

# Scalable trap technology for quantum computing with ions

by

Amira M. Eltony

S.M., Electrical Engineering and Computer Science  
Massachusetts Institute of Technology (2013)

B.A.Sc., Engineering Physics  
University of British Columbia (2010)


Submitted to the Department of  
Electrical Engineering and Computer Science  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science  
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

© Massachusetts Institute of Technology 2015. All rights reserved.



Author ... Signature redacted .....  
Department of  
Electrical Engineering and Computer Science  
May 20, 2015

Certified by Signature redacted .....  
Isaac L. Chuang  
Professor of Physics  
Professor of Electrical Engineering and Computer Science  
Thesis Supervisor

Accepted by .. Signature redacted .....  
Leslie Kolodziejski  
Chairman, Department Committee on Graduate Students



# Scalable trap technology for quantum computing with ions

by

Amira M. Eltony

Submitted to the Department of  
Electrical Engineering and Computer Science  
on May 20, 2015, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Electrical Engineering and Computer Science

## Abstract

Quantum computers employ quantum mechanical effects, such as superposition and entanglement, to process information in a distinctive way, with advantages for simulation and for new, and in some cases more-efficient algorithms. A quantum bit is a two-level quantum system, such as the electronic or spin state of a trapped atomic ion. Physics experiments with single atomic ions acting as “quantum bits” have demonstrated many of the ingredients for a quantum computer. But to perform useful computations these experimental systems will need to be vastly scaled-up.

Our goal is to engineer systems for large-scale quantum computation with trapped ions. Building on established techniques of microfabrication, we create ion traps incorporating exotic materials and devices, and we investigate how quantum algorithms can be efficiently mapped onto physical trap hardware. An existing apparatus built around a bath cryostat is modified for characterization of novel ion traps and devices at cryogenic temperatures (4 K and 77 K).

We demonstrate an ion trap on a transparent chip with an integrated photodetector, which allows for scalable, efficient state detection of a quantum bit. To understand and better control electric field noise (which limits gate fidelities), we experiment with coating trap electrodes in graphene. We develop traps compatible with standard CMOS manufacturing to leverage the precision and scale of this platform, and we design a Single Instruction Multiple Data (SIMD) algorithm for implementing the QFT using a distributed array of ion chains. Lastly, we explore how to bring it all together to create an integrated trap module from which a scalable architecture can be assembled.

Thesis Supervisor: Isaac L. Chuang

Title: Professor of Physics

Professor of Electrical Engineering and Computer Science





## Acknowledgments

First, I want to express my sincere gratitude to Professor Isaac Chuang for contributing so much to my development over these past five years. Under his expert guidance, I have learned to seek out interesting and important questions, collaborate with others to create new avenues to explore, rapidly discover and surpass the inevitable experimental challenges, and write up results with clarity and accuracy. His efforts on my behalf, and encouragement, have made my graduate education a powerful and transformative experience. A great thanks also to the other members of my thesis committee: Professor Karl Berggren, Professor Jing Kong, and Dr. Jeremy Sage for their input and helpful discussions.

Venturing into the world of experimental atomic physics put me on a steep learning curve, and the masterful guidance of my co-workers was essential to bringing me up to speed. Shannon Wang helped me get started in lab and showed me the ropes of the cryogenic ion trapping experiment. Peter Herskind introduced me to the lasers in lab and much of what I have learned about practical optics. Yufei Ge patiently taught me how to use the wirebonder (including making some MacGyver-worthy repairs), the scanning electron microscope, and the profilometer. Tony Kim shared some useful advice about cryogenic hardware, which saved me a fair bit of time in rewiring and making other modifications to the cryostat. I also benefitted greatly from knowledge shared through working with Michael Gutierrez, Molu Shi, and Hans Andersen in lab.

I am grateful to Karan Mehta, who introduced me to the world of CMOS manufacturing and integrated photonics (and mandolinist/composer Chris Thile), and who, along with Dr. Jeremy Sage, Dr. John Chiaverini, Dr. Colin Bruzewicz, and Professor Rajeev Ram, was essential to creation of the first, fully-functional ion trap fabricated using a standard CMOS process. The opportunity to collaborate with this group has been very rewarding for me. A great thanks also to Dr. Hyesung Park and Professor

Jing Kong for sharing their expertise and ingenuity so that a graphene-coated trap could become reality. I have learned a lot from our collaboration.

Rich Rines deserves a special mention for his contribution to the shuttling protocols (and corresponding 8-bit music) developed in this thesis, and long hours spent working out the SIMD QFT algorithm with me. He also added a lot of useful FPGA-based hardware to lab, in particular: the motor-controllers and photon counters used in this work. I am very grateful for his efforts, friendship, and hilarious late-night comments on g-chat.

Other members of the Quanta Lab deserve mention: Helena Zhang, who always has great coding tips (and isn't afraid to use her coding prowess for a wicked April Fool's prank), and Tailin Wu, who has helped with optics assembly. I am fortunate to have worked closely with a number of undergraduate students whose work has contributed to this thesis: Samuel Bader implemented a laser power stabilizer (and provided delicious laboratory supplies: Welch's gummies), Joaquin Maury helped to characterize electronics, Zachary Fisher wrote a lot of software for the experiment, Rena Katz built an injection locked laser, and Spencer Nichols developed control electronics. The previous members of the Polar Molecules experiment: Tony Zhu, David Meyer, Adam McCaughan, Arolyn Conwill, Anders Mortensen, Joan Smith, and Charles Herder, and the theorists in the group: Theodore Yoder and Guang Hao Low, also contributed many interesting ideas.

