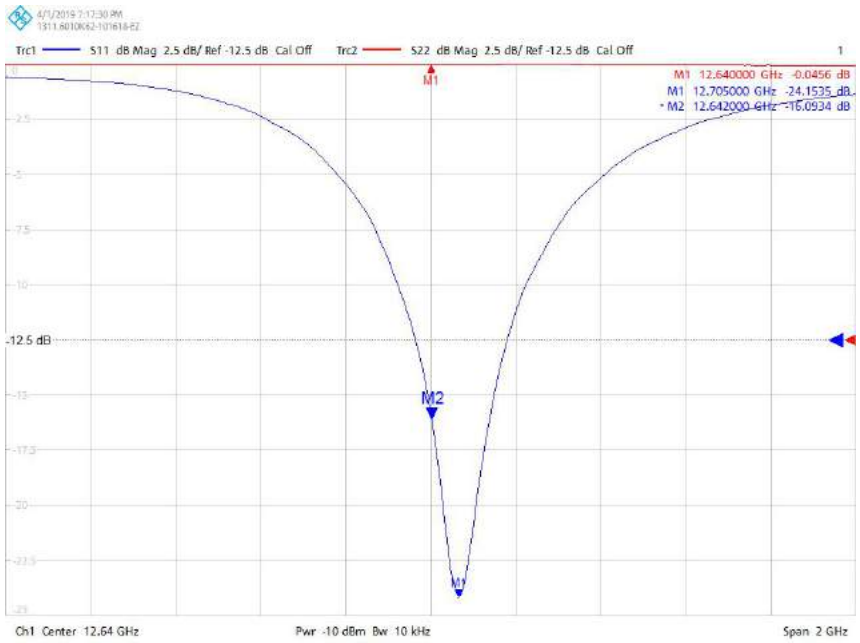

Figure 4.16: Testing of the resonant frequency of the patch antenna using a vector network analyser. The resonant frequency was 12.71 GHz. At the required frequency of 12.64 GHz the reflected signal was measured to be -16.1 dB of the input signal.

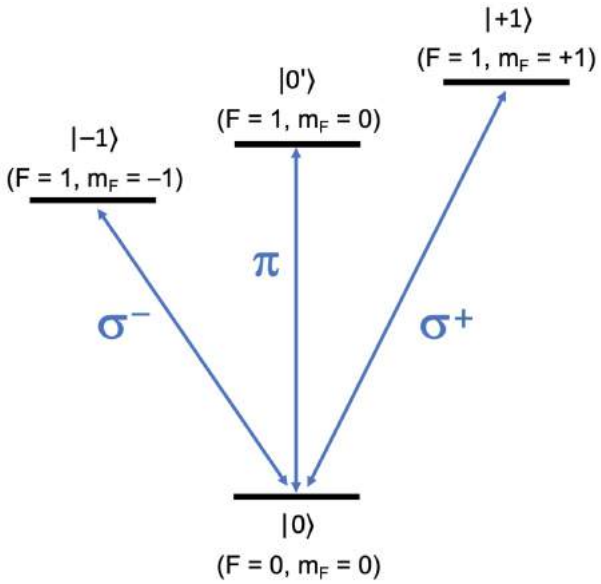

Figure 4.17: Magnetic dipole transitions in the $^2S_{1/2}$ hyperfine manifold of $^{171}Yb^+$. The polarisations shown are required to drive transitions between the states.

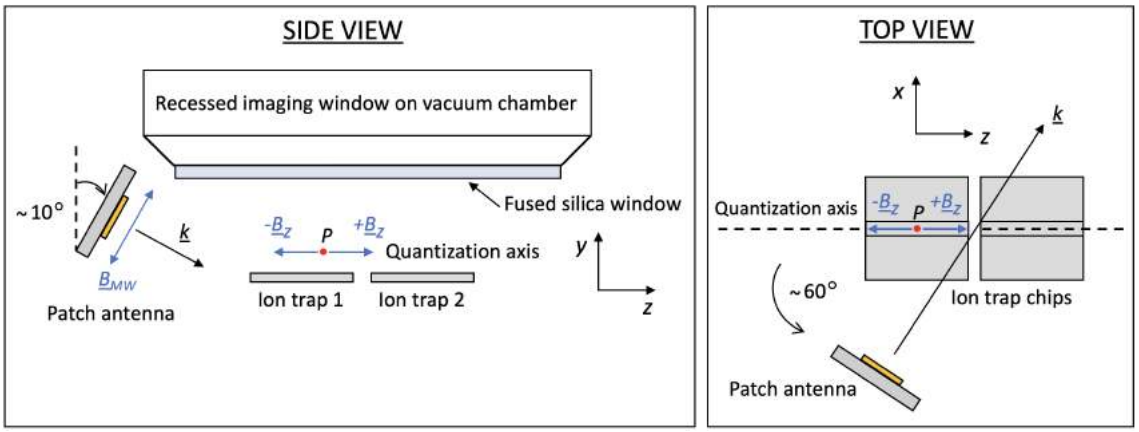

Figure 4.18: The patch antenna is mounted within the vacuum chamber as close to the ion position as possible to maximise the field strength at the ion position. The angle of the antenna is set to ensure that there is a component of the microwave B-field both parallel and perpendicular to the quantisation axis. The final angle achieved after mounting the antenna was approximately 10° to the vertical and 60° horizontally as shown.

4.10 Scalable ion-trap cooling system

Operating an ion trap quantum processor at low temperature is advantageous for device performance because it reduces the motional heating rate of trapped ions, leading to longer coherence times and improved gate fidelities. For the large-scale architecture proposed by Lekitsch *et al* [34], thermal management is a key engineering consideration just like for a conventional many-processor computing installation today. Since the large-scale ion-trap architecture calls for on-chip electronic components such as DACs, photon detectors and current carrying wires, cooling of the chips is an effective method of reducing power dissipation by improving the thermal conductivity of the silicon substrate, potentially by an order of magnitude [34]. One of the challenges is to develop a scalable method of cooling that is compatible with the modular approach of the large-scale computing architecture. For our two-module demonstrator device, we have developed a cost-effective, cold helium gas circulation system with a scalable architecture, and shown how it is extensible by connecting it to four independent ion-trap experiments in our laboratory. In this section I will describe the experimental considerations and power requirements that we needed to fulfil, followed by the selection of a suitable type of cryostat, and then a description of the completed set-up, including results achieved in cooling tests with the system.

4.10.1 Advantages of low temperature operation

For the two-module ion-trap processor, the purpose of cooling is threefold: (1) to achieve a low ion heating rate, (2) to assist in reducing the background pressure of the vacuum system by cryopumping (see explanation below), and (3) to reduce the power dissipation of the CCWs in the magnetic gradient chips. To decide upon an appropriate operating temperature for the cooling system, consideration was given to the potential benefits for each of these three criteria.

Ion heating rate

Strong suppression of electric field noise at low temperatures [82] leads to a reduction in anomalous heating of trapped ions that follows a power law [89] of the form:

$$\dot{\bar{n}}(T) = \dot{\bar{n}}_0(T)\left[1 + \left(\frac{T}{T_0}\right)^{\beta}\right], \tag{4.10}$$

where $\dot{\bar{n}}$ is the heating rate (mean number of motional quanta, n, per second), T is temperature, T_0 is the activation temperature (typically between 10 K and 75 K) and β is the

dimensionless exponent.

The power law means that most of the suppression effect occurs above the activation temperature, and an order of magnitude reduction in the heating rate can be achieved by cooling from room temperature to around 70 K [89].

Cryopumping

One of the benefits of operating an ion-trap at cryogenic temperatures, where gases such as nitrogen, oxygen, hydrogen and helium begin to condense, is the ability to perform cryopumping. Cryopumping is the process of using cold surfaces within a vacuum chamber to reduce the background pressure by condensing gases. This is normally achieved using the low temperature metal surface of a radiation shield that surrounds the ion trap [123].

The minimum partial pressure that can be achieved by condensation in this manner is the saturation pressure of the gas species at the temperature of the cold surface. Figure 4.19 shows the saturation pressures for a variety of atmospheric gases. Water and hydrocarbon species undergo a direct phase transition from gaseous to solid at ∼100 K. Most other gases present in atmosphere condense at ∼20 K. Hydrogen and neon condense at ∼ 4 K, and at this temperature it is no longer necessary to perform baking treatment to achieve UHV pressures.

Cooling of current carrying wires

The magnetic gradient chips (described in Section 3.3.2) use microfabricated, high current wires embedded in a silicon substrate to generate large magnetic field gradients to support enhanced coupling of the internal and external degrees of freedom for ion qubits during logic gate operations [74, 73]. Higher gate fidelities can be achieved by increasing the magnetic field gradient, which requires large currents to be applied to the CCWs. The maximum sustainable current can be increased by cooling the chips to reduce the power dissipated within the CCWs. Cooling from room temperature to 70 K decreases the electrical resistivity of copper by a factor of ∼10 [125], and increases the thermal conductivity of silicon by a factor of ∼11 [126], leading to a large reduction in I^2R heating and a corresponding increase in the maximum current that the chip can sustain [34].

4.10.2 Temperature and power requirements

Based on the criteria outlined above, the required operating temperature for the cooling system was set at 70 K. Given this temperature, it was then necessary to consider the

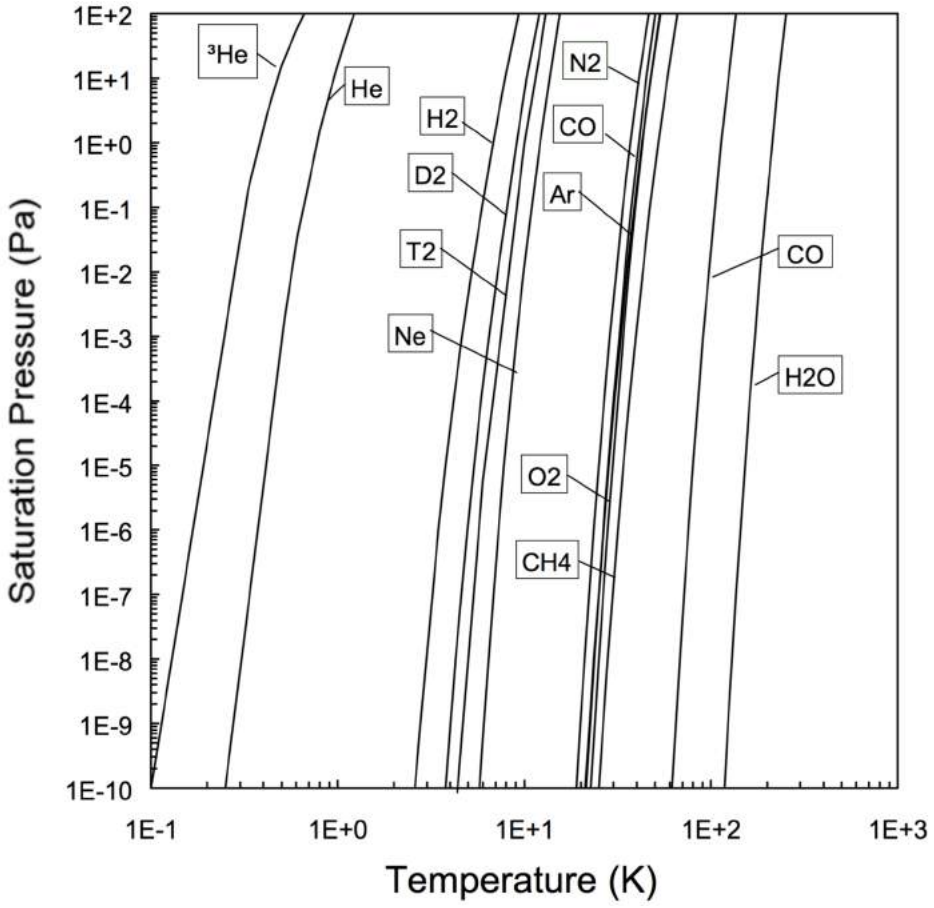

Figure 4.19: Saturation curves for gases present in atmosphere. $1\,\mathrm{Pa} = 10^{-2}\,\mathrm{mbar}$ [124].

cooling power requirements that the system would need to fulfil, which is equivalent to the heat load on the system due to our four ion-trap experiments. Contributions to the heat load can be categorised into either passive or active loads. Passive loads include feedthrough connections into the chamber, that provide a conductive path from the room temperature environment of the laboratory to the low temperature ion-traps. Active loads are dissipative processes that generate heat, which includes the high current CCWs of the magnetic gradient chips, and the application of the RF trapping voltages to the surface electrodes of the chips. Considering all of the elements that may contribute to the heat load in each of the four ion-trapping experiments led to the allocation of a power budget for each experiment. The power budget for the two-module ion-trap processor was set at 75 W (allowing headroom for future expansion of the experiment), and the power budgets for each of the other three experiments in the laboratory were set at 12 W. Therefore, the combined cooling power requirement for all four experiments was 111 W.

4.10.3 Selection of cryostat type

Consideration was given to a variety of cryostat technologies that were capable of meeting the operating temperature and power requirements of the experiments. These included liquid nitrogen systems, either as a bath cryostat or a closed-flow system; helium gas circulation systems with forced flow using centrifugal fans; and dry cryostats such as Gifford McMahon (GM) or pulse tube (PT) cryocoolers, connected to the ion-trap chips either by a low vibration interface or as part of a cryogenic flow system. A cryocooler is a device that uses repeated cycling of a gas expansion chamber to cool a metal probe, which may be placed into contact with an application to provide cooling. The probe and the chamber are housed within a 'cold head' which is fed by gas from a compressor.

Our requirement to develop a scalable cooling system, that could be extended to many quantum processor modules within separate experiments, was well aligned to the properties of a closed-cycle flowing cryogen system. Although liquid nitrogen has an advantageous specific heat capacity compared with helium gas, the boiling of nitrogen above 77 K (unless appropriately pressurised) can lead to undesired vibrations, and in a closed flow system additional engineering is required to prevent nitrogen from existing in both the liquid and gas phases. Helium gas on the other hand is easy to use, has a wide range of temperatures for which it remains stable, and its heat capacity can be increased by pressurising, which is easy to achieve using a regulator attached to a high pressure gas bottle. By using helium as a versatile cryogen, a closed-cycle circulation system can be constructed in which the cooling power is provided by one or more cryocoolers with their cold heads in contact with the flowing gas. This is an attractive solution because cryocooler technology is well established and there are many models on the market to meet a wide range of temperature and power requirements. Two models of cryocooler were identified as being compatible with our application, the Sumitomo CH110 and the Cryomech AL330, and the performance of each device is shown in the capacity curves of Figure 4.20 and Figure 4.21, which are graphs of temperature for a given heat load. For an operating temperature of 70 K (and with a 50 Hz power supply), the Sumitomo CH110 provides a cooling power of 180 W, while the Cryomech AL330 provides 190 W.

Therefore the decision was made to develop a helium circulation system, since this approach can be scaled up in size and power based on the simple principle of adding cryocooler units to increase the capacity of the system.

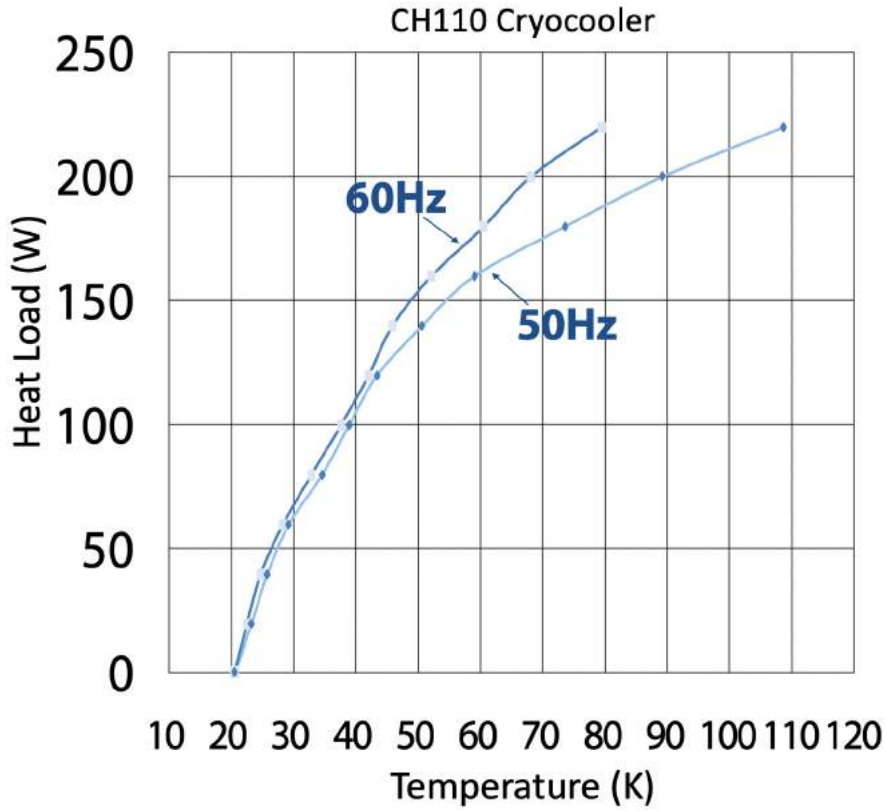

Figure 4.20: Capacity curve for the Sumitomo CH110 cryocooler, showing the temperature achieved by the cold head versus heat load from the application.

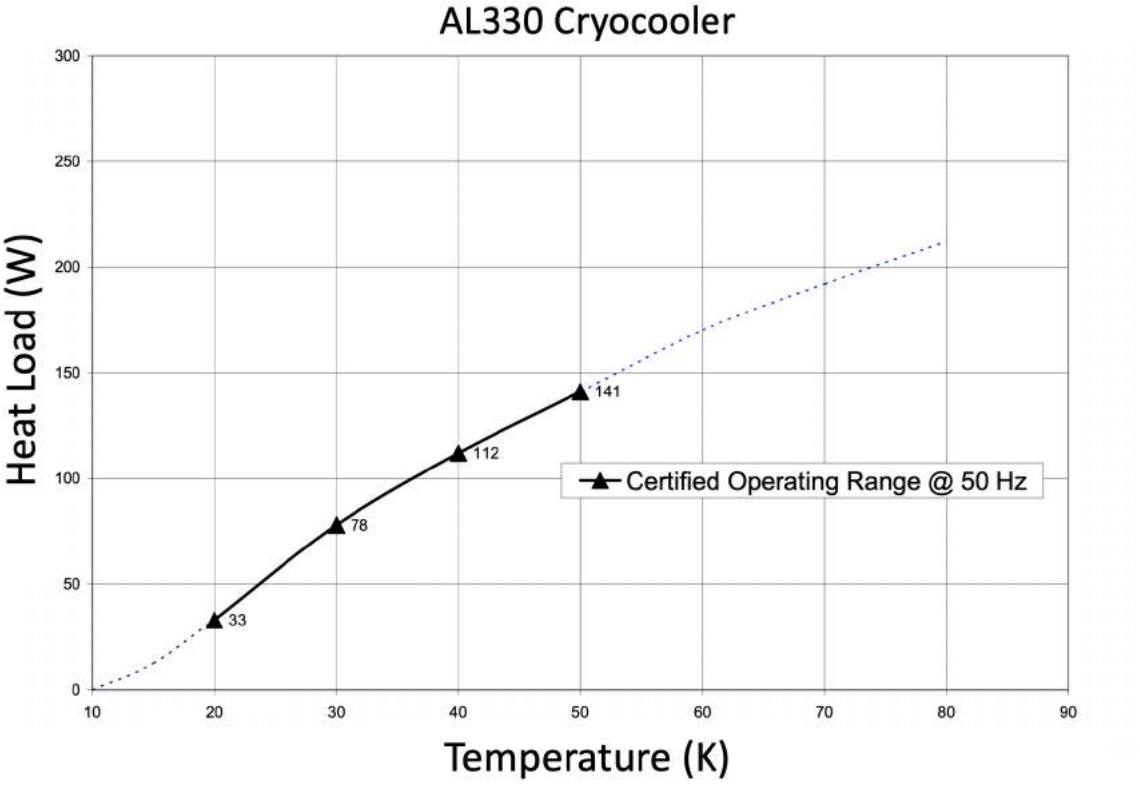

Figure 4.21: Capacity curve for the Cryomech AL330 cryocooler, showing cold head temperature versus heat load.