Outside of the group, I'm grateful for constructive discussions, and lots of loaned equipment (most of which was eventually returned, I swear!) from Professor Vladan Vuletić's ion trap team: Dorian Gangloff, Alexei Bylinskii, and Leon Karpa. Thanks to the many other members of the CUA who have also been generous with their time and equipment. I learned a lot from sharing an office with David Meyer and Adam McCaughan, along with some frequent visitors from the Berggren Lab: Faraz Najafi, and Andrew Dane. I should also acknowledge the iQuISE program and

members at large, especially Kristi Beck, which provided a broader perspective on quantum information science.

My housemates, both human: Tony Zhu and Greg Dooley, and feline: Batcat and Molly, added a lot of joy and color to my time as a graduate student. I only wish they hadn't kept me up quite so late with good conversation and incessant meowing (respectively).

Finally, my deepest thanks go to my cherished family: my parents Nagy Eltony and Laurie Douglas, and my sister Vanessa Eltony, for their love, support, and encouragement throughout every step in over two decades of schooling!



# Contents

<b>Cover page</b>	<b>1</b>
<b>Abstract</b>	<b>3</b>
<b>Acknowledgments</b>	<b>5</b>
<b>Contents</b>	<b>9</b>
<b>List of Figures</b>	<b>15</b>
<b>1 Trapped ion qubits</b>	<b>27</b>
1.1 Quantum computation . . . . .	28
1.1.1 A new computing paradigm . . . . .	28
1.1.2 Physical qubits . . . . .	29
1.2 Trapped ion qubits . . . . .	32
1.2.1 Demonstrations with ions . . . . .	32
1.2.2 Potential architectures . . . . .	34
1.2.3 Challenges to scaling-up . . . . .	35
1.3 This work . . . . .	37
1.3.1 Publications . . . . .	39
1.3.2 Contributions of co-workers . . . . .	40

<b>2</b>	<b>Introduction to ion trapping</b>	<b>41</b>
2.1	Spatial confinement . . . . .	41
2.1.1	Paul traps . . . . .	42
2.1.2	Surface-electrode traps . . . . .	45
2.2	Trapped ion qubit . . . . .	47
2.2.1	$^{88}\text{Sr}^+$ atomic ion . . . . .	47
2.2.2	Quantized ion motion . . . . .	48
2.2.3	Quantum state detection . . . . .	48
2.3	Basic laser-ion interactions . . . . .	50
2.3.1	Hamiltonian . . . . .	50
2.3.2	Sideband transitions . . . . .	51
2.4	Laser cooling . . . . .	52
2.4.1	Doppler cooling . . . . .	53
2.4.2	Sideband cooling . . . . .	55
2.4.3	Motional state detection . . . . .	56
<b>3</b>	<b>Experimental apparatus</b>	<b>59</b>
3.1	Planar ion traps . . . . .	59
3.1.1	The “Tower of London” trap geometry . . . . .	60
3.1.2	The “Quanta Classic” trap geometry . . . . .	60
3.1.3	Packaging traps . . . . .	61
3.2	Experimental chamber . . . . .	65
3.2.1	Bath cryostat . . . . .	65
3.2.2	Vacuum assembly . . . . .	66
3.3	Light delivery and collection . . . . .	67
3.3.1	Laser delivery optics . . . . .	67
3.3.2	Imaging . . . . .	69
3.4	Control electronics . . . . .	70

3.4.1	DC voltages: SuperDAC . . . . .	70
3.4.2	Trap RF voltage . . . . .	73
3.4.3	Data acquisition software . . . . .	74
3.5	Baseline measurements . . . . .	77
3.5.1	Micromotion amplitude . . . . .	77
3.5.2	Blueline spectrum . . . . .	77
3.5.3	Sideband spectrum . . . . .	78
3.5.4	Rabi oscillation . . . . .	78
<b>4</b>	<b>Transparent ion traps</b>	<b>81</b>
4.1	Fluorescence collection from ions . . . . .	81
4.2	New idea: transparent traps . . . . .	83
4.3	ITO trap fabrication . . . . .	84
4.4	Trap verification and proof-of-concept prototype . . . . .	85
4.5	Conclusion . . . . .	90
<b>5</b>	<b>Graphene-coated ion traps</b>	<b>93</b>
5.1	Motional heating of trapped ions . . . . .	93
5.2	Our approach: engineered surface . . . . .	95
5.2.1	Graphene as a protective coating for metals . . . . .	96
5.3	Trap design . . . . .	97
5.4	Trap fabrication . . . . .	98
5.5	Trap testing and heating rate measurements . . . . .	101
5.6	Results . . . . .	102
5.7	Discussion . . . . .	103
5.7.1	Surface roughening during CVD synthesis . . . . .	103
5.7.2	Related studies: modifying surface composition in situ . . . . .	104
5.8	Conclusion . . . . .	107

<b>6</b>	<b>Ion traps in CMOS</b>	<b>109</b>
6.1	CMOS technology for trapped-ion systems . . . . .	109
6.2	Trap design . . . . .	110
6.2.1	Edge-effects due to shortened RF electrode length . . . . .	112
6.2.2	Width of DC electrodes . . . . .	114
6.3	Trap fabrication . . . . .	114
6.4	Trap characterization . . . . .	116
6.4.1	Dielectric breakdown . . . . .	117
6.4.2	Laser-induced photo-effects . . . . .	117
6.4.3	Temperature-dependent nonlinearities . . . . .	118
6.4.4	Motional heating . . . . .	118
6.4.5	Excess micromotion . . . . .	118
6.4.6	Light scatter . . . . .	119
6.5	Conclusion . . . . .	120
<b>7</b>	<b>A SIMD quantum Fourier transform</b>	<b>121</b>
7.1	The discrete Fourier transform and classical parallelism . . . . .	122
7.1.1	Flynn’s taxonomy of parallelism . . . . .	122
7.1.2	The discrete Fourier transform . . . . .	123
7.1.3	The fast Fourier transform algorithm . . . . .	124
7.1.4	Hardware implementation . . . . .	124
7.2	The quantum Fourier transform . . . . .	125
7.2.1	Definition of the QFT . . . . .	125
7.2.2	Circuit for the QFT . . . . .	126
7.2.3	Approximate QFT . . . . .	127
7.2.4	Computing with the QFT . . . . .	128
7.2.5	Parallelizing the QFT . . . . .	129
7.2.6	Implementing the QFT with a single chain of ions . . . . .	130



7.3	Implementing the QFT in a scalable architecture . . . . .	132
7.3.1	Notation . . . . .	133
7.3.2	Shuttling overhead – performance bounds . . . . .	134
7.3.3	Simple QFT algorithm for $H_s = 1, V = 2$ . . . . .	135
7.4	A SIMD QFT algorithm . . . . .	139
7.4.1	SIMD algorithm for $H_s = 1, V = 2$ . . . . .	139
7.5	Conclusion . . . . .	144
<b>8</b>	<b>An integrated trap module</b>	<b>145</b>
8.1	Prototype system . . . . .	146
8.1.1	Capabilities of the Innsbruck experiment . . . . .	147
8.1.2	Laser requirements of the Innsbruck system . . . . .	147
8.2	Relevant requirements for the SIMD QFT . . . . .	150
8.2.1	Requirements for ion movement . . . . .	150
8.2.2	Requirements for qubit manipulation (gates) . . . . .	150
8.2.3	Requirements for quantum state detection . . . . .	151
8.2.4	Requirements for motional cooling . . . . .	151
8.3	Capabilities of existing hardware . . . . .	152
8.3.1	Ion movement . . . . .	152
8.3.2	Qubit manipulation – waveguides and couplers . . . . .	154
8.3.3	Qubit manipulation – modulators . . . . .	156
8.3.4	Quantum state detection . . . . .	158
8.3.5	Motional cooling . . . . .	161
8.4	Visualizing an integrated trap module . . . . .	162
8.5	Conclusion . . . . .	166
<b>9</b>	<b>Conclusion</b>	<b>167</b>

<b>A</b>	<b>Additional details about SuperDAC boards</b>	<b>171</b>
A.1	DAC board . . . . .	171
A.2	High-voltage amplifier board . . . . .	175
A.3	ADC board . . . . .	178
<b>B</b>	<b>Additional details about QFT algorithm</b>	<b>181</b>
B.1	Pulse sequences for QFT with constant circuit depth . . . . .	181
B.2	Details of simple QFT algorithm for $H_s = 1, V = 2$ . . . . .	182
	<b>Bibliography</b>	<b>187</b>

# List of Figures

1-1	The canonical four-rod Paul trap has been a workhorse for early demonstrations of quantum information processing with ions. (a) Schematic of a Paul trap consisting of four radiofrequency (RF) electrodes and two end-cap (DC) electrodes to confine ions in a linear chain. A laser beam is shown applying a gate pulse to a single ion. (b) Camera images of few ions in a Paul trap. The spacing between adjacent ions is 2 to 5 $\mu\text{m}$ . Images courtesy of University of Innsbruck. . . . .	33
1-2	Diagram of the quantum charge-coupled device (QCCD). Ions are stored in the memory region and moved to the interaction region for logic operations. Thin arrows show transport and confinement along the local trap axis. Reproduced from Ref. [KMW02]. . . . .	35
1-3	A notional quantum network composed of quantum nodes for processing and storing quantum states and quantum channels for distributing quantum information. Reproduced from Ref. [Kim08]. . . . .	36
2-1	(a) Electrode geometry of a generic 4-rod Paul trap (with the ion position indicated in blue), as discussed in the text. The voltages applied to diagonally-opposed rods are the same. (b) A cross-sectional view of the 4-rod Paul trap showing electric field lines and the coordinates used in the text. . . . .	42

2-2 (a) A generic surface-electrode ion trap electrode geometry. When  $V_{DC}^{(1)} = V_{DC}^{(3)} = V_{DC}^{(4)} = V_{DC}^{(6)}$  is positive, and  $V_{DC}^{(2)} = V_{DC}^{(5)}$  is negative, a symmetric trap is formed with the ion trapped above the surface at the center (indicated in blue). The trap dimensions,  $a, b, c$  are typically chosen to optimize properties such as the trap depth. (b) A cross-sectional view of the trap showing electric field lines and the coordinates used in the text. . . . . 46

2-3 Level structure of  $^{88}\text{Sr}^+$ , simplified to show only transitions relevant for cooling and state-detection. Air wavelengths are indicated for each transition, as well as the transition lifetime (shown in brackets). . . . 47

2-4 The tensor product of the atomic and motional states of the trapped ion result in a ladder of combined states. A transition which affects only the atomic state is known as a “carrier transition” (shown in black); a transition which increases the number of motional quanta by one is known as a “blue sideband transition” (shown in blue); and a transition which decreases the number of motional quanta by one is known as a “red sideband transition” (shown in red). . . . . 49

2-5 Relevant atomic transitions for quantum state detection using the electron shelving method. Light scattered (or not) on the  $5S_{1/2} \leftrightarrow 5P_{1/2}$  transition indicates the state of the qubit. . . . . 49

2-6 Relevant atomic transitions for Doppler cooling of  $^{88}\text{Sr}^+$ . During each cooling cycle on the  $5S_{1/2} \leftrightarrow 5P_{1/2}$  transition (at 422 nm), there is some probability ( $\approx 1/20$ ) of decay to the  $4D_{3/2}$  state, so this state is continuously repumped using a laser at 1091 nm. . . . . 53