4.10.4 Design considerations for a helium circulation system

Figure 4.22 is a schematic diagram of the basic elements of the cooling system. The system is made up of (1) a main cryostat chamber, which includes the cryocooler cold heads, (2) the helium transfer lines, which in practice are vacuum insulated, and (3) the ion trap experiments, which include the processor modules that require cooling. Within each module, the helium flows through a heat sink that interfaces with the microchips. To circulate the helium cryogen around a closed system, a pressure differential must be created within the gas, and this can be achieved using a cryogenic centrifugal fan (also referred to as a cryofan). The total pressure head, ΔP supplied by the cryofan must balance the local pressure drops throughout the system,

$$\Delta P = \Delta P_{cryostat} + \Delta P_{lines} + \Delta P_{modules}, \tag{4.11}$$

where the pressure drop is caused by dissipative friction forces between the flowing gas and the surfaces throughout the circuit.

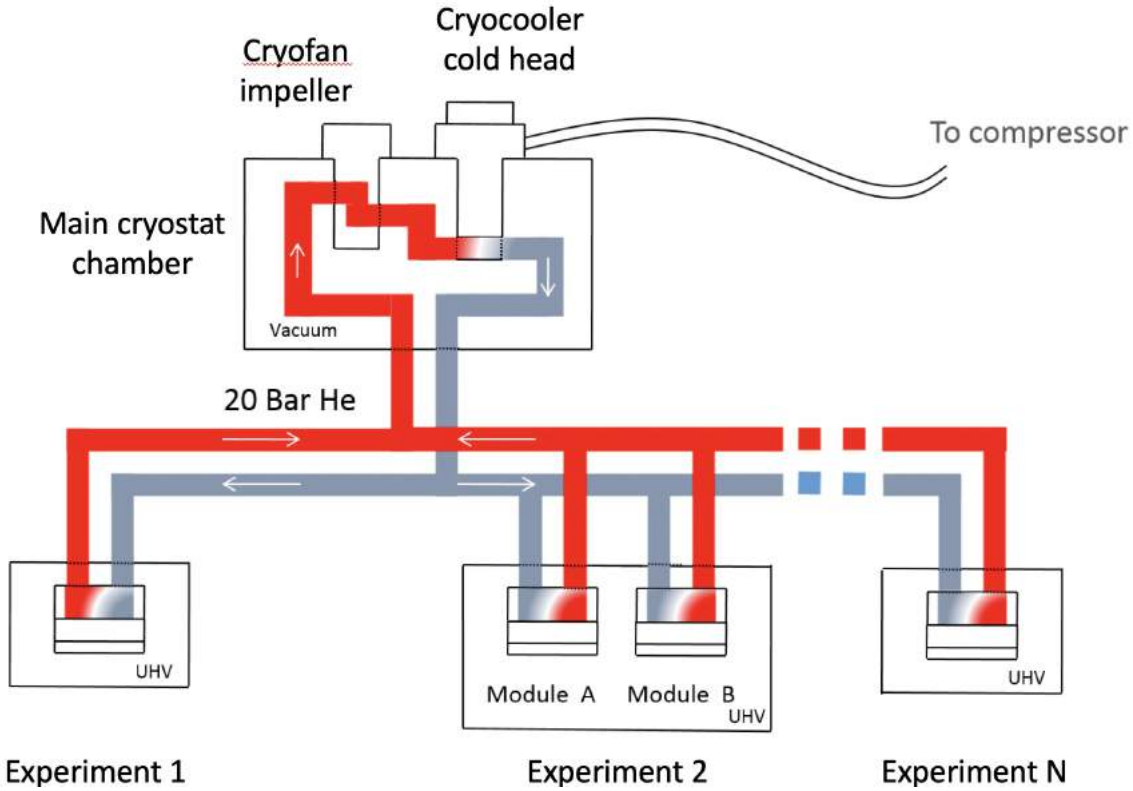

Figure 4.22: Schematic diagram of the closed-cycle helium gas circulation system architecture. The system can be scaled-up to fulfil higher cooling power requirements by installing additional cryocooler cold heads into the main cryostat chamber.

The choice of cryofan is driven by the required volume flow rate for our application, which then determines the temperature gradient between the cold helium inlet and warm

helium outlet at a module. We wish to keep this temperature difference to within a few K, to ensure a low operating temperature for the ion trap modules and a stable pressure within the system.

To understand the total cooling power requirements, the total heat load on the system must be considered. This is the sum of heat loads from all sources, including the modules, the mechanical power delivered to the helium gas from the cryofan, the frictional heating within the lines (referred to as labour), and heat ingress from room temperature into the helium lines and cryostat chamber. Therefore the total power, $\dot{Q}$, to balance the heat loads is given by,

$$\dot{Q} = \Sigma\dot{Q}_{modules} + \Sigma\dot{Q}_{lines} + \dot{Q}_{fan} + \dot{Q}_{cryostat} + \dot{Q}_{labour}. \tag{4.12}$$

Therefore the system was required to provide an overhead of cooling power above the specified 111 W that we had estimated for the four ion trapping experiments within the laboratory.

4.10.5 Cooling system laboratory set-up

To provide the necessary cooling capacity to operate all four experiments as required, both the Sumitomo CH110 and Cryomech AL330 cryocoolers were installed into the main cryostat chamber of the cooling system, which is shown in Figure 4.23. This configuration delivers a total cooling power of ∼350 W at 70 K. The large stainless steel enclosure is evacuated to $<1\text{x}10^{-6}$ mbar to provide thermal insulation for the cryocooler cold heads, the cryofan unit and the helium gas manifold that connects the devices inside the chamber. Each of the cryocooler cold heads are operated by a water cooled compressor, the Sumitomo F-70H and Cryomech CP9170, respectively, and these connect via 20 m helium transfer lines. The cryofan is a Cryozone Noordenwind model that delivers a volume flow rate of up to 2.75 m^3/hr using a 31 mm impeller that rotates at a maximum speed of 21,000 revolutions per minute. Cold helium gas pressurised to 20 bar is circulated through four parallel networks of supply and return transfer lines made of 10 mm diameter stainless steel pipework enclosed within vacuum insulated stainless steel tubing of 100 mm diameter. These extend across the laboratory to connect to the four independent ion trap experiments in parallel, covering a distance of approximately 30 m.

The cryostat chamber connects to the helium transfer lines using 1/2 inch VCR connections that can be accessed via a sealable top hat, as shown in Figure 4.23. At the location of each ion trap experiment within the laboratory, the helium supply and return lines split and terminate in flexible vacuum jacket lines with 1/4 inch transfer lines at their centre,

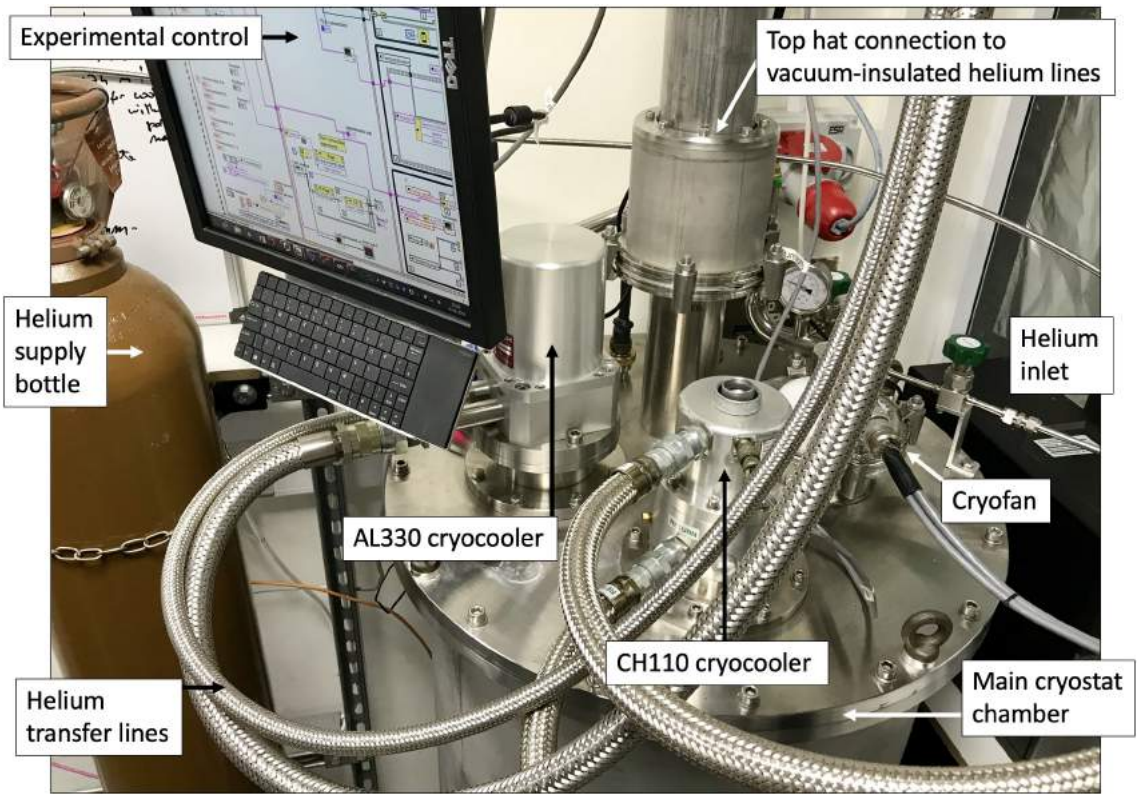

Figure 4.23: Photograph of the main cryostat chamber.

ending in a VCR connector. The surrounding jacket line has a KF flange which connects to the vacuum chamber using a top-hat sliding seal to allow access to the VCR connector on the helium line.

The vacuum insulation surrounding the main helium transfer lines is kept separate from the flexible jacket lines by a vacuum barrier, and so the flexible lines need to be pumped down independently from the main cryostat to achieve vacuum insulation. A vacuum barrier is a thin wall of stainless steel that connects between the inner helium line and the surrounding jacket, to form a closed seal. The length of the barrier is extended to provide a high thermal resistance path between the internal low-temperature environment and the external jacket at room temperature. This prevents ice from forming on the cold surface. Vacuum barriers are also installed on the ion trap vacuum chambers, to allow independent disconnection of the flexible lines, without losing ultrahigh vacuum conditions in the ion trap experiments. Combined with the fact that each experiment is supplied by a parallel helium circuit for cooling, this ensures that every ion trap experiment is completely independent of the others.

4.10.6 Cooling system performance

The performance of the cooling system was tested by Raphaël Lebrun-Gallagher. To achieve this, the cooling system was connecting simultaneously to two vacuum systems at separate locations within the laboratory. Each system was configured with an in-vacuum heat sink, and instrumented using PT100 temperature sensors and heaters to vary the thermal load.

Initial tests conducted under no active heat load showed that the first heat sink reached a steady state base temperature of 31.5 K after 3 hours of cooling. The second heat sink reached a temperature of 36 K in 2 hours.

Active heat load tests were then conducted using the heaters attached to the two heat sinks, and by monitoring the temperature across the system using thermometry fitted within the cryostat chamber, the results were used to make an estimate of the total volume flow rate and mass flow rate, based on a model of the system constructed by Raphaël Lebrun-Gallagher. This led to estimates of the heat load contributions from each of the components of the system, including the main cryostat chamber, the helium lines, the cryofan, and by inference the labour in the system arising from friction forces. The results then permitted a projection to be made of the temperature that the cooling system was capable of achieving when operating four experiments running in parallel. For a heat load contribution of 111 W from the four experiments, the system was projected to operate at a minimum temperature of $\sim$70 K.

Chapter 5

Assembly and characterisation of the demonstrator using first generation ion trap chips

5.1 Introduction

This chapter describes the steps taken to prepare the two-module demonstrator device for ion trapping experiments using the first generation microchips (described in Chapter 3). The hardware elements of the experiment have been covered in Chapters 3 and 4. I begin by describing the testing and processing treatment required to prepare the microchips for trapping. After this, I describe how the chips were installed into the experiment by die-bonding onto copper mounting blocks using a UHV-compatible epoxy, and then electrically connected to the in-vacuum circuit boards by wire-bonding. After fully assembling all in-vacuum components of the experiment and closing the vacuum chamber, the two-module system was characterised. This included an investigation of the electrical behaviour of the chips during module alignment, and testing of the travel and precision of the positioning system, including thermal drifts resulting from temperature changes during cooling of the system. The chapter ends with a discussion of the implications for achieving module-to-module connectivity based on the ion trap alignment achieved in the experiment.

5.2 Preparation of microchips for ion trapping

Batches of chips were received on 6-inch silicon wafers, and these were processed by Martin Siegele into individual dies using a diamond saw and dry etching with XeF_2 gas to produce

a suitable edge for module alignment. The next step was a visual inspection of each chip using an optical microscope, to check for defects or damage to the electrodes and thus identify the most perfect candidates to use in the experiment. After this, electrical testing of the chips was carried out by carefully positioning needle probes onto each electrode to check for short circuit connections. Shorts can occur between neighbouring electrodes or between an electrode and the underlying ground plane layer of the chip. All of the above testing was conducted by Martin Siegele. The two chips selected for installation into the experiment were a SET150-02 linear design labelled #9 and an EX150 X-junction design labelled #6. Figures 5.4 and 5.5 show the chips in their final conditions after all of the upcoming steps described in this section were completed.

5.2.1 Cleaning using solvent wash

The microchips are washed with acetone and isopropyl alcohol (IPA) before being installed into the experiment. The purpose of the cleaning is threefold: (1) it removes any oil or surface contaminants that would prevent the attainment of ultra-high vacuum (UHV) pressures in the experiment because of high outgassing rates, (2) it ensures that the electrode surfaces are clean and free of any residues that would prevent trapping or contribute to motional heating of trapped ions, and (3) it provides a clean surface for epoxy bonding chips to the copper mounting blocks and for wire bonding. The cleaning procedure is as follows:

1. Ultrasonic bath at maximum vibration setting in acetone bath for 5 mins
2. Rinse with IPA
3. Ultrasonic bath on maximum vibration setting in IPA for 5 mins
4. Blow dry with inert gas (argon or nitrogen)

The ultrasonic bath provides sufficient agitation of the solvent against the gold surface of the chips to remove contaminants, and to avoid risk of damage the duration was limited to 5 minutes. The ultrasonic bath was also found to remove some electrical shorts. This suggests that small fragments of conductive material trapped within the trenches between electrodes was being removed.

5.2.2 Die-bonding using single component epoxy

After cleaning, the chips were attached onto oxygen free high thermal conductivity (OFHC) copper mounting blocks using single part cryogenic epoxy EPO-TEK H67-MP. The epoxy

is easy to use, suitable for UHV and has a thermal conductivity of 0.5 W/mK, which is higher than other comparable epoxy products. The recommended cure conditions are 150 °C for 1 hour, and once cured the epoxy is capable of resisting baking temperatures up to 260 °C. The copper blocks provide a thermal connection to the heat sink of the cooling system, which is positioned directly underneath the mounting block of the X-junction chip. A flexible copper braid joins the heat sink to the mounting block of the linear chip, allowing for free movement of the chip by the positioning system (described in Chapter 3).

A disadvantage of the epoxy is the mismatched thermal expansion coefficient compared to that of the silicon substrate of the chip. There is a risk that temperature cycling due to baking treatment followed by cryogenic cooling and warming to room temperature may cause stresses at the interface between the silicon and epoxy owing to different rates of expansion and contraction that can damage the chip. This was investigated during characterisation of the system and the results are discussed later in this chapter, although in short no damage was observed due to this effect.

Material	**Thermal conductivity** $Wm^{-1}K^{-1}$	**Thermal expansion coefficient** $^{\circ}C^{-1}$
Silicon	150	$5\text{x}10^{-6}$
OFHC Copper	390	$16\text{x}10^{-6}$
EPO-TEK H67-MP	0.5	$16\text{x}10^{-6}$
Indium	86	$25\text{x}10^{-6}$

Table 5.1: Properties of thermal interface materials used in the experiment

For the first generation ion trap chips which do not feature magnetic gradient coils, the power dissipation was expected to remain below ∼10 W (considerations of power dissipation were discussed as part of the cooling system description in Chapter 4), and in this regime the epoxy is a suitable choice of thermal interface material. However, for experiments using second generation chips, the increased power dissipation of the on-chip coils requires a material with higher thermal conductivity to minimise the temperature gradient that arises between the hot chip surface and the cold heat sink. A potential replacement material is indium metal foil, which has improved thermal conductivity compared to the epoxy. Table 5.1 compares the properties of the materials discussed in this section. Indium has already been used successfully as a thermal interface material in a previous cryogenic experiment within the group [123]. Although indium has a thermal expansion coefficient that is also mismatched to both silicon and OFHC copper, its malleability allows it to deform at the interface between materials. However, recent attempts at using indium have encountered poor adhesion between the chip and the copper mounting block, so we chose to use epoxy for

the experiments with first generation chips. To prepare the mounting blocks for bonding, the top surface of each block was polished to a mirror finish using sand paper of grit 1200 to remove oxide and create a smooth bonding surface, and the blocks were cleaned by sonication in an IPA bath for 20 minutes, then blow dried using argon gas. Each chip was bonded to its mount using a small amount of epoxy that was added to the polished top surface of each block. To avoid overspill near the alignable edge of the chip, the epoxy was applied in a V shape with the angle of the V at the centre of the block and the two arms of the V extending towards the corners of the block at the non-alignable edge. The chip was placed onto the surface and positioned with the edge overhanging the block by approximately 0.5 mm, to allow for unimpeded chip alignment when the two modules are brought together. To cure the epoxy, a die bonder[41] was used. The copper mount was secured in position by the vacuum base, and the placement head (with suction disabled) was lowered onto the chip to apply a downward force of 8 N. A small amount of excess epoxy was expelled and wiped away with a fibre-free cleanroom cloth. The magnified camera of the die-bonder was used to carefully check the position of the chip on the mounting block, to minimise any angle of rotation that could prevent the two chips from being closely aligned during the experiment. Heating was then applied to both the base and head with a ramp profile of 1 °C/s to reach a temperature of 150 °C, which was maintained for the 1 hour duration of the cure. After this, the heaters were turned off and the block-chip assembly was left to cool at room temperature. When cool, the chip was rinsed with IPA and dried using argon gas. The results of the epoxy bonding process are shown in Figure 5.1 for the linear chip and Figure 5.2 for the X-junction chip. After a visual inspection, the surface of the linear chip was found to have been contaminated with a black substance that formed a circular deposit approximately 500 µm in diameter. This was most likely the result of dust or dirt falling onto the chip when the die bonder head was lowered into position, which was then compressed and heated during the cure. Multiple rounds of sonication in acetone only partially removed the substance, and so it was carefully wiped from the surface using acetone and a cleanroom cloth. Figure 5.1 shows the area of contamination on the chip surface before and after cleaning. It is expected that the residue will contribute to significant motional heating heating of ions trapped on the chip.

The X-junction chip was bonded using the same procedure as for the linear chip, after cleaning the die-bonder head with IPA. A visual inspection after the bonding procedure showed no contamination on the surface.