2-7	This cartoon spectrum illustrates what we mean by the “resolved sideband regime”. Two peaks in the spectrum due to sideband transitions (corresponding to a trap mode with secular frequency $\omega_x$ ) are distinguishable from the peak due to the carrier transition (at frequency $\omega_0$ ). By adjusting the laser frequency, it is possible to address an individual sideband transition (not drawn to scale). . . . .	54
2-8	Illustration showing sideband cooling of a trapped ion. Red sideband pulses (shown in red) excite the atom to the upper state, while reducing the number of motional quanta by one. Then, spontaneous emission occurs (shown in black), which does not change the motional state. After repeated cycles, the ion is left in the motional ground state. . .	56
3-1	(a) Schematic showing the trap electrodes in the Tower of London geometry. The RF electrodes are shown in red, the ground electrodes are shown in blue, and the eight DC electrodes are shown in green. (b) The central trapping region of the Tower of London geometry with key dimensions indicated (for an ion height of about $150 \mu\text{m}$ ). Dimensions are shown net of gaps between electrodes (which are $15 \mu\text{m}$ ). The small marking visible in the outer ground electrode (in the bottom right of this image) is an alignment marking, used to align imaging optics to the trap during installation in the vacuum chamber. . . . .	61
3-2	Computed pseudopotential in the $\hat{x} - \hat{y}$ plane at the trap center. An ion would ideally be trapped at $(x, y, z) = (0, 150, 0) \mu\text{m}$ . . . . .	63

3-3 (a) Schematic showing the trap electrodes in the Quanta Classic geometry. The RF electrode is shown in red, the ground electrodes are shown in blue, and the DC electrodes are shown in green. The labels LM, RM, BE, FE are the names given to the four DC electrodes. (b) The central trapping region of the Quanta Classic geometry with key dimensions indicated (for an ion height of about  $100 \mu\text{m}$ ). Dimensions are shown net of gaps between electrodes (which are  $10 \mu\text{m}$ ). . . . . 63

3-4 A surface-electrode trap packaged in a CPGA and installed in the cryostat. A copper tab connected to a copper spacer sitting below the trap can be seen, which provides a thermal contact to the 4 K baseplate. Bypass capacitors are also visible which are connected to each of the 8 DC electrodes with wirebonds. . . . . 64

3-5 (a) Schematic of the bath cryostat, with liquid helium (LHe) and liquid nitrogen ( $\text{LN}_2$ ) tanks labeled. The 77 K shield attached to the  $\text{LN}_2$  tank encloses the LHe tank and the 4 K baseplate. The working area (where the trap is mounted) is built around the 4 K baseplate, directly below the LHe tank. (b) Photograph of the 4 K baseplate with major components labeled. . . . . 66

3-6 Diagram showing the vacuum apparatus for the cryostat (all the indicated valves are gate valves). . . . . 67

- 3-7 Diagram showing optics for laser beam delivery to the cryostat. The orientation of the trap axis within the chamber is indicated by a green line. A single Thorlabs achromatic doublet pair (FAC1025-A) is used to focus the two Doppler cooling beams (at 422 nm) to a waist at the position of the ion. A half-wave plate (HWP) is used to adjust the power ratio between the two beam paths, and a single-axis translation stage in the 422-only path (the photoionization light is filtered out) is used to assure the path lengths are equal. Additional half-wave plates in the other 422 nm beam path and the 674 nm beam path assure that the correct polarization is set. Another achromat (FAC1025-A) is used to focus the 674 nm beam to a waist at the trap. The 1033/1092 nm beam is reduced in size and collimated using two lenses with focal lengths of 45 and 500 mm, respectively, while the 1064 nm beam is unaltered from its free-space, collimated waist size. Each beam enters the cryostat via a periscope adjustable using two motor-controlled stages. 68
- 3-8 Photographs of a SuperDAC unit with 8 outputs. (a) The front of the box, showing the BNC connections for power ( $\pm 25$  V), the BNC connections for the outputs (1 to 8) and the USB connection for computer control. (b) The inside of the box, showing the top of the two 3-PCB stacks (the amplifier board is on top), the regulator board (which is used to convert the  $\pm 25$  V input power to the various voltages required by the other boards, and the Arduino MEGA microcontroller. . . . . 72
- 3-9 Screenshot showing the user-interface for SuperDAC. For each channel, the left box shows the ADC reading, and the right box allows the user to set the voltage (the hardware had not been calibrated yet, which is why the ADC readings are not as close as expected). . . . . 73

3-10	Noise spectrum of SuperDAC unit (channel 1); the noise floor is shown in blue, and the signal is shown in red. The resolution bandwidth was set to 3 kHz and the input impedance was set to 50 $\Omega$ . . . . .	76
3-11	A typical scan obtained while compensating micromotion in the $\hat{x}$ -direction. The micromotion amplitude is normalized to the total number of photon counts. . . . .	78
3-12	Spectrum of the $5S_{1/2} \leftrightarrow 5P_{1/2}$ transition (at 422 nm). When the frequency offset from the atomic resonance becomes positive, the ion is heated by the laser. The data points (shown in black) are fitted to a Lorentzian (shown in red), which has a linewidth (full width at half maximum) of 42 MHz. . . . .	79
3-13	Sideband spectrum of the qubit transition, $5S_{1/2} \leftrightarrow 4D_{5/2}$ (at 674 nm) following sideband cooling of all three modes close to the ground state. The lowest frequency sidebands correspond to the axial mode, and the two higher frequency sidebands correspond to the radial modes. Because the tilt of the trap axes is small, there is only a small projection of the laser beam onto the highest frequency radial mode, which is why this line is much narrower. . . . .	80
3-14	Rabi oscillations on the qubit (carrier) transition, $5S_{1/2} \leftrightarrow 4D_{5/2}$ (at 674 nm), following sideband cooling. Each data point (shown in black) represents 100 experiments, and is fitted to a sinusoid (shown in red), which has a Pi time of 5.3 $\mu$ s and a contrast of 99%. . . . .	80
4-1	Diagram outlining the conventional approach to fluorescence collection, which places a high numerical aperture objective near to the ion, and detects light emitted with a photomultiplier or charge-coupled detector located outside of the vacuum chamber. Dimensions shown are typical for experiments with surface traps. . . . .	82



4-2	Conceptual drawing showing how an ion trap with transparent electrodes would allow fluorescence typically obstructed by the trap to be transmitted to a photodetector below. For the (typical) dimensions shown, the solid angle capture possible approaches 50%. . . . .	83
4-3	(a) ITO-4K trap mounted in a CPGA; 50 nm of Au is visible on the RF electrodes and contact pads for wire bonding. (b) ITO-PD trap mounted on photodiode in a CGPA; with only 5 nm of Au on the RF electrodes, the photodiode is visible through the trap. (c) Diagram of trap geometry showing RF electrodes (red), ground electrodes (blue), and DC electrodes (green). . . . .	85
4-4	Apparatus for detection of ion fluorescence through a transparent trap using a photodiode mounted below. The acousto-optic modulator (AOM) is used to modulate the repumper for lock-in detection. . . .	86
4-5	Schematic for the preamplifier circuit. . . . .	87
4-6	(a) The PCB artwork for the preamplifier circuit. The top layer routing is shown in red, the bottom layer routing is shown in blue, and the drill locations are shown in green. (b) The populated preamplifier circuit board before installation in the cryostat. . . . .	88
4-7	(a) Photograph showing how the preamplifier circuit is mounted to the 77 K radiation shield. The top of the chip packages make thermal contact to the copper mount. (b) Close-up view showing how the preamplifier circuit board is attached. . . . .	89
4-8	(a) Photodiode voltage and photomultiplier count rate, both background subtracted, during loading of an ion cloud. Each point is averaged over 30 seconds. (b) Histogram of photodiode voltages over a period of several minutes without ions (red), and after loading an ion cloud (blue). . . . .	90

4-9	Diagram of proposed compact entanglement unit (see text). . . . .	90
5-1	(a) Diagram of trap geometry showing RF electrodes (red), ground electrodes (blue), and DC electrodes (green). A single ion is confined 100 $\mu\text{m}$ above the center of the ground electrode adjacent to the notch in the RF electrode. (b) Photograph showing one of the graphene-coated ion traps tested (MLG-1a) mounted in a CPGA (all traps tested are from a single fabrication batch and therefore look very similar). . . . .	98
5-2	Raman spectra of a copper trap following graphene synthesis (with an integration time of 30 seconds). The broad hump is the Cu background and the indicated peaks correspond to graphene. The small G/2D ratio and small D peak suggest that the graphene is primarily monolayer. . . . .	100
5-3	(a) A single, trapped $^{88}\text{Sr}^+$ ion above MLG-c as seen on the experiment CCD camera (false colors). (b) Shelving rate (averaged over 20 measurements, each with 100 trials per point) on the red sideband (shown in red) and the blue sideband (shown in blue). (c) Plot showing the inferred average phonon number $\bar{n}$ as a function of time. The linear fit indicates a heating rate of $\dot{\bar{n}} = 1020 \pm 30$ quanta/s. . . . .	103
5-4	(a) Profilometer scan across 20 $\mu\text{m}$ in a central area of an evaporated copper trap. (b) Profilometer scan across 20 $\mu\text{m}$ of a trap from the same batch following graphene synthesis (notice the axis scale has changed). (c) Profilometer scan across 400 $\mu\text{m}$ of the same graphene-coated trap. . . . .	106
6-1	Trap mask. The mask is 3 mm long, and 1.5 mm wide. . . . .	111
6-2	Estimated RF field along the trap axis due to edge-effects ( $\mathbf{E}_{RF}(z) \cdot \hat{\mathbf{z}}$ ). . . . .	113
6-3	Estimated excess micromotion due to edge-effects along the trap axis. . . . .	113
6-4	Angle subtended (in the x-direction) by the electrode from the ion as a function of electrode width. . . . .	114

6-5	Die and process cross-section. A micrograph of the fabricated $3 \times 3$ mm <sup>2</sup> die is shown in the upper left panel. The lower left panel shows a perspective rendering of the top aluminum trap layer and the meshed ground plane in copper below, as designed; the gaps in the trap electrodes here are 5 $\mu$ m, and the ground mesh is formed of 600 nm wires with 350 nm gaps along $x$ and 10 $\mu$ m gaps along $y$ . A chip cross section is diagrammed at right, with approximate relevant dimensions labeled (“pSi” is polysilicon and metal interconnect layers are labelled m1 through m8). Vias shown between metal layers are only representative. . . . .	115
6-6	Micrograph of trap chip diced from die and mounted on the sapphire interposer of a cryogenic vacuum system. Aluminum wirebonds are used to make contact from the aluminum trap electrodes to the gold interposer leads. The chip is 2.5 mm long and 1.2 mm wide. The inset shows two ions trapped 50 $\mu$ m from the surface of the trap chip. The ions are approximately 5 $\mu$ m apart. . . . .	116
6-7	Representative measurement of heating rate in a CMOS-foundry-fabricated ion trap. Average occupation of the axial mode of vibration in the linear trap is plotted as a function of delay time after preparation in the ground state at a trap temperature of 8.4 K. A linear fit (line shown) gives a heating rate for these data of 80(5) quanta/s where the uncertainty is due to statistical errors propagated through the fit. An average of five such measurements gives a heating rate of 81(9) quanta/s where the uncertainty is due to run-to-run variability. The inset shows all five measurements with a line indicating the weighted average value. . . . .	119
7-1	Circuit for a 3-qubit QFT, up to a permutation of the output qubits.	127