[41] Finetech FINEPLACER pico ma

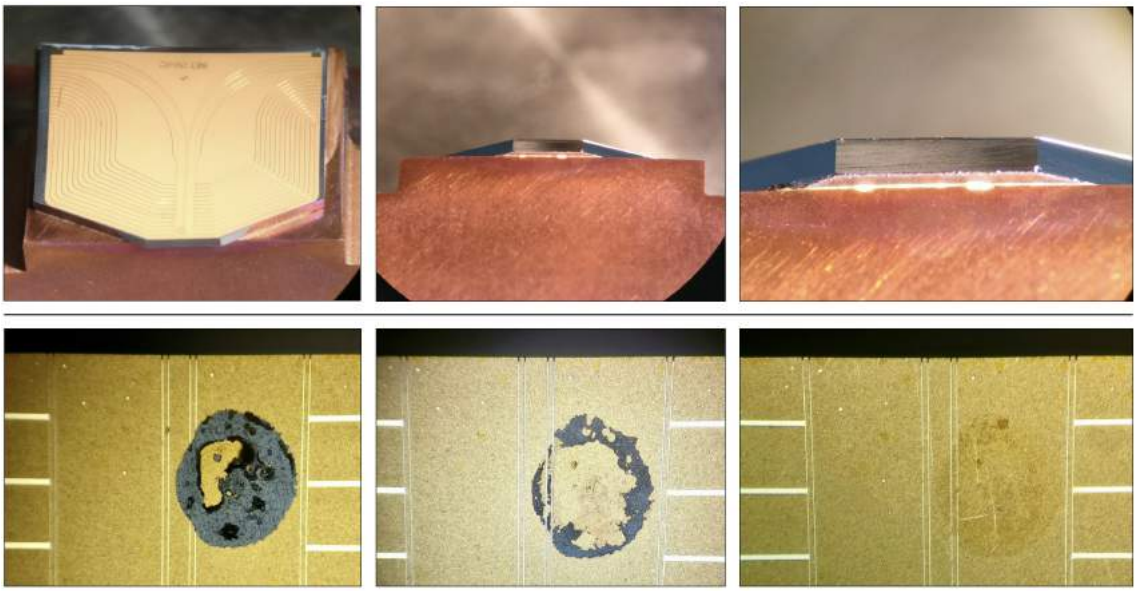

Figure 5.1: Linear chip after die-bonding to copper mounting block using EPO-TEK H67-MP epoxy. *Top row*: The chip is positioned so that it overhangs the edge of the block by approximately 0.5 mm to permit unimpeded alignment with the adjacent module. The small amount of epoxy overspill posed no problems in operating the chip. *Bottom row*: Contamination from a particle of dirt discovered on the chip surface after using the diebonder. Cleaning with acetone and IPA gradually removed most of the material but a surface residue remained (bottom right image).

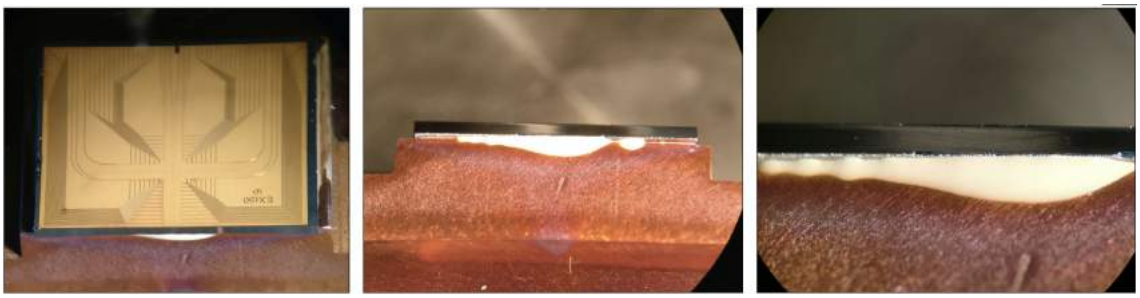

Figure 5.2: X-junction chip after being epoxied to its copper mounting block using the same process as for the linear chip. The angle of the first image (*left*) is somewhat deceptive, since the extent of the epoxy overspill is in fact smaller than the 0.5 mm distance by which the chip overhangs the copper block. The residual epoxy caused no problems in operating the chip or aligning the ion-trap modules together.

The next step in the preparation of microchips for ion-trapping was wire bonding the chips to their respective circuit boards.

5.2.3 Wire bonding

Wire bonding is a technique used to make electronic connections onto integrated circuit microchips by precisely welding thin conducting wires onto electrode surfaces using downward pressure from an applicator head, in combination with an ultrasonic pulse of vibrational

energy. The wire bonder used in the laboratory[42] is specifically a wedge bonder, which refers to the shape of the head of the applicator needle. The wire bonder was used to connect the gold surface electrodes of the linear and X-junction chips to the PCB traces using 25 µm aluminium wire. The wire-bonded chips are shown in Figure 5.3.

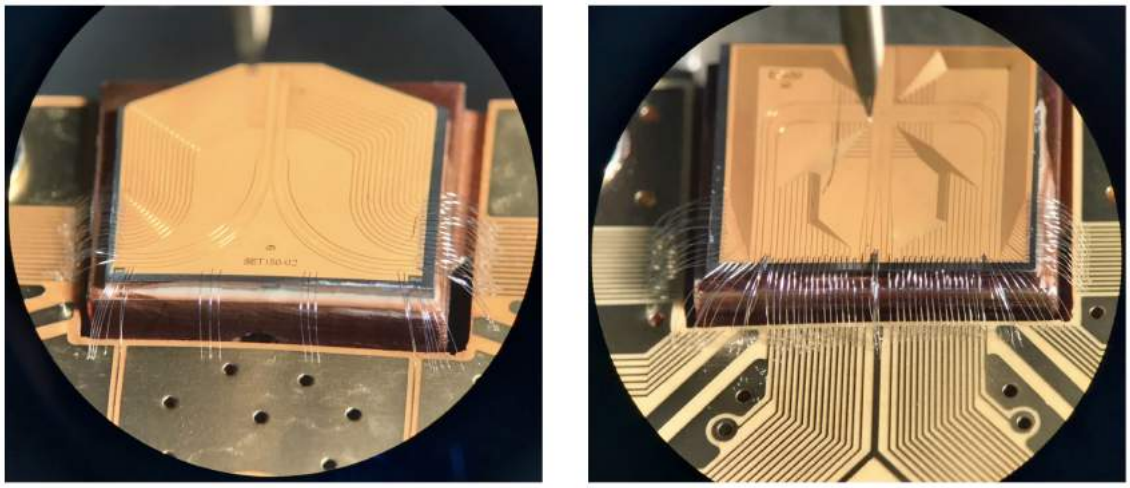

Figure 5.3: Linear (*left*) and X-junction (*right*) ion trap microchips after wire-bonding to their respective circuit boards. To increase the reliability and robustness of the electrical connections, at least two wire bonds were made for each DC electrode, and 5 wire bonds are used to connect to the RF electrodes.

To increase the reliability and robustness of the electrical connections, each of the DC electrodes were connected to the PCB using two wire bonds, where possible. The number of bonds achievable was limited by space constraints and by the angle of bonding. The RF electrodes were each connected by five bond wires, and this configuration has previously been shown to be capable of withstanding an RMS power $> 10\,\mathrm{W}$ in testing of chip breakdown voltages (described in Chapter 3).

After wire bonding, the chips were tested for shorted connections between electrodes using a multimeter applied to the PCB tracks. No shorts were detected on the linear chip. The X-junction was discovered to have a $2.2\,\Omega$ short from the left RF track (denoted RF2) to ground. This was successfully removed (at a later stage, once the chips were installed into the vacuum system) by applying a pulse of current at 2 A for $\sim 2\,\mathrm{s}$ using a precision regulated power supply[43]. After the pulse, measurements showed open circuit between the RF electrode and the ground plane, indicating that the conducting material causing the short had been burned away.

During handling of the X-junction chip under the applicator needle of the wire bonder, the surface was scratched across multiple electrodes on the left hand side of the chip. The

[42] 4523A digital wedge bonder
[43] PSD 30/3B High performance regulated power supply

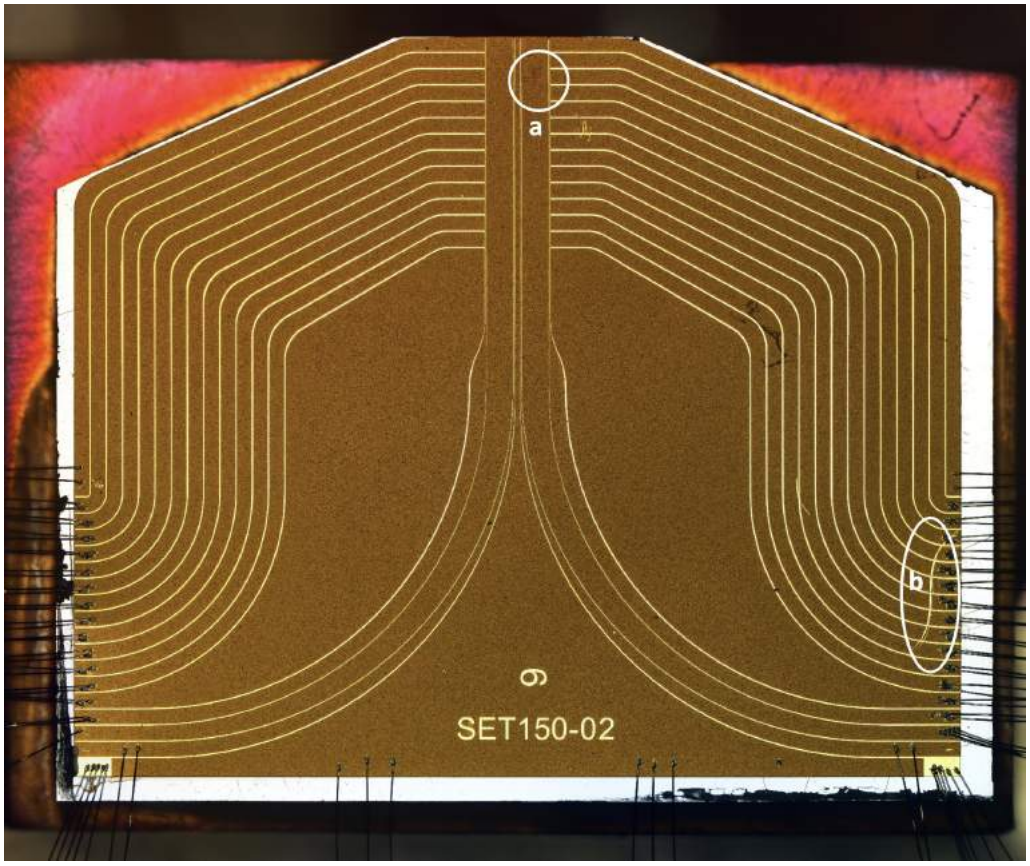

Figure 5.4: Linear surface trap microchip SET150-02 #9, prior to installation onto the piezo side of the experiment. Defects are circled and labelled in white. Location (a) on the right hand RF track is the site of contamination from the die-bonder, while location (b) is a superficial scratch across eight DC electrodes that did not result in electrical shorts.

scratch is visible in Figure 5.3 and Figure 5.5, and in the close detail photographs of Figure 5.6. A visual inspection using a microscope revealed that the scratch was not so deep as to penetrate through the gold surface layer into the underlying dielectric layer, but small amounts of gold flakes from the surface layer had been pushed into the trenches between electrodes, causing electrical shorts to occur. In an attempt to remove the gold flakes, the chip was sonicated for 4 minutes in IPA and then sprayed directly with IPA from a wash bottle. This was successful at clearing away gold from the trenches and removing the shorts, and so the chip was deemed suitable for use in the experiment.

A visual inspection of both the linear and X-junction chips revealed a number of other non-critical defects on their surfaces, prior to installation into the vacuum chamber. These defects are labelled in Figure 5.4 and Figure 5.6. The X-junction, in particular, exhibited a range of minor surface defects, including a breaking-up of the gold surface on DC electrodes #62 and #63, that appeared to result from faults during microfabrication, and these are shown in Figure 5.6. None of these defects were considered to be serious enough to prevent

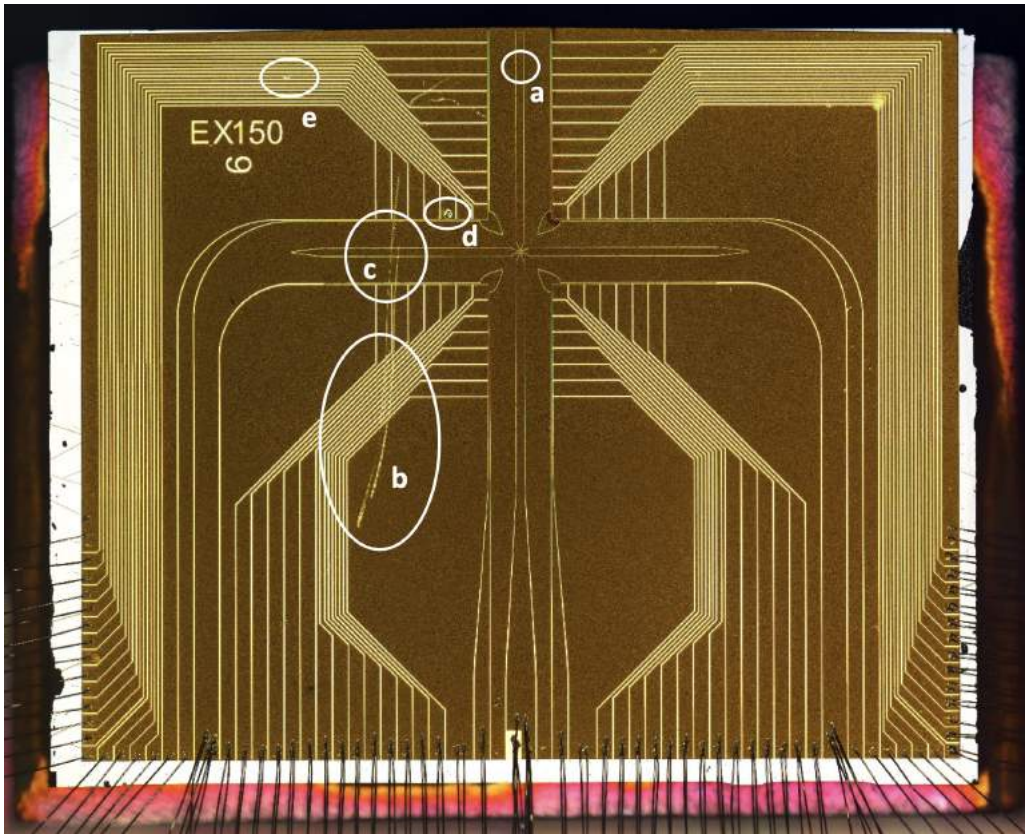

Figure 5.5: X-junction surface trap EX150 #6 after wire bonding. This chip was installed onto the heat sink side of the experiment. Defects are circled in white - see detailed descriptions in figure 5.6 below.

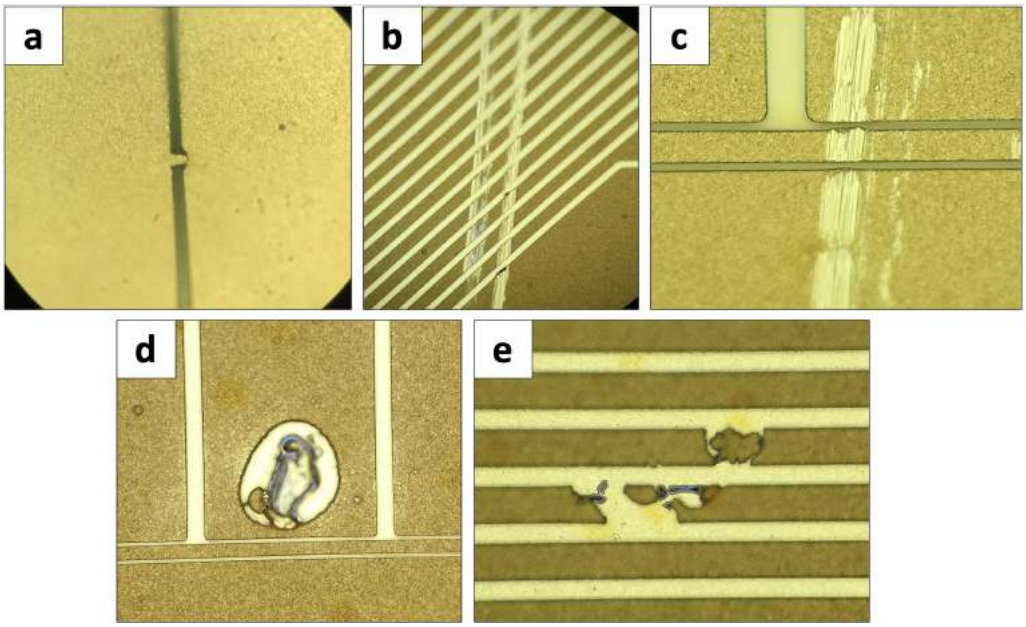

Figure 5.6: Site (a) is a bridge of gold between the right RF track and central ground electrode, estimated to be 3 to 5 µm wide. A long scratch crosses DC electrodes at site (b) and RF electrodes at site (c) but causes no shorts. Location (d) is a microfabrication defect on DC #59. At site (e) the gold surface is broken on DCs #62 and #63.

the chips from being used for trapping and shuttling ions.

5.2.4 Installation and manual pre-alignment of microchips

The next step in the assembly of the experiment was to install the two PCBs (including copper mounting blocks and wire-bonded chips) onto the titanium support chassis that was used to mount the ion trap modules within the vacuum chamber. On completion, the fully assembled chassis was installed into the chamber using groove-grabber clamps to attach onto the CF160 flange of the Kimball Physics spherical square. Feedthrough connections were then made for the RF and DC electronics, thermometry, and cryogenic cooling. Connection of the helium supply lines to the heat sink required two VCR connectors to be tightened by applying a substantial torque, which resulted in movement of the two ion trap modules towards each other. Inspection of the chips afterwards showed that damage had occurred to the electrodes at the edge of the X-junction (DC electrodes #1 and #72, and the end of the RF electrodes at the edge), confirming that the chips made contact along their alignable edges. The linear chip was not damaged. The movement of the support structure when the torque was applied also caused the modules to become misaligned, as shown in Figure 5.7.

The collision lifted the electrodes at the edge of the X-junction chip to form a ramp with a height of $\sim 100\,\mu\text{m}$. The lack of damage on the linear chip can be explained by a height difference during the collision, so that the electrodes of the X-junction would have impacted the silicon layers that lie underneath the electrodes of the linear chip.

Since the chips were misaligned as a result of the movement of the support structure, an attempt was made to adjust the position of the linear chip using the manual pre-alignment stage (described in Chapter 3). Unfortunately the stage had become seized and could not be released. This was due to friction between a retaining lug and a guide track, which caused undesired rotation during translation. A new improved design was developed by Raphaël Lebrun-Gallagher but the part was not ready in time for the experiments described in this chapter.

The lateral misalignment of the two chips was estimated to be $\sim 400\,\mu\text{m}$, which is around half of the full width across the RF tracks, as can be seen in Figure 5.7. This starting misalignment later led to challenges when attempting to align the chips, although the problem was overcome as described later in this chapter.