7-2	A circuit for the QFT, up to reversing the order of the output qubits. Note that in this figure, the notation $R_k$ is used to refer to $R_\theta$ with $\theta = \frac{2\pi}{2^k}$ as defined in Equation 7.8. . . . .	128
7-3	Circuit for the 5-qubit QFT (Hadamard gates not shown) with the “staircases” through which gates are not commutable (see description in text) indicated in blue. In this diagram, the label $\pi/2$ indicates a phase rotation gate $R_{\pi/2}$ as defined in Section 7.2.2. . . . .	130
7-4	Schematic showing the general distributed-ion-chain layout we consider in this chapter. It is a 2D array of ion chains of height $H$ and overall width $W$ . Each of the grey rectangles labelled 1 through $W$ correspond to a single trap site (electrodes not shown) with $H$ ions shown in blue; connection between chains is also not shown. . . . .	131
7-5	The first six shuttling steps for a SIMD QFT with $H = 3$ , $W = 5$ , $H_s = 1$ and $V = 2$ (see text for details). Rotation gates and in-chain <i>SWAP</i> steps are not shown, but progress through the QFT circuit is indicated schematically on the right as completed gates turn from white to grey. Input (binary) ions are shown in blue, and transformed (Fourier) ions are shown in green. Permutation $P_0$ is shown in pink (applied to chains 0, 1, 2, and 3 simultaneously), and permutation $P_1$ is shown in orange (applied to chains 1, 2, 3, and 4 simultaneously). In this simple case (of small $W$ ), parallelism is apparent in the four ion movements occurring in each shuttling step. . . . .	143
8-1	Level scheme for $^{40}\text{Ca}^+$ (only relevant transitions are shown). . . . .	149

8-2	(a) SEM image showing a vertical grating coupler designed by Orcutt et al. [OKH <sup>+</sup> 11]. (b) TEM image showing the cross-section of a CMOS-integrated waveguide. The waveguide has a $670 \times 80$ nm polysilicon core, clad with a conformal 50 nm silicon nitride liner, surrounded by oxide. (Images reproduced from Ref. [OKH <sup>+</sup> 11]). . . . .	155
8-3	Diagram showing the silicon modulator developed by Liu et al. [LLR <sup>+</sup> 07]. It is a Mach-Zehnder interferometer containing two pn junction based phase shifters. The waveguide splitter is a $1 \times 2$ multi-mode interference (MMI) coupler. (Image reproduced from Ref. [LLR <sup>+</sup> 07]). . . . .	158
8-4	(a) Micrograph of SPAD structure fabricated by Field et al. [FLC <sup>+</sup> 10]. (b) Illustration of the expected pn diode structure after fabrication of the SPAD. (Images reproduced from Ref. [FLC <sup>+</sup> 10]). . . . .	159
8-5	Illustration showing the scheme proposed by Streed et al. [SNJ <sup>+</sup> 11] for highly parallel readout of trapped-ion qubits. Fluorescence from ions trapped at multiple sites on a surface trap is efficiently coupled into single optical modes by an array of microfabricated phase Fresnel lenses on a single substrate. (Image reproduced from Ref. [SNJ <sup>+</sup> 11]).	160
8-6	Optical micrograph of 10 waveguide-integrated detectors (D1 – D10) assembled on the same photonic chip by Najafi et al. [NMH <sup>+</sup> 15]. The waveguides are marked by red arrows. The length of the scale bar (shown in blue) is $100 \mu\text{m}$ . (Image reproduced from Ref. [NMH <sup>+</sup> 15]).	161

8-7	Top view of a hypothetical trap zone for 5 ions. Ions would be trapped $\sim 50 \mu\text{m}$ above the in-coupler sites, with an inter-ion spacing of $\sim 5 \mu\text{m}$ . Integrated optical components are shown in place below trap electrodes (see text for details). Picture is not drawn to scale, and square bends (not practical for waveguides) are used only for clarity. In a real system, the width of the central ground electrode might be $\sim 50 \mu\text{m}$ and the grating couplers and EO modulators can be expected to have a scale of $\sim 10 \mu\text{m}$ . . . . .	164
8-8	Conceptual drawing showing cross-sectional view of a hypothetical trap zone (not drawn to scale). A $^{88}\text{Sr}^+$ ion is confined about $50 \mu\text{m}$ above the trap surface. Light emitted by two CMOS-compatible grating couplers comes to a focus at the ion position, delivering light for Doppler cooling and qubit control. Integrated waveguides carry light to these focusing grating couplers, and away from an in-coupler placed below the ion to collect fluorescence. Modulators are also integrated to direct light and apply pulses. Beneath a ground plane formed in a buried metal layer, additional metal layers are used to bias optical devices, and supply control voltages to the trap DC electrodes. . . . .	165
A-1	Schematic for the DAC board (see text for details). . . . .	173
A-2	Board layout for the DAC board (see text for details). . . . .	174
A-3	Schematic for the high-voltage amplifier board (see text for details). . . . .	176
A-4	Board layout for the high-voltage amplifier board (see text for details). . . . .	177
A-5	Schematic for the ADC board (see text for details). . . . .	179
A-6	Board layout for the ADC board (see text for details). . . . .	180
B-1	A set of simultaneous controlled rotations with $H = 4$ . . . . .	182

# Chapter 1

## Trapped ion qubits

We are living in the Information Age, a time when lives and economies are driven by computerized information. A single unit of information is a binary state: “true” or “false”, “0” or “1”, known as a binary digit, or “bit”. Physically, a bit is captured by the state of a two-state device, most often a transistor.

The Digital Revolution has brought us exponential growth in computing power, with transistors—the basic building blocks of digital technology—shrinking towards the nanoscale. What happens when these basic building blocks have miniaturized to just one atom per bit? At this scale, the bits behave according to quantum mechanics, and a new model of computation, “quantum computation”, becomes possible, expanding capabilities for simulation and for new, potentially more efficient algorithms.

This thesis documents the development of technology to enable a large-scale quantum computer, focusing on the electronic state of a trapped atomic ion as a basic unit of information. This first chapter provides some context for the work presented in the thesis. First, the basic idea of quantum computation is introduced in Section 1.1, along with a description of some of the different physical systems currently under study. Next, experimental progress with trapped ion qubits is described, and poten-

tial architectures for building up a large-scale computer are discussed in Section 1.2. Finally, a summary of the results documented in this thesis is presented in Section 1.3.

## 1.1 Quantum computation

A quantum computer is a machine that performs calculations through the control and manipulation of quantum objects and their quantum dynamics. By harnessing unique properties of quantum mechanics such as entanglement and superposition, a quantum computer can process information in a profoundly different way than a familiar “classical” computer. For certain problems, like factoring large numbers into primes [Sho94] or searching large databases [Gro96], quantum algorithms exist which are asymptotically more efficient than the best known classical algorithms for the same problems.

### 1.1.1 A new computing paradigm

A classical bit of information exists in one of two possible states. A quantum bit of information (known as a “qubit”) is a quantum mechanical two-level system; it exists in a two-dimensional, complex Hilbert space, in a state represented by a unit vector. Typically, the qubit levels are denoted  $|0\rangle$  and  $|1\rangle$ , so the state of a qubit can be represented as  $\alpha_1|0\rangle + \alpha_2|1\rangle$  (where  $\alpha_1, \alpha_2$  are complex numbers satisfying  $|\alpha_1|^2 + |\alpha_2|^2 = 1$ ). A quantum gate, like a classical logic gate, can act on one or more qubits, and is represented by a unitary matrix. Unlike a classical bit, a qubit can exist in a superposition state ( $\alpha_{1,2} \neq 0$ ); and as quantum mechanical objects, multiple qubits can form entangled states.  $n$  qubits combine (via a tensor product) to form a Hilbert space of dimension  $2^n$ . This exponential growth in the size of the system description is required to capture all the potential correlations between the



amplitudes and phases of the different qubits.

The unique properties of qubits (superposition, entanglement) have inspired entirely new kinds of algorithms (quantum algorithms), which take advantage of these properties to solve problems more efficiently. This began in 1981, with Feynman's suggestion that a quantum machine might be more effective than a classical computer at simulating quantum physics [Fey82]. Feynman, and later Deutsch [Deu85], realized that a "quantum computer" could potentially surpass a classical computer at solving certain problems [DJ92]. In 1994, a quantum algorithm to factor large numbers into primes, known as Shor's factoring algorithm [Sho94], was introduced. This algorithm outperforms the best known classical algorithm for factorization with an exponential improvement in speed. Prime-factorization is important for cryptography; its difficulty (for a classical computer) is responsible for the security of many current encryption schemes [Bru96]. Other quantum algorithms providing a speed up over the best known classical solution include Grover's algorithm for unstructured search [Gro96], and a recent algorithm for solving linear systems of equations [HHL09].

### 1.1.2 Physical qubits

Modern computers based on integrated circuits have transistor counts in the billions. The greatest transistor count in a commercially-available CPU chip is over 5.5 billion (the Intel 18-core Xeon Haswell-EP). To build a quantum computer capable of tackling intractable problems for a classical machine, the qubit implementation chosen needs to be scalable to thousands of qubits at least [Sho95, Ste02], with each qubit manipulable and measurable with sufficient fidelity. The physical system must also be designed such that interactions with the environment do not erode the quantum information stored in the qubits before computation is complete. The essential ingredients for a quantum computer are expressed succinctly in DiVincenzo's criteria [DiV00]; they are:

1. A well-defined, scalable qubit array
2. The ability to initialize the qubits to a well-defined state (such as the  $|000\dots\rangle$  state, where all qubits are in the  $|0\rangle$  state)
3. A universal set of quantum gates (meaning the available gate operations on the qubits can be combined to yield any arbitrary quantum operation)
4. The ability to “readout” (measure) the state of each qubit
5. Coherence times much longer than the gate time (the coherence time is the delay before the quantum state of a qubit is lost due to interaction with its environment)

For a classical computer, the concept of a Turing machine [Tur37] can be invoked to determine an analogous list of requirements. A Turing machine is a hypothetical device, made up of four components:

1. A “tape” made up of cells, each containing a symbol from some finite alphabet
2. A “head” that can read/write symbols on the tape, and move the tape (or itself) cell by cell.
3. A “state register” that stores the state of the Turing machine
4. A finite table of instructions that control the action of the head

Any classical computer algorithm can in principle be simulated on a Turing machine. The elements of a Turing machine can be recast to be parallel to DiVincenzo’s criteria, creating an analogous list of requirements for a general-purpose classical computer:

1. A physical bit implementation scalable to a larger architecture
2. The ability to take in instructions/information from the outside world

3. A universal set of logic gates (such as {NAND, FANOUT})
4. The ability to transmit information to the outside world
5. The ability to store information (memory)

Viewed this way, it is clear that the requirements for a quantum computer are similar to those for a classical computer, with the exception of details related to quantum effects.

In essence, a qubit can be any quantum two-level system or two-level subspace of a system. Early implementations used nuclear magnetic resonance (NMR), in which the nuclear spins of molecules in liquid solution form the qubits, controlled using radio frequency (RF) pulses [VSB<sup>+</sup>00]. Other early demonstrations employed optical photons, in which the polarizations or spatial modes of the photons form the qubits, controlled using optical elements (such as beamsplitters or phase shifters) [PMO09]. However, neither of these implementations show promise for scaling-up to a large computer. In NMR quantum computing, the ability to prepare the qubits in a well-defined initial state decreases exponentially with the number of qubits, and in photon-based quantum computing, it is prohibitively difficult to interact the qubits (photons), precluding a universal gate set.