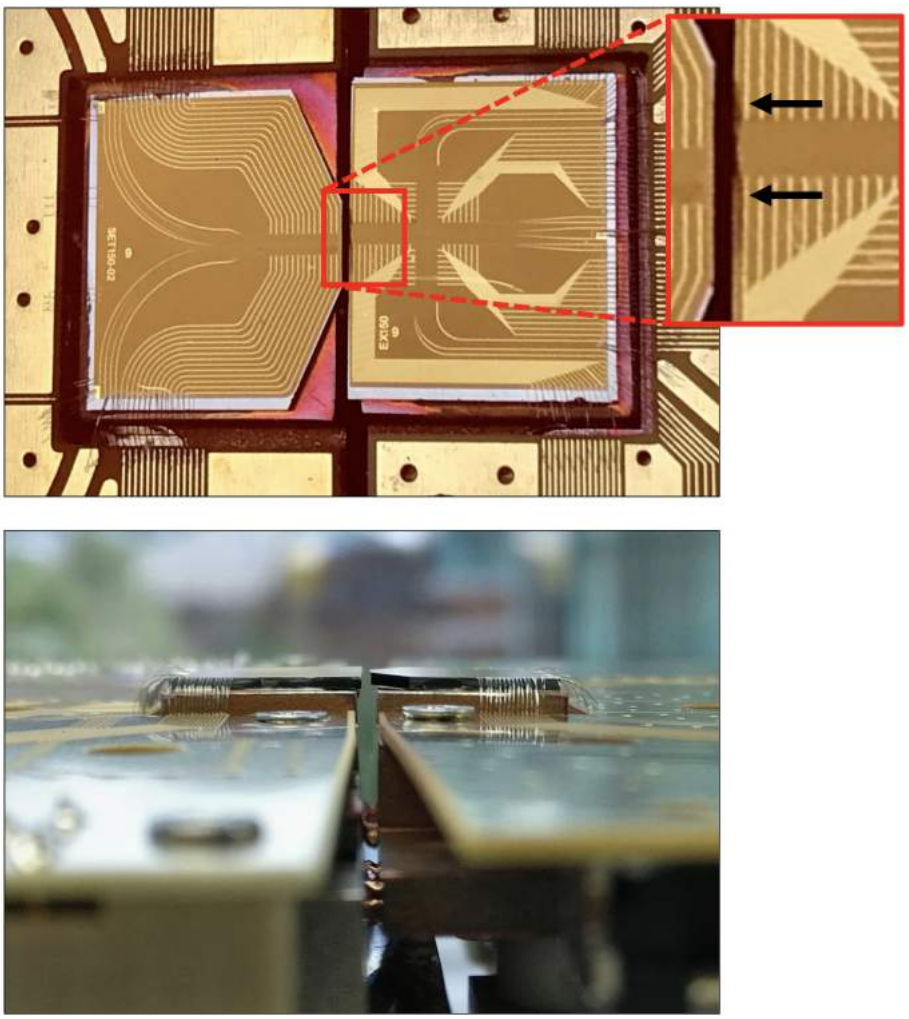

Figure 5.7: The top image shows the separation and lateral misalignment of the microchips. The inset zooms into the damage at the edge of the X-junction, where the electrodes had lifted to form a ramp. The bottom image is a side view of the gap between the two modules, showing the engineered overhang of each chip on its mounting block, and the copper braid between the two modules (just visible).

Status of the microchips after installation into the experiment

A summary of the electrical shorts and significant defects for each ion-trap chip, following installation into the experiment, is given below.

Linear chip:

- Localised surface contamination due to plastic residue on right-hand RF track close to alignable edge (after die-bonding).

X-junction chip:

- RF short: between right RF track and central ground electrode.
- DC #8: 112 Ω to ground.

- DC #59: 13.3 kΩ to ground.
- DCs #62 and #63 disconnected.
- Electrodes along edge of chip lifted to form ≈ 100 µm ramp.

At this stage the vacuum chamber was closed-up and evacuated using a scroll pump and turbomolecular pump (TMP). The vacuum chamber was leak tested by dispensing helium close to the ports and monitoring the output from a residual gas analyser (RGA) connected to the input port of the TMP, to monitor any helium ingress into the chamber. On completion, the chamber was moved from the cleanroom in which the experiment was assembled, and set-up on the optical table in the laboratory in preparation for imaging, measurement and module alignment tests.

5.3 Imaging system testing and calibration

This section describes the tests of the microchip imaging and measurement system, which was used to determine the positions of the modules in three dimensions. The imaging system used for the tests was previously described in Chapter 3. Calibrated distance measurements were used to quantify the degree of separation between microchips and guide the subsequent alignment of the two modules. The reference axes used to describe the positions of the modules are as defined in Figure 5.8. Movement in the x direction closes the gap between chips, while the y direction points parallel to the alignable edges of the chips and z points vertically upwards[44].

5.3.1 Image optimisation

A bright, detailed image of the chip surfaces was achieved by illuminating the chamber using an LED light source, and adjusting the incidence angle to maximise light scattered from the electrodes and into the objective lens. The quality of the image was further improved by adjusting the settings of the pixel histogram in the camera software. Images of the chips during alignment experiments are shown later in this chapter in section 5.5.

5.3.2 Distance measurement

The distance measurement feature of the camera was calibrated using known widths of electrodes on the linear chip surface. A distance of 640 µm across the span of three DC

[44]Note that this definition differs from the conventional axes shown in Chapter/,2 for surface traps, where z pointed along the surface in the axial direction of the trap.

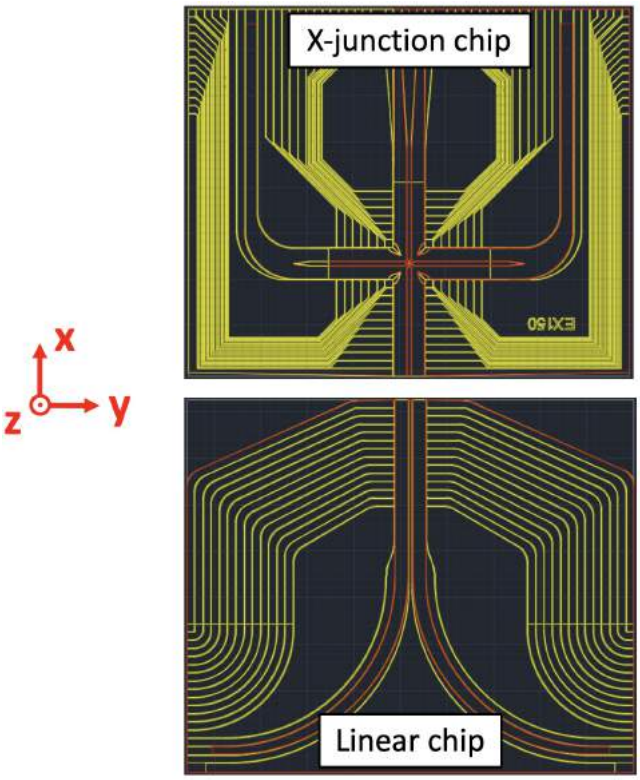

Figure 5.8: Reference axes used for chip alignment experiments.

electrodes including the gaps between them corresponded to 1638 pixels on the sCMOS sensor. By comparing the apparent size of one pixel on the image to the actual pixel size of 6.5 µm, we determined a magnification of 16.67. After inputting this value into the software, lines drawn using the measurement function showed a distance in micrometres rather than pixels. This method would be used to measure the translation of the modules in the x and y directions. To measure the positions of the modules in the z direction, a digital micrometer was used to control the height of the objective lens, which was carefully adjusted to achieve an in-focus image of the microchip surfaces. The depth of field of the lens system was 1.3 µm.

The optics were then used to observe the damage at the edge of the X-junction chip (DC electrodes #1 and #72, and the end of the RF electrodes). We saw that the damaged electrodes appeared brighter than the other electrodes, because they were raised into a ramp that scattered more of the incident light. Measurements were made using the focal plane of the optics to estimate a height difference of 120 ± 14 µm between the damaged electrodes of the X-junction and the other features on the surface of the chip. Imaging of the linear chip at the alignable edge revealed no such damage.

5.3.3 Measurement of initial positions of modules

The initial separation of the microchips was measured to be:

- Gap in x direction: 271 µm
- Misalignment in y direction: >600 µm
- Height difference in z direction: 272 µm (the X-junction being higher)

By comparing the results above to the maximum expected travel range of the module positioning system, which is around 600 µm in each axis, it was concluded that the starting separation of the modules was sufficiently small to permit alignment in the x and z axes, but that the misalignment in the y direction was at the limit of what could be achieved. This eventually required opening of the system to manually reset the positions of the modules, as detailed later in this chapter.

The next section describes the characterisation of the motion of the positioning system prior to the alignment of the two ion-trap modules. After alignment, the final positions of the modules are measured, and the offset in three dimensions is determined.

5.4 Positioning system characterisation at room temperature

The positioning system is based around an XYZ piezo translation stage as described in Chapter 3. The stage is designed to operate within the UHV environment needed for ion trapping, while also being compatible with the cryogenic temperatures of the modules during cooling. Movement of the stage was controlled using the E-727 controller, via the PI MikroMove software installed on a Windows 10 laptop. Operating the positioning system using the software involves setting a voltage of between -20 V and 120 V that drives a corresponding displacement of the piezo actuators (typically from less than a micrometre to a few hundred micrometres). The actuators operate in an open-loop configuration, without position feedback from a sensor, which was not available for the UHV variants of the actuators (due to outgassing concerns relating to an epoxy used for its construction). Therefore, it was necessary to characterise the positioning system to understand the displacement versus voltage behaviour of the actuators and the full extent of the travel range in all axes.

At the start of the characterisation process, the initial positions of the modules were as with a separation in the x direction of 271 µm. The control voltage applied to each axis was initially set to zero, and then each axis was characterised from a starting voltage of

-20 V up to the maximum voltage of 120 V. Movement was stepped using increments of 10 V. The stage was first translated forwards to the limit of its travel range, and then backwards to return to the starting voltage of -20 V. The results are shown in Figure 5.9.

5.4.1 Observations

The forward travel range of the positioning system is $162 \pm 2\,\mu m$ in the x axis, $240 \pm 2\,\mu m$ in the y axis and $283 \pm 10\,\mu m$ in the z axis. This is a reduction in travel compared to previous measurements of the piezo actuators when unloaded (as described in chapter 4), and before they were mounted as part of the positioning system.

A hysteresis loop is visible in the characterisation curves of the x and y axes (and less so for the z axis). For the y axis in particular, the backwards displacement of the stage is substantially less than the forwards displacement, and so the curve does not form a fully closed loop. This results in a position offset of 80 µm after the round trip. This hysteresis will affect the repeatability of the motion of the positioning system, and it will need to be considered for any future plans to automate the alignment process. However, for the current alignment method which uses human-in-the-loop adjustments, the effect of hysteresis is less important so long as there is sufficient travel range in each axis to bring the modules together.

A substantial off-axis drift was observed for the x and z axes (less so for the y axis). One explanation for this could be the off-centre loading of the z actuator, which supports a mass of $\sim$ 350g at a distance of 20 mm between the centre-of-mass and the point of attachment. The load is made up from a titanium bracket attached to the moving plate of the z actuator, which supports a titanium platform that holds the copper chip mount and the microchip itself. Although the weight force of ~ 3.5 N is within the operating limit of the actuator, which is 10 N, the applied torque may contribute to the observed drift. Additionally, the copper chip mount is connected to the heat sink by a flexible copper braid that links the two modules and which may contribute a small torque to the stage.

5.5 Alignment of ion-trap modules at room temperature

This section describes the alignment of the modules, which was achieved by positioning the linear ion-trap module relative to the X-junction module, so that the RF electrodes of the microchips closely lined up. Following the alignment, the distances between the RF electrodes of the two microchips were measured in all three axes.

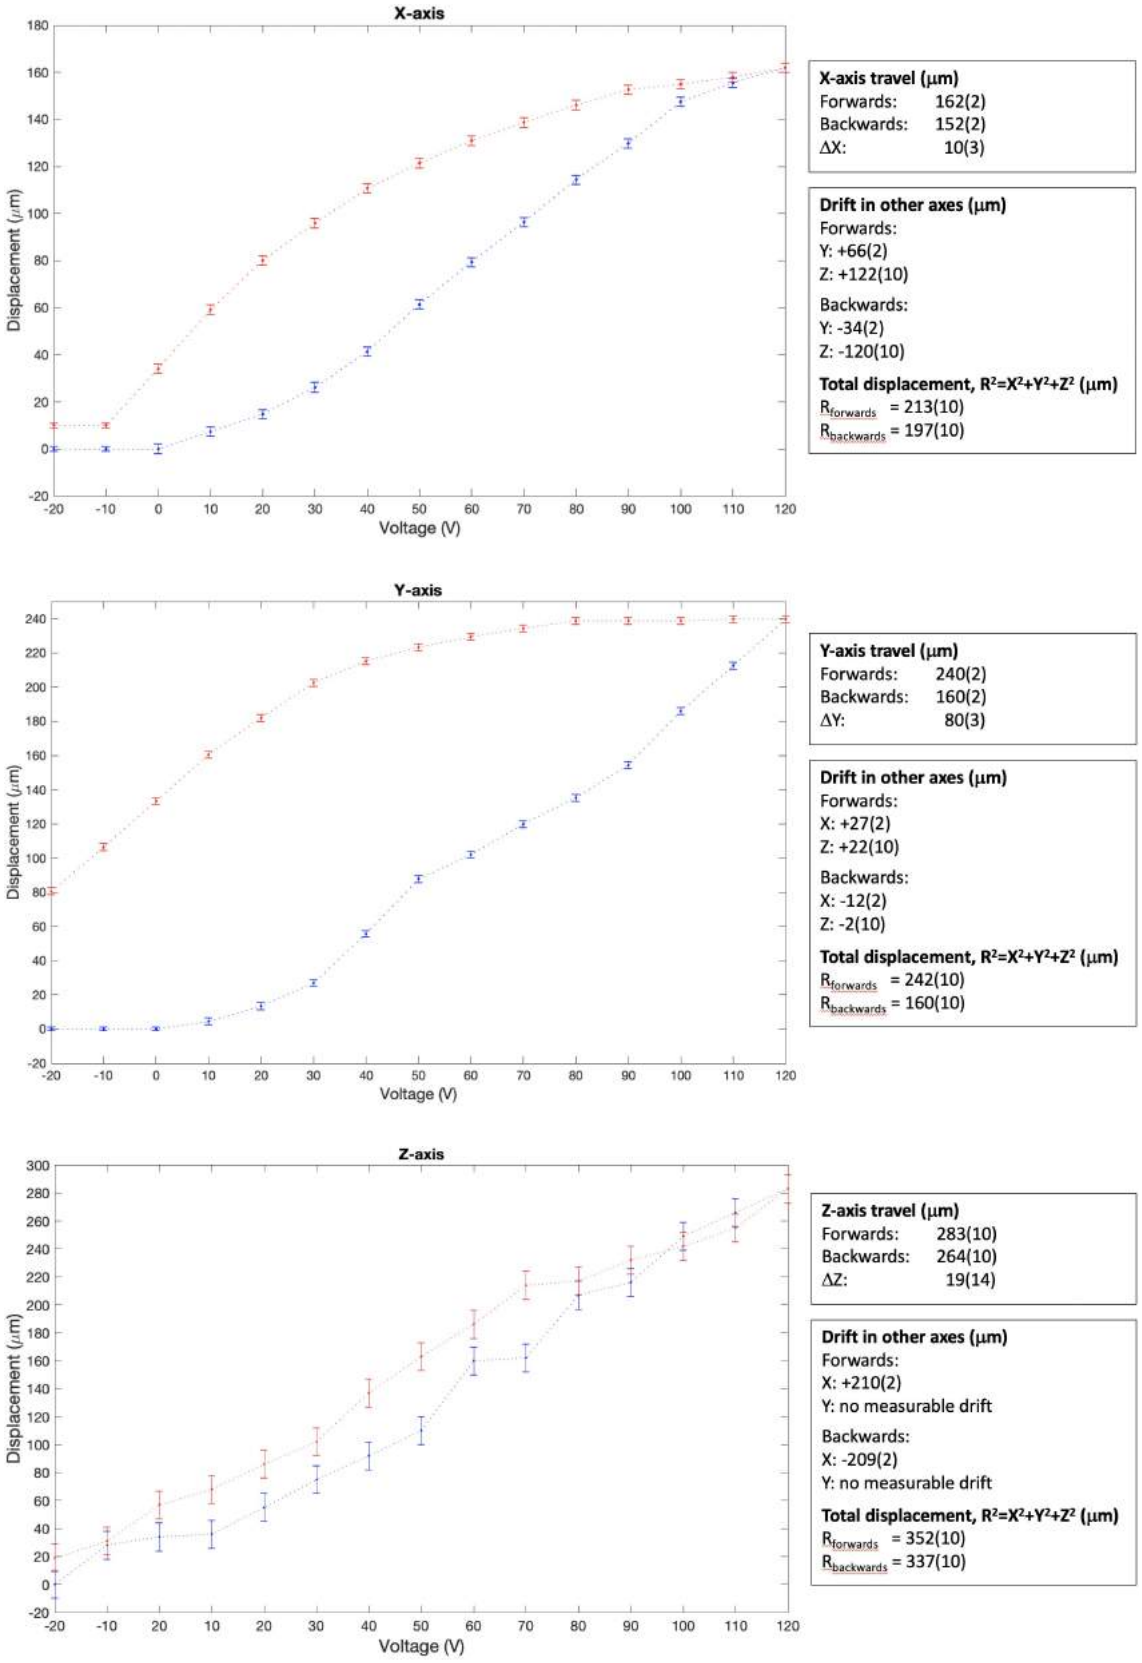

Figure 5.9: Characterisation of the module positioning system at room temperature (295 K). Graphs show stepped displacement against applied voltage for the in-vacuum XYZ piezo translation stage moving forwards (blue) and backwards (red).

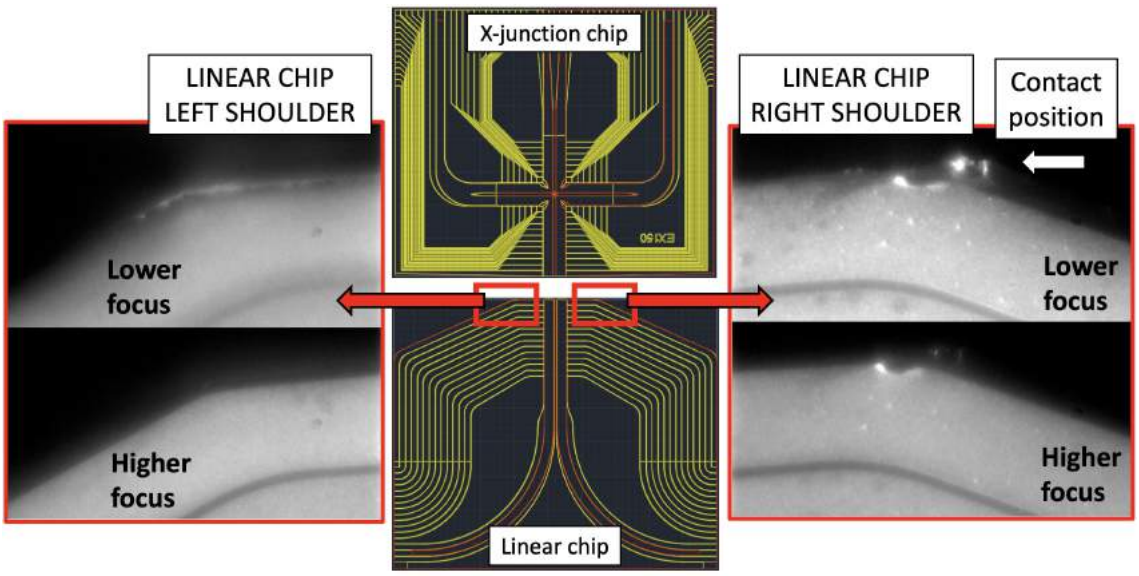

Figure 5.10: Outcrops of silicon substrate at the shoulder regions of the linear chip. The two modules came into contact at the right shoulder.

5.5.1 Initial tests

The modules started with a misalignment in the y direction of $\sim 600\,\mu\text{m}$ which was too large to be overcome by the available travel range of the XYZ stage. Nevertheless the first alignment was undertaken as an initial test to identify what final alignment was possible and to evaluate how closely we could position the microchips in the x and z directions.

As the linear module was moved close to the x-junction module we observed that the linear chip made contact with the X-junction chip at the position shown in Figure 5.10. The position of contact appeared to be a protruding layer of silicon substrate below the layer of electrodes at the right shoulder of the linear chip. After contact was made, additional forward movement of the linear module was seen to push the X-junction module backwards, and so the test was stopped to avoid damaging the chips. The positions of the modules in the region of the RF electrodes at the time of contact are shown in Figure 5.11, which is a composite of four photographs. Measurement of the final position gave the following results:

- Angle between modules: 0.56 degrees
- Gap in x direction at RF electrodes: between 31 µm and 46 µm
- Misalignment in y direction: 655 µm
- Height difference in z direction: 297 µm (the X-junction microchip was higher)

We concluded that the angle between chips was sufficient to cause the contact between modules at the right shoulder of the linear microchip (further testing demonstrated that this contact did not prevent close alignment - see description below). On completion of the first test, the vacuum chamber was opened and the modules were repositioned by hand to reduce the starting misalignment of the microchips. This was necessary because the module pre-alignment stage remained seized. Although adjusting by hand was effective, the resulting positions of the modules were difficult to predict precisely, and so iterative repositioning and retesting was required to identify a suitable starting position for the modules. In this starting position we made a deliberate misalignment in the negative y direction, so that the linear chip was to the left of the x-junction chip, as shown in the top image of Figure 5.12. This was necessary to overcome a substantial drift that was observed to occur in the positive y direction as the linear module was translated towards the x-junction module. The starting position was measured to have a 278 µm gap in the x direction, a -412 µm misalignment in the y direction, and a 24 µm height difference in the z direction (the X-junction microchip was higher).

From this position the modules were aligned as closely as possible, using the Mikromove software to adjust the step size of the XYZ piezo stage. During alignment the modules were again observed to come into contact at the right shoulder of the linear chip. However, by continuing to translate the linear module towards the x-junction module using small displacement steps, the gap between chips was eliminated, and no damage was observed. One effect of pressing the chips together in this way was that the linear chip lined up with the shoulder of the x-junction chip, which was fabricated at an angle of approximately 1 degree. Therefore a small angle was apparent between the two modules, even in the final

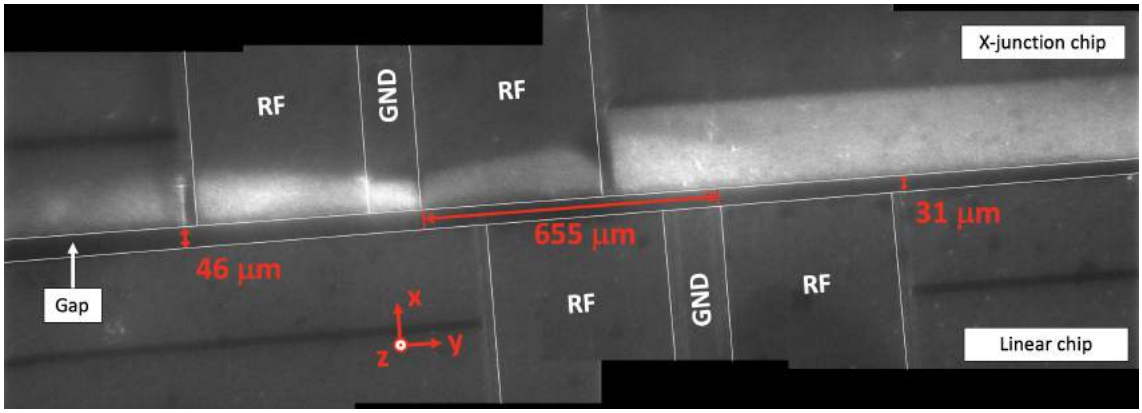

Figure 5.11: The image shows the separation between chips at the instant when contact was detected between the modules, which occurred at the right shoulder of the linear chip (out of picture). White outlines have been added for clarity. The 655 µm misalignment between RF electrodes occurred during initial installation of the chips, and was later overcome.

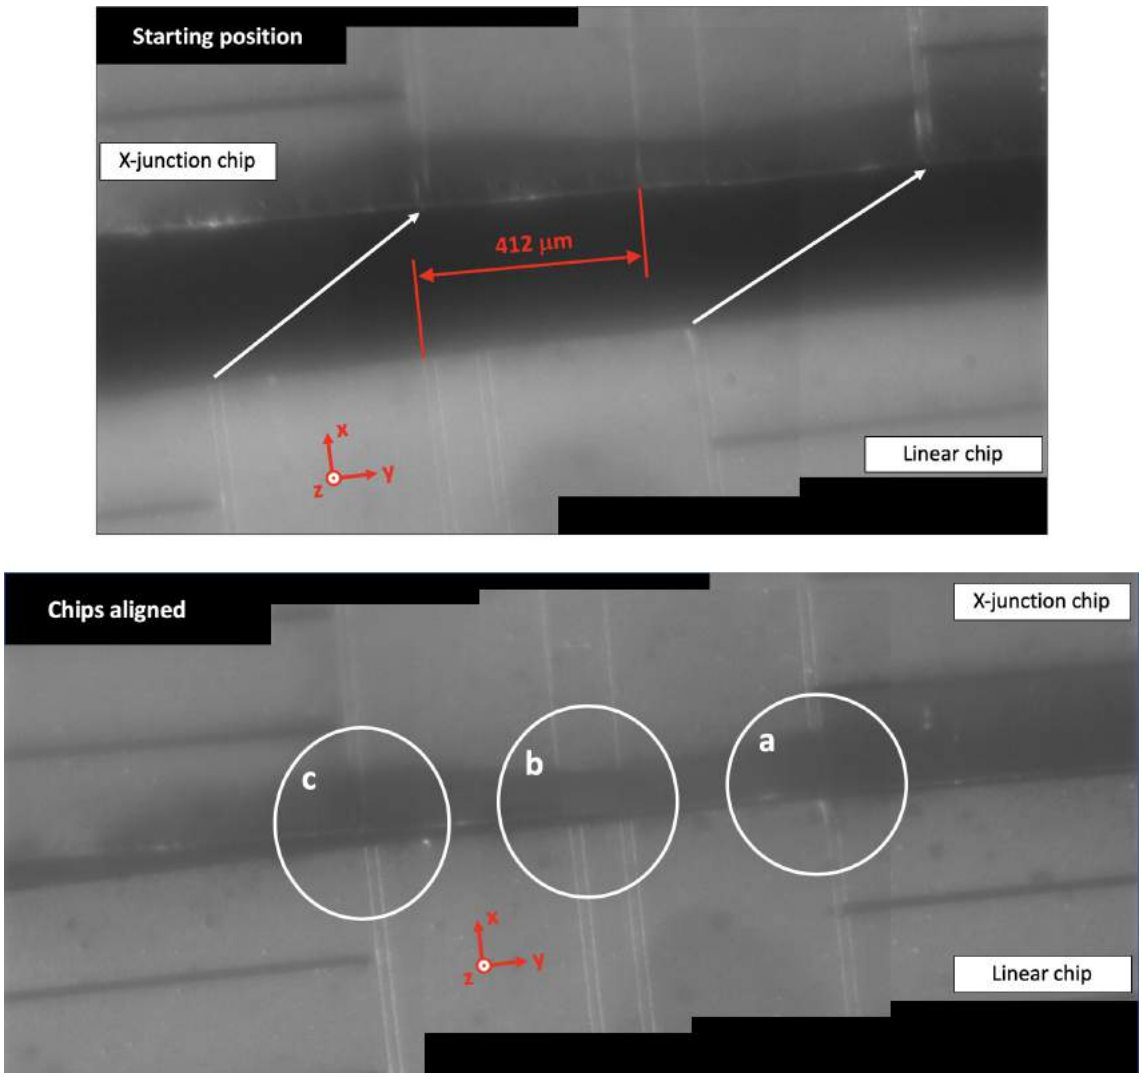

Figure 5.12: Top image shows the initial starting positions of the chips before alignment. Bottom image shows the chips after having been moved into the aligned position, with the RF-GND-RF structures closely matched-up. The shadow running along the edge of the x-junction chip was cast by the raised electrodes.

aligned position. The final position is shown in the bottom image of Figure 5.12.

5.5.2 Results of module alignment

The final positions of the chips were measured to be:

- Gap in x direction:

  At positions (a) $7 \pm 4\,\mu m$, (b) $11 \pm 4\,\mu m$, (c) $21 \pm 4\,\mu m$

- Misalignment in y direction:

  At positions (a) $6 \pm 4\,\mu m$, (b) $10 \pm 4\,\mu m$, (c) $20 \pm 4\,\mu m$

- Height difference in z direction:

  $16 \pm 14\,\mu m$ (the X-junction microchip was higher)

There was again an angle between the edges of the chips due to the contact between the shoulders of the linear and the x-junction chip. The damage at the edge of the x-junction chip made measurement more difficult due to the shadows on the image cause by the raised electrodes. For distance measurements of less than around 10 µm using the line drawing software, the lines were very small the position resolution was coarse, leading to uncertainties of a few micrometres. The uncertainty in the z direction was caused by the apparent range over which the image of the chip appeared in focus, when the height of the objective lens was adjusted using a micrometer screw.

The issue of the angle between chips can be mitigated by trimming the shoulders of the linear trap to make the leading edge narrower, which will improve the maximum permissible angle before contact occurs. Alternatively, a rotation stage may in future be integrated into the positioning system to control the angle. The rotation stage would need to meet the constraints of physical size, loading weight, UHV compatibility to 10^{-11} mbar, and low minimum operating temperature. During research at the planning stage of this project, limited options were found to be available for off-the-shelf products to meet these requirements.

This results of this section demonstrate the capability of the positioning system for aligning the two ion-trap modules, given an initial starting separation of a few hundred micrometres. In the next section I describe how the tests were repeated at low temperatures, by cooling the modules using the cryogenic helium gas circulation system, to measure the positional drift during thermal cycling and characterise the performance of the positioning system when cold.

5.6 Position drift during cryogenic cooling

To prepare for the operation of the cooling system, the main cryostat vacuum chamber and the vacuum jacket surrounding the transfer lines were pumped down to a pressure below 10^{-5} mBar, to ensure adequate thermal insulation from room temperature. The helium lines were evacuated and flushed through with helium gas to remove air. During thermometry checks, the status of the PT100 temperature sensors were found to be as follows:

Although sensors T3 and T4 were unavailable, the readings from the Helium in and Helium out positions were sufficient to infer the temperatures of the ion-trap modules to within a small interval of uncertainty ($\sim$ a few K), and since the the XY piezo stage was thermally connected to the Z piezo stage by a short OFHC copper braid, the difference in

Sensor	Location	Status
T0	XY piezo stage	OK
T1	Helium out	OK
T2	Helium in	OK
T3	Z piezo stage	Unavailable
T4	X-junction ion trap	Unavailable

Table 5.2: Status of PT100 temperature sensors during cooling test

temperature was also expected to be small.

The initial positions of the chips were measured at room temperature (295 K) and then again after the cooling system had been running for 1 hour, at which time the temperature readings were:

- Helium in: 76 K
- Helium out: 79 K
- XY piezo stage: 260 K

The XY piezo stage remains at a higher temperature than the modules because it is insulated from the mounting structure using narrow PEEK spacers, and also thermally anchored to the vacuum chamber using a copper base. This is needed to ensure that the piezo stages remain above their minimum operating temperature of -20 °C (253 K).

The drift in the positions of the modules is shown in table 5.3. We observed that the gap between modules (in the x direction) decreased due to the change in temperature, and so did the height difference between the modules. However, the misalignment in the y direction increased.

Thermal drift: separation between ion-trap modules			
	Before cooling (μm)	**After cooling** (μm)	**Change** (μm)
Distance in x axis	263	215	-48
Distance in y axis	49	74	+25
Distance in z axis	240	150	-90

Table 5.3: Thermal drift in the position of the linear ion-trap module relative to the X-junction module after cooling from room temperature to $\sim$ 79 K

5.7 Investigation of RF behaviour during chip alignment

Given that the ion trap chips in the two-module processor require to be closely aligned to within around 10 µm, and that we wish to apply substantial RF voltages (Vpeak $\sim$ 200 V)

to drive the chips, I conducted a number of experimental tests to determine the electrical behaviour of the chips during alignment of the edge electrodes.

Simulations by Martin Siegele determined that for a trap separation of 1 µm in x (and with zero misalignment in y and z) there would be a capacitance of $\sim 1{\times}10^{-15}$ F between the edge electrodes of the two chips. This is a very small capacitive coupling, and it implies that we should not expect the voltage signals applied to one chip to affect the other.

To check this result, RF voltages were applied to the chips and the voltage signals were monitored while the two modules were aligned to within $\sim$ 10 µm in all axes. The experimental setup for the tests was as follows. A dual channel signal generator[45] was used to provide sinusoidal voltage signals to two amplifiers[46] and then to two helical resonators (described in Chapter 4). The output voltages from the resonators were applied to the chips via a capacitive divider at the output of each resonator. The capacitive dividers are set-up to send a small ratio of the voltage signal from the resonator to an oscilloscope to measure the voltage applied to the chip without the additional capacitance of the oscilloscope affecting the resonant circuit.

The initial separation of the chips was 280 µm in x, 110 µm in y, and 290 µm in z. In this position, the resonant frequencies were measured for each chip, and these were found to be 16.86 MHz for the linear and 12.06 MHz for the X-junction chip. The difference in resonant frequency was due to a significant difference between the capacitances of the chips - measured to be 15 pF for the linear and 40 pF for the X-junction chip (including the in-vacuum cable connections). For the first test, the chips were driven at an equal frequency of 13.8 MHz and the amplitudes of the signals were both 34 Vpp (peak-to-peak). The chips were then aligned to within $\sim$ 10 µm in each axis. During the alignment, no change was seen to occur in the amplitudes or phases of the signals. After alignment, the voltages were increased to 96 Vpp for the linear and 56 Vpp for the X-junction chip. These were the maximum achievable voltages since the resonators were both offset from their resonant frequencies. The power dissipated was around 5 W per chip. The phase difference between the signals was almost exactly 180 degrees, and this was adjusted on the signal generator until the signals were in phase. No change in the voltage amplitudes was observed to occur for the two signals as the phase was adjusted, nor was there a change in the measured power. No phase locking of the signals was observed to occur.

At this point, with the chips in the aligned position, the resonant frequencies were remeasured to check if the close alignment of the chips caused any effect, but no change

[45] Rigol DG4062
[46] Mini Circuits LZY-22+

was observed. This was expected, since the estimated capacitive coupling between the two chips was so small ($\sim 1\,\mathrm{fF}$).

Finally, the chips were driven at their resonant frequencies of 16.86 MHz for the linear and 12.06 MHz for the X-junction chip. During this test the voltages were independently increased to 385 Vpp for the linear and 491 Vpp for the X-junction chip, with no effect caused by one signal on the other.

This led to the conclusion that the capacitive coupling between the two chips was immeasurably small at a separation of $\sim$10 µm, and no detrimental effect was observed to occur when the chips were operated at realistic ion-trapping voltages of $\sim 400\,\mathrm{Vpp}$.

5.8 Conclusion

This chapter has described the construction and characterisation of a two module ion trap processor, designed for connecting surface traps using electric fields to permit ion transport across their interface. The individual microchips used for this experiment endured a range of damage during preparation, but it was also shown that a cleaning treatment using solvents and sonication was an effective means to remove shorted connections by removing fragments of conductive material from between the electrodes. The chips were functional for ion trapping, although sub-optimal for ion shuttling between the two chips due to the lifting of the electrodes at the edge of the X-junction chip, and the surface contamination at the edge of the linear chip. However, new batches of single layer chips and magnetic gradient chips were being fabricated at the time of writing, for use in future experiments.

The experiments described in this chapter have demonstrated the effectiveness of the imaging, measurement and positioning systems that were used to align the ion trap modules to within approximately 10 µm in all axes. The experimental issues arising from the alignment tests were attributable to (1) the travel range of the piezos and the ability to set the initial starting position of the linear module with respect to the x-junction module and (2) the angle between microchips causing contact at the shoulders of the chips. The issue of the starting position will shortly be addressed by a newly designed pre-alignment stage which is currently being manufactured. The issue of the angle between chips can be mitigated by narrowing the leading edge of the linear chip, or alternatively, attempting to integrate a rotation stage into the positioning system in future, to control the angle.

Simulations of the trapping potential have indicated that the chip alignment precision achieved in testing is within the constraints of trap depth and RF barrier necessary to achieve ion confinement in the space between modules. Therefore we conclude that ion

transport across the interface between aligned modules is achievable given the precision of the module positioning system. However, the height of the barrier and the ion heating rate are particularly sensitive to z axis misalignment of the microchips due to exposure of the ion to the underlying ground plane below the surface of the traps, which may be resolved by making adjustments to the trap structure during microfabrication.

The microchips were tested at realistic RF trapping voltages, and an investigation into the electrical behaviour of the chips during alignment showed that there was no measurable capacitive coupling between their edges, and no crosstalk was detected between the signals applied to the chips.

Chapter 6

Ion trapping and transport between independent surface-trap modules using second generation chips

6.1 Introduction

This chapter describes the ion trapping and transport experiments that were performed with ^{174}Yb+ ions using the second generation ion trap chips. After discussing the installation of the chips into the demonstrator experiment, I provide details of successful trapping of ^{174}Yb+ ions and initial characterisation of the axial and radial secular frequencies. I then describe shuttling operations with single ions, in which the ion was trapped and moved to a stable location close to the edge of one chip, and then reliably transported between chips for many round trips.

6.2 Preparation of chips for ion trapping

The second generation chips were produced for the demonstrator experiment by Martin Siegele and Seokjun Hong using facilities at the Center of MicroNanoTechnology at EPFL. As an additional microfabrication step, the backside of each chip was coated with a 10 nm chromium adhesion layer and a 0.5 µm layer of gold was deposited by physical vapour deposition (PVD). This backside gold layer was added to improve the effectiveness of indium die-bonding used to attach the chips to mounting blocks for installation into the experiment as discussed in the next section.

Two EL125 linear surface traps labelled #13 and #14 were selected for use in the

experiment after a visual inspection of the trap surfaces under an optical microscope at 50x magnification and preliminary electrical testing using a needle probe station to check for shorts between electrodes. Section 6.2.3 provides full details of the electrical status of the chips as they were installed into the experiment, after die-bonding and wire bonding was completed.

6.2.1 Microchip die-bonding using indium

The ion trap chips were bonded to gold-plated copper mounting blocks using indium, which provides a strong mechanical connection and is an excellent thermal interface material owing to its high thermal conductivity as described in Chapter 6. Plating the copper blocks with a 3 µm layer of gold improves thermal contact between the block and the top surface of the cryogenic heat sink, which is similarly plated with a 3 µm gold layer. Additionally, the gold surface coating avoids the issues of wettability of molten indium to a copper surface which is hindered by the formation of an oxide layer on copper [127]. Indium wire with 1.2 mm diameter was melted onto the gold surface of the block using a soldering ion at 450 °C while simultaneously heating the block using a heat gun at 500 °C. Care was taken not to remove the gold layer on the surface of the block with the soldering ion during this process. After cooling, an uneven layer of indium oxide formed on top of the indium

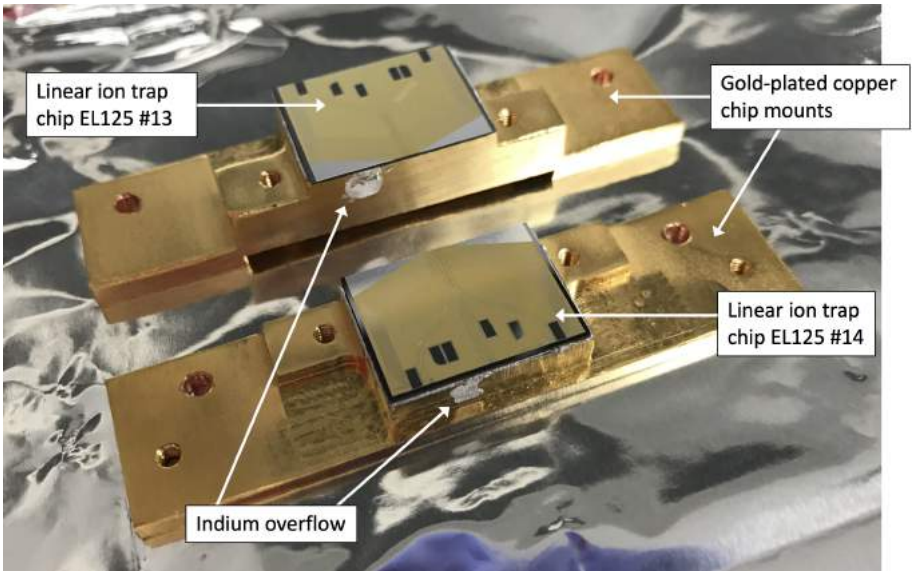

Figure 6.1: Second generation EL125 linear ion trap chips after die bonding to gold-plated copper mounting blocks. Each chip is placed so that the alignable edge overhangs the mounting block by approximately 0.3 mm to aid in chip alignment. Small amounts of indium overflow result from the pressure applied by the die-bonder placement head, and these were easily removed using a scalpel, leaving behind a small residue which does not affect the experiment.

layer. To improve bonding of the microchip to this surface, it is pressed against a highly polished aluminium plate which transfers a flat mirror finish to the indium oxide surface. The mounting block with indium layer is then immersed into a 10% aqueous solution of hydrochloric acid for 30 s to remove indium oxide, and the microchip is attached to the indium surface using the placement head of a die-bonder. The die-bonder was programmed to perform a temperature ramp at $1\,^{\circ}\mathrm{C\,s^{-1}}$ to a target temperature of 220 °C, chosen to be well above the indium melting point of 156.6 °C. After holding for 1 min to ensure indium flow, the temperature was lowered at a rate of $-2\,^{\circ}\mathrm{C\,s^{-1}}$. Figure 6.1 shows the microchips bonded to the mounting blocks.

6.2.2 Magnetic gradient coil connections

The magnetic gradient coils were connected on the heat sink side of the experiment only, by using the gold-plated beryllium copper clamps. The magnetic coils on the piezo side were not connected, so as to reduce the significant heat load due to high currents in the on-chip microfabricated wires. For quantum logic experiments with the two-module processor this is sufficient to provide a single gate zone which is made accessible to all qubits via ion transport between the trap modules.

6.2.3 Wire bonding

The ion trap chips were electrically connected to PCBs by wire bonding performed by Sam Hile using a 4523A digital wedge bonder with 25 µm diameter aluminium wire. The DC electrodes closest to the edge, which are most critical for ion transport, were each double wire-bonded to two traces on the PCB, which allowed for testing of the connections after closing the vacuum system by applying a voltage to the first trace and measuring on the second trace. Chip #13 provides additional DC electrodes beyond the capacity of the 2x 15-pin DSUB connectors on the PCB, so electrodes furthest from the edge were grouped together by wire bonds. Details for each chip are as follows:

- Chip #13 (piezo side): DC electrodes 1-5 and 31-35 were double wire-bonded. DCs 6-10, 11-14, 22-25, and 26-30 were grouped.
- Chip #14 (heat sink side): DCs 1-10 and 26-35 were double wire-bonded. DCs 30 and 31 shorted by microfabrication defect. Outer rotation DC15 connected to DC14, and outer rotation DC21 connected to DC22 in order to make wire bonding accessible due to positioning of clamps for CCWs.

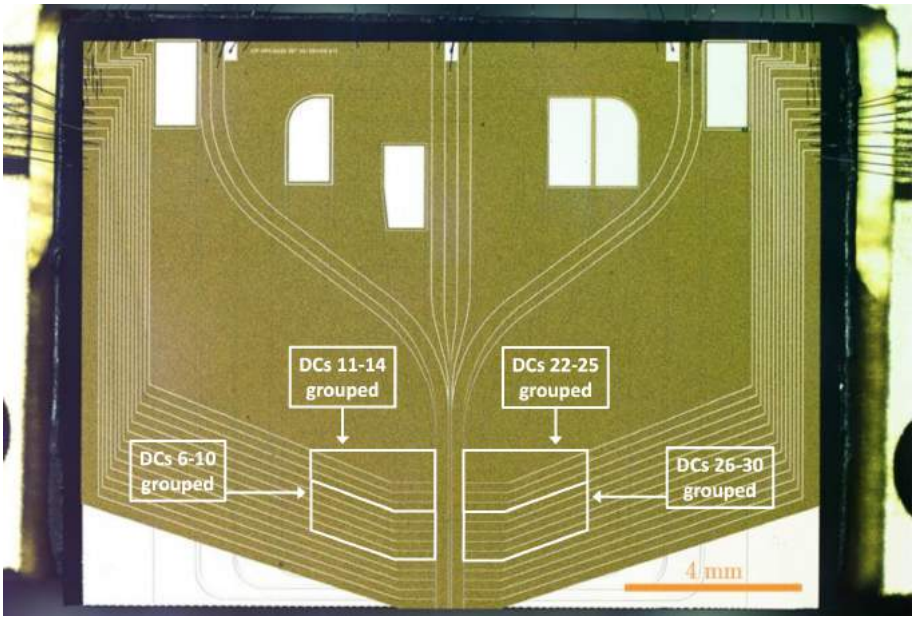

Figure 6.2: Microscope image of EL125 linear surface trap chip #13 from the production wafer of second generation traps. This chip was installed into the module on the piezo side of the experiment. The chip is shown after being die-bonded to a gold plated copper mounting block and wire-bonded to the surrounding PCB. Electrical testing identified DC electrodes 11, 12 and 13 were shorted together due to imperfect lithography. These electrodes are furthest from the alignable edge and were not required during experiments.

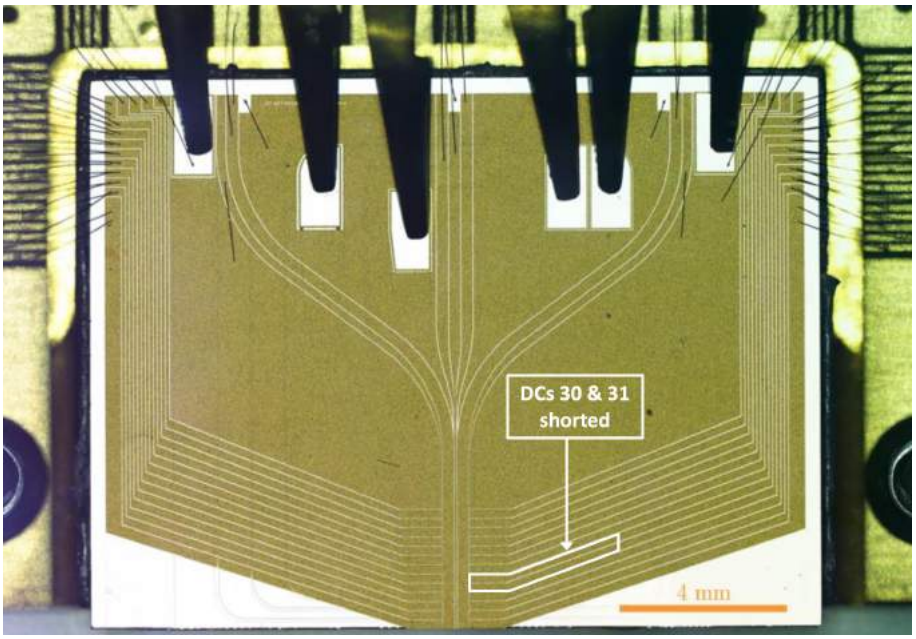

Figure 6.3: Microscope image of EL125 linear surface trap chip #14 from the production wafer of second generation traps. This chip was installed into the ion trap module on the heat sink side of the experiment. As shown, the chip has been die-bonded to a gold plated copper block and wire bonded to the surrounding PCB. High current connections to the on-chip magnetic gradient coils are made using six beryllium copper clamps. Electrical testing identified that DC electrodes 27 and 28 were shorted together by a gold filament joining the electrodes with a resistance of 4.9 MΩ.

Figure 6.2 and Figure 6.3 show the status of the chips after wire bonding and electrical testing was completed and immediately before installation onto the ion trap modules of the experiment.

6.2.4 Pre-alignment of ion trap modules

The ion-trap chip, mounting block and PCB assemblies were installed onto the ion-trap module support structure within the vacuum chamber as shown in the top left image of Figure 6.4. Prior to closing the vacuum chamber, the modules were manually positioned to ensure there was sufficient separation between chips to prevent contact due to thermal drifts in position during baking and cooling. These drifts were estimated in Chapter 3. The pre-aligned separation between modules was estimated to be $276 \pm 5\,\mu m$ in the x-direction and $89 \pm 5\,\mu m$ in the y-direction, as shown in Figure 6.4. These distances are within the achievable 300-600 µm displacement range of the piezo stages used for in-vacuum alignment of the modules.

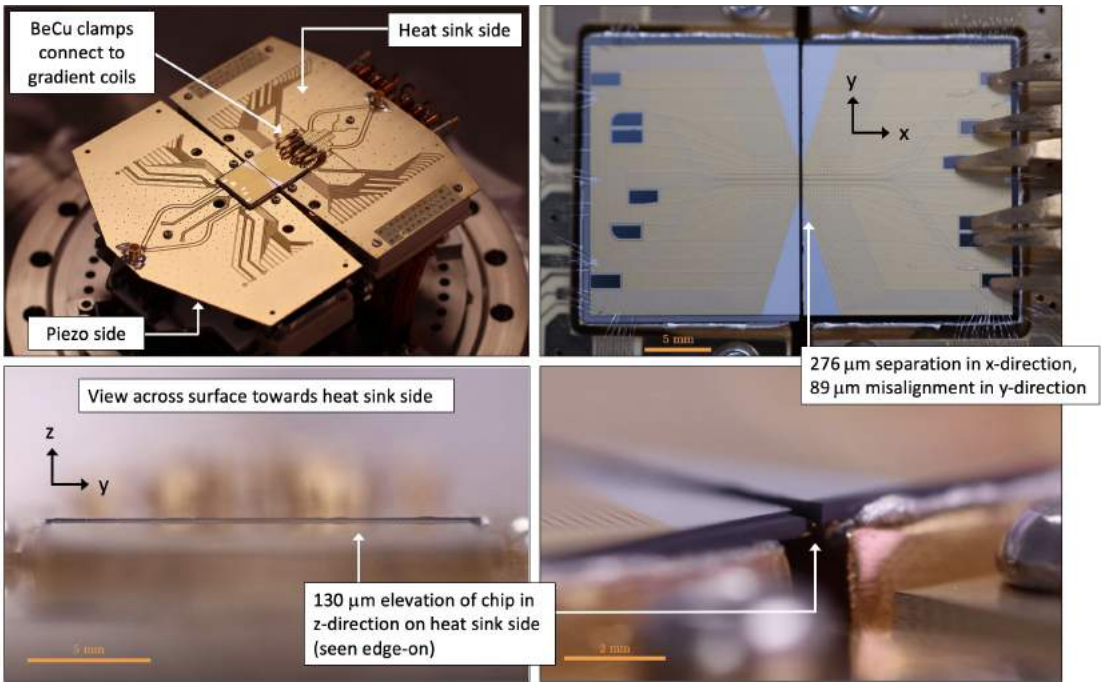

Figure 6.4: *(Top left)* Manual alignment of the ion trap modules immediately before closing the vacuum chamber. The trap on the heat sink side has beryllium copper clamps installed to connect the on-chip magnetic gradient coils, while the chip on the piezo side does not. The modules were set up with an initial separation to prevent thermal drifts in position from causing the chips to collide. *(Top right)* Close-up view of the ion trap surfaces showing trap separation of $276 \pm 5\,\mu m$ in the x-direction and $89 \pm 5\,\mu m$ in the y-direction. *(Bottom left)* Edge-on view showing elevation of the chip on the heat sink side by $130 \pm 20\,\mu m$ in the z-direction. *(Bottom right)* Oblique side view showing the slight elevation of the module on the heat sink side.

6.2.5 Experimental adjustments to achieve appropriate trap stability parameter

Measurements of the capacitance of the in-vacuum part of the RF circuit (including the chips, PCBs and RF cabling) via the SMA feedthrough flange of the vacuum chamber were significantly higher than anticipated with values of 96 pF for the heat sink module and 95 pF for the piezo module. This affected the RF setup used to supply high voltage, low noise RF potentials to the traps and resulted in resonant frequencies at 11.3 MHz for the ion trap modules. To understand if a trap frequency, $\Omega = 11.3\,\mathrm{MHz}$ was suitable for producing stable confinement of ^{174}Yb ions, the trap stability parameter, q, was calculated at a voltage amplitude, $V = 100\,\mathrm{V}$, using two different methods. The first method used Equation 2.13 (previously given in Chapter 2) which is restated here for perfect hyperbolic electrodes,

$$q_x = -q_y = \frac{2eV}{mr_0^2\Omega^2},$$

where e is the electron charge, m is the mass of the $^{174}Yb^+$ ion and r_0 is the ion-electrode distance of 125 µm.

This resulted in $q = 1.5$ (assuming hyperbolic electrodes). Typically, a value between 0.1 to 0.3 is required for stable ion trapping within a surface trap [128]. Since surface trap electrodes deviate significantly from a perfect hyperbolic shape, a second method was used to estimate the q-parameter for a surface trap with a gapless five-wire geometry using the analytic equations derived by House in Ref [129],

$$q_{\mathrm{surf}} = \frac{eV}{m\Omega^2}\frac{8(b-a)}{\pi\sqrt{ab}(a+b)^2}, \tag{6.1}$$

where the electrodes are assumed to be infinitely long, a is defined as the distance from the centre to the edge of the ground electrode, and b is defined as the distance from the centre to the outer edge of the RF electrodes. For a real, non-gapless geometry the distances are measured to the centre of the gaps between electrodes. For the electrode dimensions of the second generation ion trap chips, $a = 42\,\mu\mathrm{m}$ and $b = 318\,\mu\mathrm{m}$. This resulted in a value of $q_{\mathrm{surf}} = 0.54$. Both results for the stability parameter are above the typical values of q between 0.1 and 0.3 for stable confinement in a surface trap [128].

To reduce the q parameter below 0.3, the trap frequency was increased towards a target value of 20 MHz while keeping the voltage amplitude unchanged at 100 V. This required adjustment to the helical resonators within the RF setup to increase the resonant

frequency of the RF filtering to match the trap frequency. Since the resonant frequency has the relationship $\omega = \frac{1}{\sqrt{LC}}$, an increase was achieved by reducing the inductance of the main pickup coil of the helical resonator to compensate for the high capacitance of the ion trap chips. By reducing the number of turns on the coil from 9.5 to 4.5, and increasing the winding pitch from 8 mm to 17 mm, the coil inductance decreased from an estimated 1.9 µH to an estimated 0.6 µH. Once these modifications were complete, the helical resonators were re-connected to the ion trap chips via the SMA feedthrough of the chamber, and the following measurements were taken for resonant frequency, $f = \omega/2\pi$:

- For the heat-sink module: Resonant frequency, $f_{\mathrm{hs}} = 19.343\,\mathrm{MHz}$
- For the piezo module: Resonant frequency, $f_{\mathrm{pz}} = 19.317\,\mathrm{MHz}$

The centre frequencies of the two helical resonators were then matched by using a variable capacitor to tune the resonator on the piezo side of the experiment. The final measurements of resonant frequency and quality factor, Q, were:

- For the heat-sink module: Resonant frequency, $f_{\mathrm{hs}} = 19.331\,526\,\mathrm{MHz}$, $Q_{\mathrm{hs}} = 60$
- For the piezo module: Resonant frequency, $f_{\mathrm{pz}} = 19.331\,426\,\mathrm{MHz}$, $Q_{\mathrm{pz}} = 68$

The stability parameter was re-calculated for a trap frequency of 19.33 MHz and voltage amplitude of 100 V, to give $q = 0.48$ for hyperbolic electrodes and $q_{\mathrm{surf}} = 0.18$ for the surface trap approximation. This result for q_{surf}, which is the most appropriate parameter, falls within the expected range for stable ion confinement within the trap.

6.3 Trapping of $^{174}Yb^{+}$

This section describes the methods used to calculate the required DC potentials to apply to ion trap electrodes, the initial parameters used to successfully trap $^{174}Yb^{+}$ ions and the measurement of secular frequencies using the tickling method.

6.3.1 Calculation of DC potentials

The DC electrodes are used to create a harmonic trapping well along the axial direction of the chip. The procedure for calculating the required DC potentials can be outlined in three steps: simulate the trap, impose physical constraints on the output potential to achieve a harmonic well, minimise the necessary electrode voltages. First the trap was simulated using FEM methods (in Comsol) to generate the potential above the chip surface due to

each DC electrode individually, with all others set to 0 V. To achieve this, the simulation was run N times for each of the N electrodes of the trap, starting with electrode 1 which was set to 1 V and all others were set to 0 V. This was repeated for electrode 2, 3 and so on, each time with 0 V applied to all other electrodes. The resulting potentials, $\tilde{\phi}_n(\boldsymbol{r})$ are termed basis functions and they can be multiplied by the actual voltage applied to the electrode, V_n and then summed to give the total potential:

$$\Phi(\boldsymbol{r}) = \sum_{n=1}^{N} V_n \tilde{\phi}_n(\boldsymbol{r}) + \phi_p(\boldsymbol{r}) \tag{6.2}$$

where we have also added the pseudopotential, $\phi_p(\boldsymbol{r})$. The pseudopotential can be calculated by first modelling a static potential of 1 V applied to the RF electrodes and then using the time independent approximation:

$$\phi_p(\boldsymbol{r}) = \frac{q}{4m\Omega_{RF}^2}(V_{RF}\boldsymbol{\nabla}\tilde{\phi}_n(\boldsymbol{r}))^2 \tag{6.3}$$

where V_{RF} is the voltage amplitude applied to the RF electrode, and q and m are the charge and mass of the ion, respectively. The next step in determining the DC voltages is to apply the physical constraints to the total potential, $\Phi(\boldsymbol{r})$ to ensure that the result is a confining potential suitable for trapping. First we impose that there exists a minimum at the ion position $\boldsymbol{r_0}$,

$$\boldsymbol{\nabla}\Phi(\boldsymbol{r_0}) = 0. \tag{6.4}$$

Next we wish to constrain the potential well to have an axial curvature that is related to our desired axial secular frequency, ω_z, by

$$\frac{\partial^2}{\partial z^2}\Phi(\boldsymbol{r}) = \frac{m}{q}\omega_z^2. \tag{6.5}$$

Finally, since there are an infinite number of solutions which meet the criteria set out above, we constrain the voltages to be of a reasonable magnitude by imposing that

$$V_n^2 < 10\,\text{Volts}. \tag{6.6}$$

A numerical solver was developed by Sam Hile to read-in FEM basis functions and implement the constraints described above and output DC voltages suitable for trapping.

6.3.2 Trapping position

Prior to setting up for trapping runs, fluorescence testing was performed with the natural ytterbium oven to identify the exact position at which the atomic beam intersected the RF rails of the traps. This defines the trapping position on the chip surface. By scanning the position of the 399 nm laser across the trap at the ion height of 125 µm, fluorescence was observed to occur on the heat sink side of the experiment at the position of DC electrodes 10 and 26, as shown schematically in Figure 6.5. Once this position was identified, the chip was simulated using the method described in the previous section to calculate the required DC trapping voltages at this location.

6.3.3 Trapping parameters

The RF voltage was measured in two ways. The first was a calibrated directional power meter placed between the RF amplifier and helical resonator. A power of 700 mW corresponded to an RF amplitude of 100 V applied to the chip. This was corroborated by a capacitive divider setup at the output of the helical resonator which was used to measure a ratio of 1/54 that of the voltage applied to the chip (the setup was described in Chapter 4), and which read an amplitude of 3.62 Vpp corresponding to 195 Vpp at the chip. The DC potentials applied to the chip are shown in Figure 6.6. All rotation electrodes were initially set to 0 V and after initial trapping the inner rotations were updated to +1 V and -1 V.

The 369 nm, 399 nm and 935 nm laser beams were overlapped and aligned to the 125 µm ion height above the trap surface. The laser wavelengths were:

- 369 nm laser: 739.05022-25 nm (before frequency doubling) at a power of 50 µW
- 399 nm laser: 398.91136 nm at a power of 200 µW
- 935 nm laser: 935.17984-86 nm, 3 mW (power broadened)

Trapping of 14 ions was achieved after firing the ovens using a target temperature of 300 °C for 100 s. After initial trapping the ovens were adjusted to a target temperature of 276 °C which resulted in two trapped ions after 20 s. Figure 6.7 shows sCMOS camera images of one, two and four ions within the trap.

Ion lifetime measurements were performed by interrupting the Doppler cooling lasers (by blocking the laser beams using a motorised shutter) for increasing time intervals (during which the ion is dark) and then re-applying Doppler cooling and re-acquiring the ion on the sCMOS camera. The ion lifetime is determined by the timescale at which the ion can

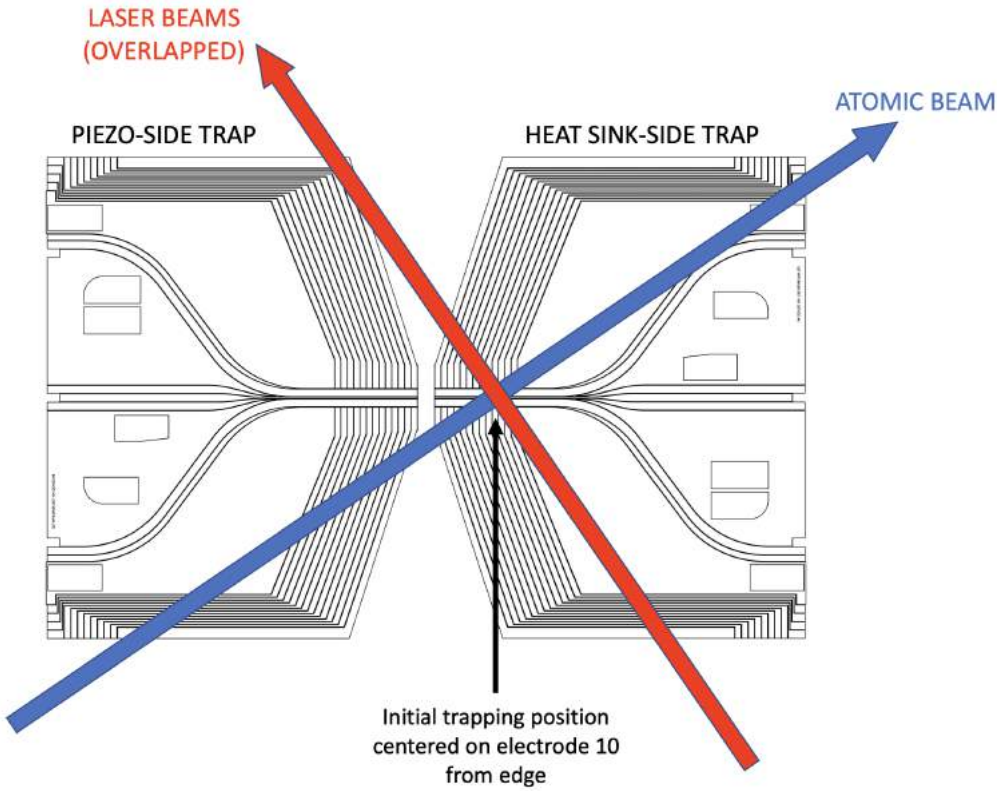

Figure 6.5: The initial trapping position was centred on electrode 10 on the heat sink side of the experiment. This was defined by the location at which the RF rails were intersected by the neutral atomic beam from the fixed atomic oven loaded with natural ytterbium.

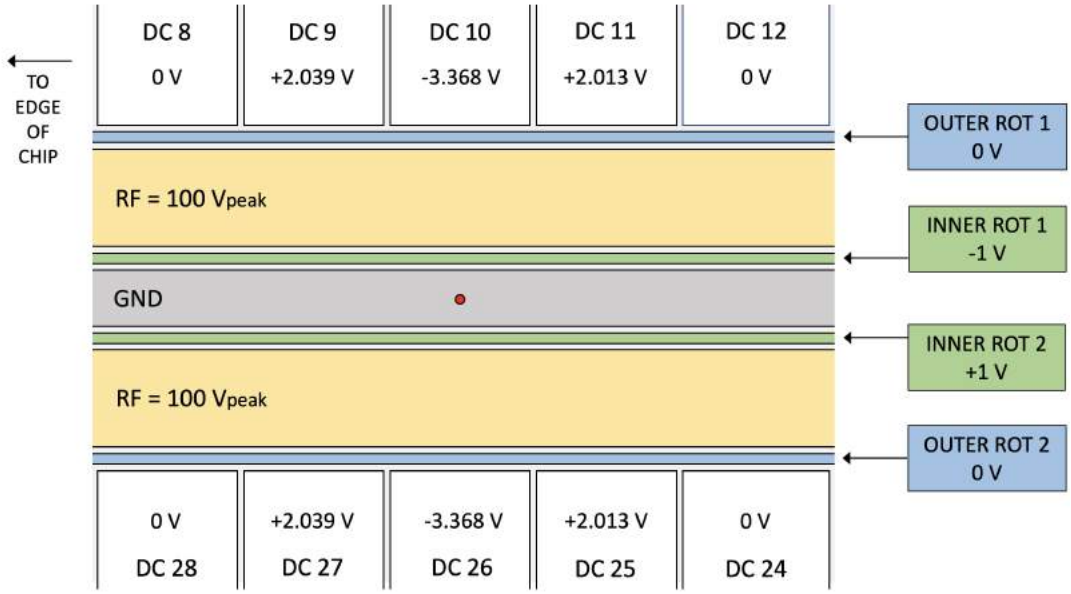

Figure 6.6: Potentials applied to EL125 chip for trapping $^{174}Yb^{+}$ ions. The DC potentials were calculated using FEM simulations of the trap in combination with a numerical solver that imposed physical constraints on the resultant potential to generate a harmonic well with zero electric field at the ion position and a defined axial secular frequency.

no longer be reacquired. Initial ion lifetime was measured to be 42 minutes, and the longest lifetime recorded was upwards of 5 hours, which was limited by laser stability.

Initial trapping was achieved at room temperature with a background pressure in the vacuum chamber of 3.3×10^{-10} mbar. For subsequent trapping runs the cooling system was activated and a temperature of 40 K was measured from the PT100 sensor attached to the chip mount on the heat sink side of the experiment. Activating the cooling system lowered the temperature at the chips to 40 K and reduced the background pressure in the chamber to 6.2×10^{-11} mbar.

6.3.4 Measurements of trap secular frequencies

Measurements of the secular frequencies of ions within the trap were made by monitoring the perturbing effect of applying a 'tickle' frequency to the trap. Scanning of the applied frequency over a range of values results in excitation of the ion's motion when approaching resonance with the axial or secular frequencies of the trap. To implement this method, a signal generator was used to apply an RF tickle signal to one of the inner rotation electrodes of the trap (designated DC19) which was chosen for its proximity to the trap centre. A BNC to IDC connector was used to connect from the signal generator to the DC19 channel of the Yenesei external filter board. The motion of the ion was observed to stretch out along the principal direction when excited at the corresponding secular frequency. The axial secular frequency was measured to be 120 kHz and this required a voltage amplitude of 5 mV. The radial secular frequency was observed at around 1 MHz, and this required a voltage amplitude of 100 mV to be applied. The difference in voltage amplitudes was attributed to the combined attenuating effect of the filtering from the Yenesei filter box, which has two-stage low-pass RC filtering with a cut-off frequency of 72 kHz, and the

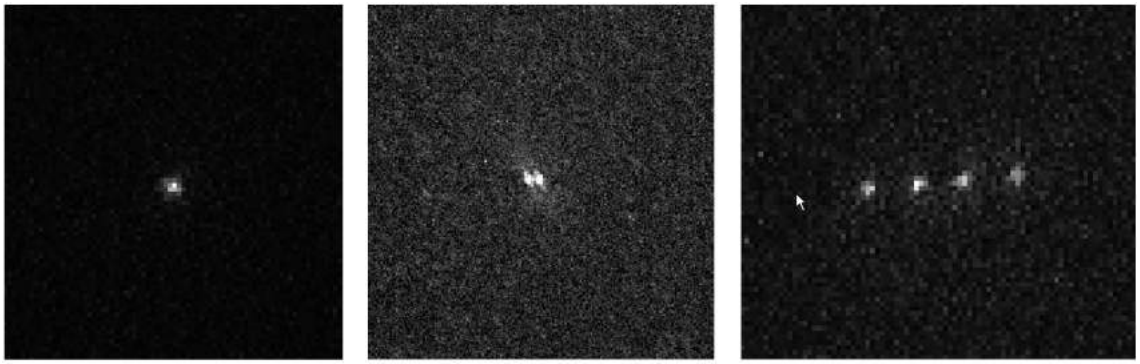

Figure 6.7: Image captured from EMCCD camera showing one, two and four $^{174}Yb^{+}$ ions confined on the EL125 ion trap microchip. The ion separation is approximately 6 µm for two ions and 4 µm for four ions when applying the DC potentials shown in Fig. 6.6

.

in-vacuum first-order low-pass filters with cut-off frequency of 257 kHz.

6.4 Transport operations with $^{174}Yb^{+}$ ions

In this section I describe the work that was conducted by Mariam Akhtar and Falk Bonus to transport ions on the second generation chips. Ions were first transported from the initial trapping position on the heatsink side of the experiment to a location close to the edge on the same trap. After this the ions were successfully transported between ion trap modules, proving the capability of the demonstrator system to connect modules using electric fields.

6.4.1 Generating transport waveforms

The geometry of the two aligned ion traps was simulated using FEM methods, taking account of the discontinuous electrodes at the trap edges and the separation between chips. The simulation outputs a grid of values for the basis potentials and RF pseudopotential with a resolution of 2 µm in the region where the ion passes through. The numerical solver described in section 6.3.1 was used to generate the time dependent DC potentials, or waveforms, for ion transport. The same physical constraints that were imposed on the electric field, secular frequency and maximum DC value for the static trapping voltages are also imposed for the time dependent waveforms. The harmonic potential well is translated along the trap axis and the ion remains stably trapped even when traversing the pseudopotential barrier in the gap region between modules.

6.4.2 Initial transport operations on a single ion trap module

The first implementation of ion transport using the above methodology was to move the ion from its initial trapping position centred on electrode 10 (on the heat sink side of the experiment) to the edge of the trap at a stable position centred on electrode 2. This was a distance of 1,840 µm. The shorted electrodes 30 and 31 were used as part of the voltage waveform and were constrained by the solver to have the same potential owing to the shorted connection. The two DC electrodes (6 and 5) directly opposite the shorted pair, on the other side of the RF rails, were also constrained by the solver in the same manner.

6.4.3 Stable trapping at the edge using DC endcaps of both chips

The trapping location near to the edge of the chip on the heat sink side of the experiment is shown in Figure 6.8. Stable trapping in this position was achieved using a 5-point trap

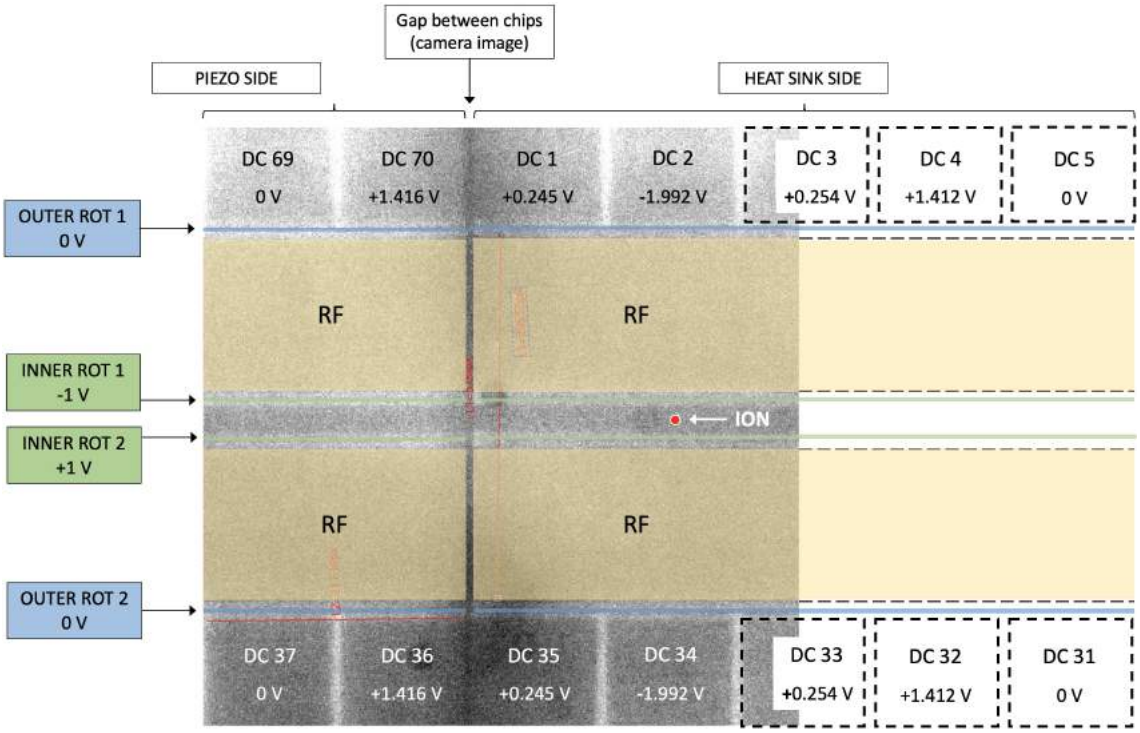

Figure 6.8: Image from the sCMOS camera showing the gap (dark vertical line) between chips after alignment of the ion trap modules. The colour overlay is used to highlight RF electrodes in yellow, and inner and outer rotation electrodes in green and blue respectively. The voltages shown were used to trap an ion (indicated by the red dot) on the heat sink side such that the minimum of the DC potential was centred between DC2/DC34. This location was then used as the starting position for the ion when implementing transport round trips between the two chips.

(5 pairs of DC electrodes) that was formed by the DC electrodes of both chips, located on either side of the gap between modules. It can be seen in Figure 6.8 that the axial gap between the modules is approximately the same width as the microfabricated trenches separating individual DC electrodes, which are 10 µm wide.

6.4.4 Ion transport between aligned modules

The ion was transported from the stable position centred on DC electrode 2 on the heat sink side of the experiment, over to the equivalent position on the piezo side of the experiment (DC37 as shown in Figure 6.8) and back again. The shuttling waveforms translated the ion in a potential well formed by a 5-point trap. Doppler cooling was applied to the ion on the heat sink side of the experiment only. The ion was transported back and forth between the modules for 2×10^7 round trips without loss, which is a transport fidelity of 99.999995 %. This was limited by the ion lifetime of 3 hours. The ion transfer rate between modules was 1.212 kHz for round trip transport and 2.424 kHz for one way transport.

6.5 Conclusion

The demonstrator experiment has shown reliable transport of ion qubits between separate surface-trap modules connected by electric fields. This is the first example of ion transport between independent ion trap chips. High transfer rates of 2.4 kHz were achieved for one way transport between the trap modules with a near unity success probability. This is an order of magnitude higher rate than has been achieved by photonic interconnects [56] for distributing quantum entanglement between ion trap modules.

Chapter 7

Conclusion

The experimental work described in this thesis is aimed at developing a scalable architecture for building a modular, microwave-based quantum computer using ion traps connected by electric fields. Towards this end, my PhD project has developed a two-module ion trap processor that serves as a demonstrator device for implementing ion transport between chips to realise a high-rate quantum link. The demonstrator incorporates microfabricated surface-electrode ion-trap chips with alignable edges, a precision positioning system for module alignment, a scalable cooling system with extensible cryostat, and a versatile vacuum system for ion trapping experiments.

Chapter 1 provided an introduction to quantum computing, along with a summary of current research towards achieving a scalable, universal architecture for building a quantum computer based on trapped ions.

Chapter 2 explained the theoretical background to ion trapping, which included a description of the fundamentals of microfabricated ion trap technology. After this, the magnetic gradient scheme was described, which permits quantum logic operations to be performed using long-wavelength radiation. The creation of dressed state qubits in the hyperfine ground state manifold of $^{171}Yb^{+}$ was described, with an explanation of how this technique is used to extend the coherence times of magnetically sensitive qubits.

Chapter 3 introduced the two-module ion trap processor and set out the experimental requirements for module alignment and ion shuttling. I described the individual component systems of the ion trap modules.

In Chapter 4, the experimental set-up was given including the vacuum chamber for experiments. A system was developed for delivering low noise RF trapping voltages with impedance matching to the ion-trap chips, with a method for tuning the circuit by $\sim$2 MHz to optimise for a single RF frequency. This chapter also discussed the delivery system for

microwave and RF radiation. A scalable cryogenic cooling system was presented, able to achieve an operating temperature of 70 K with an application cooling power of 111 W, which was connected to four independent ion trap experiments within the laboratory.

Chapter 5 of the thesis describes the assembly and characterisation of the two-module processor, including post-fabrication testing and treatment of the surface trap chips to prepare them for ion-trapping. This was followed by calibration of the imaging optics to permit characterisation of the module positioning system, enabling the alignment of the surface traps to within ∼10 µm in all axes. The implications for ion trapping and shuttling between modules were considered, with reference to simulations of the trapping potential and ion dynamics. The RF behaviour of the chips during alignment was investigated at voltages of up to 380 V (peak to peak) for the linear and 490 V (peak to peak) for the X-junction chip, and showed no capacitive crosstalk between the two modules.

Chapter 6 presented the latest results of ion trapping and shuttling experiments to connect the ion traps of the two-module quantum processor.

The demonstrator experiment is a step towards the realisation of a large-scale quantum computer based on trapped ions.

Bibliography

[1] B. J. Copeland. *Colossus: The First Electronic Computer: The secrets of Bletchley Park's code-breaking computers*. Popular Science Series. OUP Oxford, 2006.

[2] H. H. Goldstine. *The Computer from Pascal to Von Neumann*. History e-book project. Princeton University Press, 1980.

[3] A. G. Bromley. Charles Babbage's Analytical Engine, 1838. *IEEE Annals of the History of Computing*, 20(4):29–45, Oct 1998.

[4] UK National Quantum Technologies Programme. `http://uknqt.epsrc.ac.uk/`. Accessed: 2019-05-03.

[5] European Union Quantum Flagship Research Programme. `https://qt.eu/`. Accessed: 2019-05-03.

[6] C. Monroe, M. G. Raymer, and J. Taylor. The U.S. National Quantum Initiative: From Act to action. *Science*, 364(6439):440–442, 2019.

[7] Google AI Quantum Research Programme. `https://ai.google/research/teams/applied-science/quantum-ai/`. Accessed: 2019-05-03.

[8] IBM Q - Quantum Computing Research. `https://www.research.ibm.com/ibm-q/`. Accessed: 2019-05-03.

[9] Microsoft Quantum Computing Research. `https://www.microsoft.com/en-gb/quantum/`. Accessed: 2019-05-03.

[10] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum Machine Learning. *Nature*, 549(7671):195–202, 2017.

[11] R. Orús, S. Mugel, and E. Lizaso. Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4:100028, 2019.

[12] P. J.J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis. Scalable quantum simulation of molecular energies. *Physical Review X*, 6(3), 2016.

[13] R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven, and G. K. L. Chan. Low-Depth Quantum Simulation of Materials. *Physical Review X*, 8(1), 2018.

[14] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge University Press, New York, NY, USA, 10th edition, 2011.

[15] TOP500 Lists. `https://www.top500.org/lists/`. Accessed: 2021-07-19.

[16] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018.

[17] S. Aaronson and L. Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference*, CCC '17, pages 22:1–22:67, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[18] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, Jun 1982.

[19] Smite-Meister - Own work. *CC BY-SA 3.0*, Downloaded 2019.

[20] D. E. Deutsch and R. Penrose. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818):97–117, 1985.

[21] D. E. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[22] D. E. Deutsch and R. Penrose. Quantum Computational Networks. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 425(1868):73–90, 1989.

[23] P. W. Shor. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.

[24] L. K. Grover. A fast quantum mechanical algorithm for database search. *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 212–219, 1996.

[25] P. W. Shor. Fault-tolerant Quantum Computation. *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 56–, 1996.

[26] A. M. Steane. Error correcting codes in quantum theory. *Physical Review Letters*, 77(5):793–797, 1996.

[27] J. Preskill. Reliable quantum computers. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 454(1969):385–410, 1998.

[28] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75(25):4714–4717, 1995.

[29] J. I. Cirac and P. Zoller. Quantum Computation with Cold Trapped Ions. *Phys. Rev. Lett.*, 74(20):4091–4094, 1995.

[30] D. P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.

[31] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. A programmable two-qubit quantum processor in silicon. *Nature*, 555(7698):633–637, 2018.

[32] Y. Wu, Y. Wang, X. Qin, X. Rong, and J. Du. A programmable two-qubit solid-state quantum processor under ambient conditions. *npj Quantum Information*, 5(1):9, 2019.

[33] G. Wendin. Quantum information processing with superconducting circuits: A review, 2017.

[34] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger. Blueprint for a microwave trapped ion quantum computer. *Science Advances*, 3(2), 2017.

[35] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2), 2016.

[36] Y. Nam, J. Chen, N. C. Pisenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, A. M. Ducore, K. M. Hudek, M. Keesan, S. M. Kreikemeier, J. Mizrahi, P. Solomon, M. Williams, J. D. Wong-Campos, C. Monroe, and J. Kim. Ground-state energy estimation of the water molecule on a trapped ion quantum computer. *arXiv e-prints*, Feb 2019.

[37] I. G. Ryabinkin, S. N. Genin, and A. F. Izmaylov. Constrained Variational Quantum Eigensolver: Quantum Computer Search Engine in the Fock Space. *Journal of Chemical Theory and Computation*, 15(1):249–255, 2019.

[38] G. E. Santoro and E. Tosatti. Optimization using quantum mechanics: quantum annealing through adiabatic evolution. *Journal of Physics A: Mathematical and Theoretical*, 41(20):209801, may 2008.

[39] V. S. Denchev, S. Boixo, S. V. Isakov, N. Ding, R. Babbush, V. Smelyanskiy, J. Martinis, and H. Neven. What is the computational value of finite-range tunneling? *Physical Review X*, 6(3):1–17, 2016.

[40] J. Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.

[41] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical Review Letters*, 113(22), 2014.

[42] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas. High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves. *Physical Review Letters*, 117(14):140501, 2016.

[43] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J. N. Zhang, L. M. Duan, D. Yum, and K. Kim. Single-qubit quantum memory exceeding ten-minute coherence time. *Nature Photonics*, 11(10):646–650, 2017.

[44] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A - Atomic, Molecular, and Optical Physics*, 62(2):11, 2000.

[45] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt. 14-qubit entanglement: Creation and coherence. *Physical Review Letters*, 106(13), 2011.

[46] N. Friis, O. Marty, C. Maier, C. Hempel, M. Holzäpfel, P. Jurcevic, M. B. Plenio, M. Huber, C. Roos, R. Blatt, and B. Lanyon. Observation of Entangled States of a Fully Controlled 20-Qubit System. *Physical Review X*, 8(2), 2018.

[47] G. Poulsen, Y. Miroshnychenko, and M. Drewsen. Efficient ground-state cooling of an ion in a large room-temperature linear paul trap with a sub-hertz heating rate. *Phys. Rev. A*, 86:051402, Nov 2012.

[48] S. Weidt, J. Randall, S. C. Webster, E. D. Standing, A. Rodriguez, A. E. Webb, B. Lekitsch, and W. K. Hensinger. Ground-state cooling of a trapped ion using long-wavelength radiation. *Phys. Rev. Lett.*, 115:013002, Jun 2015.

[49] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. Sympathetic cooling of trapped ions for quantum logic. *Physical Review A*, 61(3):1–8, 2000.

[50] M. Guggemos, D. Heinrich, O. A. Herrera-Sancho, R. Blatt, and C. F. Roos. Sympathetic cooling and detection of a hot trapped ion by a cold one. *New Journal of Physics*, 17(10), 2015.

[51] D. S. Wang, A. G. Fowler, and L. C. L. Hollenberg. Surface code quantum computing with error rates over 1%. *Physical Review A - Atomic, Molecular, and Optical Physics*, 83(2):1–4, 2011.

[52] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(3), 2012.

[53] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L. M. Duan, and J. Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A - Atomic, Molecular, and Optical Physics*, 89(2):022317, 2014.

[54] D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. Monroe. Modular entanglement of atomic qubits using photons and phonons. *Nature Physics*, 11(1):37–42, 2015.

[55] T. Walker, K. Miyanishi, R. Ikuta, H. Takahashi, S. Vartabi Kashanian, Y. Tsujimoto, K. Hayasaka, T. Yamamoto, N. Imoto, and M. Keller. Long-Distance Single Photon Transmission from a Trapped Ion via Quantum Frequency Conversion. *Physical Review Letters*, 120(20), 2018.

[56] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance. High-rate, high-fidelity entanglement of qubits across an elementary quantum network. *Phys. Rev. Lett.*, 124:110501, Mar 2020.

[57] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–11, 2002.

[58] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259–328, 1997.

[59] M. A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. Transport of quantum states and separation of ions in a dual rf ion trap. *Quantum Info. Comput.*, 2(4):257–271, June 2002.

[60] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation. *Applied Physics Letters*, 88(3):034101, 2006.

[61] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. High-Fidelity Transport of Trapped-Ion Qubits through an X-Junction Trap Array. *Physical Review Letters*, 102(15):153002, 2009.

[62] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. H. Wesenberg, M. J. Biercuk, D. Leibfried, and D. J. Wineland. Near-ground-state transport of trapped-ion qubits through a multidimensional array. *Phys. Rev. A*, 84:032314, Sep 2011.

[63] Z. D. Romaszko, S. Hong, M. Siegele, R. K. Puddy, F. R. Lebrun-Gallagher, S. Weidt, and W. K. Hensinger. Engineering of microfabricated ion traps and integration of advanced on-chip features. *Nature Reviews Physics*, 2(6):285–299, 2020.

[64] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Surface-electrode architecture for ion-trap quantum information processing. *Quantum Info. Comput.*, 5(6):419–439, September 2005.

[65] K. Wright, J. M. Amini, D. L. Faircloth, C. Volin, S. Charles Doret, H. Hayden, C. S. Pai, D. W. Landgren, D. Denison, T. Killian, R. E. Slusher, and A. W. Harter. Reliable transport through a microfabricated X-junction surface-electrode ion trap. *New Journal of Physics*, 15(3):033004, 2013.

[66] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. Controlling fast transport of cold trapped ions. *Phys. Rev. Lett.*, 109:080501, Aug 2012.

[67] M. Palmero, R. Bowler, J. P. Gaebler, D. Leibfried, and J. G. Muga. Fast transport of mixed-species ion chains within a paul trap. *Phys. Rev. A*, 90:053408, Nov 2014.

[68] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. Coherent diabatic ion transport and separation in a multizone trap array. *Phys. Rev. Lett.*, 109:080502, Aug 2012.

[69] P. Kaufmann, T. F. Gloger, D. Kaufmann, M. Johanning, and C. Wunderlich. High-fidelity preservation of quantum information during trapped-ion transport. *Phys. Rev. Lett.*, 120:010501, Jan 2018.

[70] M. W. van Mourik, E. A. Martinez, L. Gerster, P. Hrmo, T. Monz, P. Schindler, and R. Blatt. Coherent rotations of qubits within a surface ion-trap quantum computer. *Phys. Rev. A*, 102:022611, Aug 2020.

[71] Y. Wan, R. Jördens, S. D. Erickson, J. J. Wu, R. Bowler, T. R. Tan, P. Hou, D. J. Wineland, A. C. Wilson, and D. Leibfried. Ion transport and reordering in a 2d trap array. *Advanced Quantum Technologies*, 3(11):2000028, 2020.

[72] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. Demonstration of the trapped-ion quantum CCD computer architecture. *Nature*, 592(7853):209–213, 2021.

[73] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger. Trapped-Ion Quantum Logic with Global Radiation Fields. *Physical Review Letters*, 117(22):1–12, 2016.

[74] F. Mintert and C. Wunderlich. Ion-trap quantum logic using long-wavelength radiation. *Physical Review Letters*, 87(25):257904–1–257904–4, 2001.

[75] W. Paul. Electromagnetic traps for charged and neutral particles. *Rev. Mod. Phys.*, 62:531–540, Jul 1990.

[76] H. G. Dehmelt. Radiofrequency spectroscopy of stored ions I: Storage. volume 3 of *Advances in Atomic and Molecular Physics*, pages 53 – 72. Academic Press, 1968.

[77] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. *Rev. Mod. Phys.*, 87:637–701, Jun 2015.

[78] P. K. Ghosh. *Ion Traps.* International Series of Monographs on Physics. Clarendon Press, 1995.

[79] L. Ruby. Applications of the mathieu equation. *American Journal of Physics*, 64(1):39–44, 1996.

[80] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A - Atomic, Molecular, and Optical Physics*, 61(6):8, 2000.

[81] J. H. Wesenberg. Ideal intersections for radio-frequency trap networks. *Phys. Rev. A*, 79:013416, Jan 2009.

[82] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous heating in ion traps. *Physical Review Letters*, 97(10):1–4, 2006.

[83] I. A. Boldin, A. Kraft, and C. Wunderlich. Measuring Anomalous Heating in a Planar Ion Trap with Variable Ion-Surface Separation. *Physical Review Letters*, 120(2), 2018.

[84] J. A. Sedlacek, A. Greene, J. Stuart, R. McConnell, C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Distance scaling of electric-field noise in a surface-electrode ion trap. *Physical Review A*, 97(2), 2018.

[85] K. Y. Lin, G. H. Low, and I. L. Chuang. Effects of electrode surface roughness on motional heating of trapped ions. *Physical Review A*, 94(1):1–11, 2016.

[86] D. A. Hite, K. S. McKay, S. Kotler, D. Leibfried, D. J. Wineland, and D. P. Pappas. Measurements of trapped-ion heating rates with exchangeable surfaces in close proximity. *MRS Advances*, 2(41):2189–2197, 2017.

[87] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-Fold Reduction of Electric-Field Noise in an Ion Trap Cleaned With Insitu Argon-Ion-Beam Bombardment. *Physical Review Letters*, 109(10):7, 2012.

[88] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Reduction of heating rate in a microfabricated ion trap by pulsed-laser cleaning. *New Journal of Physics*, 13(12):123023, dec 2011.

[89] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. Measurement of ion motional heating rates over a range of trap frequencies and temperatures. *Physical Review A - Atomic, Molecular, and Optical Physics*, 91(4), 2015.

[90] D. F. Murgia. Microchip ion traps with high magnetic field gradients for microwave quantum logic. *PhD Thesis, Imperial College London*, 2017.

[91] M. D. Hughes, B. Lekitsch, J. A. Broersma, and W. K. Hensinger. Microfabricated ion traps. *Contemporary Physics*, 52(6):505–529, 2011.

[92] D. Cho, S. Hong, M. Lee, and T. Kim. A review of silicon microfabricated ion traps for quantum information processing. *Micro and Nano Systems Letters*, 3(1):2, 2015.

[93] D. L. Moehring, C. Highstrete, D. Stick, K. M. Fortier, R. Haltli, C. Tigges, and M. G. Blain. Design, fabrication and experimental demonstration of junction surface ion traps. *New Journal of Physics*, 13, 2011.

[94] D. T. C. Allcock, T. P. Harty, C. J. Ballance, B. C. Keitch, N. M. Linke, D. N. Stacey, and D. M. Lucas. A microfabricated ion trap with integrated microwave circuitry. *Applied Physics Letters*, 102(4), 2013.

[95] C. Highstrete, S. Scott, J. D. Sterk, E. J. Heller, P. Maunz, D. Nordquist, J. E. Stevens, C. P. Tigges, and M. G. Blain. Technology for On-Chip Qubit Control with Microfabricated Surface Ion Traps. *Sandia National Laboratories Report SAND2013-9513*, 2013.

[96] J. T. Merrill, C. Volin, D. Landgren, J. M. Amini, K. Wright, S. C. Doret, C. S. Pai, H. Hayden, T. Killian, D. Faircloth, K. R. Brown, A. W. Harter, and R. E. Slusher. Demonstration of integrated microscale optics in surface-electrode ion traps. *New Journal of Physics*, 13, 2011.

[97] M. Aymar, A. Debarre, and O. Robaux. Highly excited levels of neutral ytterbium. Multichannel quantum defect analysis of odd- and even-parity spectra. *Journal of Physics B: Atomic and Molecular Physics*, 13(6):1089–1109, 1980.

[98] J. McLoughlin. Development and Implementation of an Yb Ion Trap Experiment Towards Coherent Manipulation and Entanglement. *PhD Thesis, University of Sussex*, 2011.

[99] D. J. Wineland, Wayne M. Itano, J. C. Bergquist, and R. G. Hulet. Laser-cooling limits and single-ion spectroscopy. *Phys. Rev. A*, 36:2220–2232, Sep 1987.

[100] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Graduate Texts in Contemporary Physics book series. Springer, New York, NY, 1999.

[101] C. J. Foot. *Atomic Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2007.

[102] J. Randall. High fidelity entanglement of trapped ions using long wavelength radiation. *PhD Thesis, Imperial College London*, 2016.

[103] R. Noek, G. Vrijsen, D. Gaultney, E. Mount, T. Kim, P. Maunz, and J. Kim. High speed, high fidelity detection of an atomic hyperfine qubit. *Optics Letters*, 38(22):4735, 2013.

[104] M. Kasevich and S. Chu. Atomic interferometry using stimulated raman transitions. *Phys. Rev. Lett.*, 67:181–184, Jul 1991.

[105] N. Akerman, N. Navon, S. Kotler, Y. Glickman, and R. Ozeri. Universal gate-set for trapped-ion qubits using a narrow linewidth diode laser. *New Journal of Physics*, 17(11):113060, Nov 2015.

[106] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504, Aug 2016.

[107] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-fidelity universal gate set for $^{9}Be^{+}$ ion qubits. *Phys. Rev. Lett.*, 117:060505, Aug 2016.

[108] J. Randall, S. Weidt, E. D. Standing, K. Lake, S. C. Webster, D. F. Murgia, T. Navickas, K. Roth, and W. K. Hensinger. Efficient preparation and detection of microwave dressed-state qubits and qutrits with trapped ions. *Physical Review A - Atomic, Molecular, and Optical Physics*, 91(1):1–8, 2015.

[109] N. Timoney, I. Baumgart, M. Johanning, A. F. Varon, Ch. Wunderlich, M. B. Plenio, and A. Retzker. Quantum Gates and Memory using Microwave Dressed States. *Nature*, 476(7359):185–188, 2011.

[110] S. C. Webster, S. Weidt, K. Lake, J. J. McLoughlin, and W. K. Hensinger. Simple manipulation of a microwave dressed-state ion qubit. *Physical Review Letters*, 111(14):1–5, 2013.

[111] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.

[112] T. Navickas. Towards high-fidelity microwave driven multi-qubit gates on microfabricated surface ion traps. *PhD Thesis, University of Sussex*, 2017.

[113] D. Hucul, M. Yeo, S. Olmschenk, C. Monroe, W. K. Hensinger, and J. Rabchuk. On the transport of atomic ions in linear and multidimensional ion trap arrays. *Quantum Info. Comput.*, 8(6):501–578, July 2008.

[114] Z. D. Romaszko. Design and control of microfabricated ion traps for quantum computing. *PhD Thesis, University of Sussex*, 2020.

[115] A. Bermudez, X. Xu, R. Nigmatullin, J. O'Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation. *Phys. Rev. X*, 7:041061, Dec 2017.

[116] R. Patterson, A. Hammoud, and S. Gerber. Performance of various types of resistors at low temperatures. Technical report, NASA Glenn Research Center, Cleveland, Ohio, 2001.

[117] F. Teyssandier and D. Préle. Commercially available capacitors at cryogenic temperatures. Technical report, Ninth International Workshop on Low Temperature Electronics-WOLTE9, 2010.

[118] B. C. Wadell. *Transmission line design handbook*. Artech House, Boston, Mass., 1991.

[119] H. Muzaffar and K. S. Williams. Evaluation of complete elliptical integrals of the first kind at singular moduli. *Taiwanese Journal of Mathematics*, 10(6):1633 – 1660, 2006.

[120] H. C. Miller. Inductance formula for a single-layer circular coil. *Proceedings of the IEEE*, 75(2):256–257, 1987.

[121] J. D. Siverns, L. R. Simkins, S. Weidt, and W. K. Hensinger. On the application of radio frequency voltages to ion traps via helical resonators. *Applied Physics B: Lasers and Optics*, 107(4):921–934, 2012.

[122] J. R. James and P. S. Hall. *Handbook of microstrip antennas*. P. Peregrinus on behalf of the Institution of Electrical Engineers, London, U.K., 1989.

[123] A. Grounds. Cryogenic technologies for scalable trapped-ion quantum computing. *PhD Thesis, University of Sussex*, 2018.

[124] C. Day. Basics and applications of cryopumps. In *Vacuum in Accelerators: Specialized Course of the CERN Accelerator School*, pages 241–274, 5 2006.

[125] R. A. Matula. Electrical resistivity of copper, gold, palladium, and silver, 1979.

[126] C. J. Glassbrenner and G. A. Slack. Thermal conductivity of silicon and germanium from 3 K to the melting point. *Physical Review*, 134(4A), 1964.

[127] D. M. Jacobson and G. Humpston. Gold coatings for fluxless soldering. *Gold Bulletin*, 22(1):9–18, 1989.

[128] Seokjun Hong, Minjae Lee, Hongjin Cheon, Taehyun Kim, and Dong-il "Dan" Cho. Guidelines for designing surface ion traps using the boundary element method. *Sensors*, 16(5), 2016.

[129] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Phys. Rev. A*, 78:033402, Sep 2008.