Quantum computers based on solid-state technologies, in which the qubits are built from physical systems such as superconducting circuits [RDN<sup>+</sup>12] and semiconductor quantum dots [PJT<sup>+</sup>05], are being widely explored. The ability to leverage fabrication techniques used for classical computing technology makes these systems nominally scalable, but imperfections in fabrication make repeatability a challenge. Because they are embedded in solid-state systems, the qubits are also difficult to isolate from the environment, limiting their coherence times (although this has improved dramatically for superconducting qubits over the past few years) [DS13]. So far, quantum computing systems based on trapped ions, in which the qubits are either

the electronic or the nuclear spin states of atomic ions, have made the most progress towards satisfying the DiVincenzo criteria.

## 1.2 Trapped ion qubits

The electronic or nuclear spin states of trapped atomic ions make for good qubits because of their long coherence times (ranging from ms to s) relative to the gate times (typically  $\mu\text{s}$  or less), strong inter-ion interactions (via the Coulomb force), long lifetimes (single ions can remain trapped for hours or days), and natural reproducibility (for example: to the best of our knowledge, all  $^{88}\text{Sr}^+$  ions are identical).


### 1.2.1 Demonstrations with ions

Many of the essential ingredients for quantum computing have been demonstrated in toy systems of several trapped ions. In pioneering experiments [WMI<sup>+</sup>97], a small qubit register was formed by a chain of ions confined within a trap of oscillating electromagnetic fields (called a Paul trap), constructed from precisely-machined metal rods or blades in ultrahigh vacuum (see Figure 1-1). The ions' internal (electronic) and external (motional) states were manipulated using laser or microwave pulses to carry out quantum gates [NC02], which were then read-out via the fluorescence emitted by the ions. With this hardware, quantum protocols such as teleportation [BCS<sup>+</sup>04, RCB<sup>+</sup>07, OMM<sup>+</sup>09], state mapping from an ion to a photon [SCB<sup>+</sup>13], and quantum simulation [KCK<sup>+</sup>10, BMS<sup>+</sup>11, LHN<sup>+</sup>11, GLK<sup>+</sup>11, BSK<sup>+</sup>12, BR12, ISC<sup>+</sup>13] have been realized.

A complete methods set for scalable quantum computing with ions has been described [HHJ<sup>+</sup>09] and toy computers based on trapped ion qubits have provided proof-of-principle demonstrations of the essential ingredients [Hug01], including: high-fidelity quantum state preparation [RZR<sup>+</sup>99, MMK<sup>+</sup>95b, KWM<sup>+</sup>98], single qubit ro-

tations [RZR<sup>+</sup>99, NLR<sup>+</sup>99], multi-qubit operations [MMK<sup>+</sup>95a, SKK<sup>+</sup>00, LDM<sup>+</sup>03, SKHR<sup>+</sup>03, MKH<sup>+</sup>09b], and state read-out [RZR<sup>+</sup>99, RKM<sup>+</sup>01, MSW<sup>+</sup>08]. Additionally, a number of quantum algorithms including the quantum Fourier transform [CBL<sup>+</sup>05, SNM<sup>+</sup>13], the Deutsch-Josza algorithm [GRL<sup>+</sup>03], Grover search [BHL<sup>+</sup>05], and quantum error correction [SBM<sup>+</sup>11] have been realized.

Yet, it is a formidable technical challenge to extend these results much beyond the current experimental record of 14 qubits [MSB<sup>+</sup>11]. In these demonstrations, the qubits are chains of ions trapped in machined 4-rod style Paul traps (as shown in Figure 1-1), which are likely limited to computations with 10s of ions [WMI<sup>+</sup>97]. Designing a massive, versatile ion system that can retain coherence through a quantum computation necessitates exploiting new manufacturing platforms and devices, and understanding the microscopic behaviors of different trap materials. Hence, the challenge faced by researchers in recent years has been in crafting cohesive ion trap *systems*: combining diverse parts into a more powerful whole which achieves the ultimate performance from each component simultaneously.




**Figure 1-1:** The canonical four-rod Paul trap has been a workhorse for early demonstrations of quantum information processing with ions. (a) Schematic of a Paul trap consisting of four radiofrequency (RF) electrodes and two end-cap (DC) electrodes to confine ions in a linear chain. A laser beam is shown applying a gate pulse to a single ion. (b) Camera images of few ions in a Paul trap. The spacing between adjacent ions is 2 to 5  $\mu\text{m}$ . Images courtesy of University of Innsbruck.

## 1.2.2 Potential architectures

An important development towards scaling-up ion systems was the invention of the surface-electrode ion trap (or “surface trap”) [CBB<sup>+</sup>05], in which a single ion or chain of ions can be confined 20  $\mu\text{m}$  to 2 mm above a trap chip [SCR<sup>+</sup>06]. Surface traps can be constructed using standard printed circuit board manufacturing [BCL<sup>+</sup>07], or microfabrication techniques [SHO<sup>+</sup>06], making them far simpler to build than their precisely-machined, three-dimensional counterparts. As such, surface traps are more amenable to being shrunk in size and replicated to create an array of trapping sites. This compactness comes at the cost of shallower trap depths and greater anharmonicity of the trapping potential. Nonetheless, this convenient platform has allowed researchers to incorporate a variety of useful devices into ion traps.

Building on surface trap technology, a scalable quantum computer could potentially be constructed by interconnecting a large number of traps—which are individually limited to computations with 10s of ions [WMI<sup>+</sup>97]—to form a densely-populated chip, or “quantum charge-coupled device” (QCCD) [KMW02], capable of computations with 100s or 1000s of ions. In such an architecture, there are a number of modular “trap zones” in which a small number of interacting ions are confined, and gates can be performed between all possible ions by shuttling ions between the different zones to mediate interactions, as pictured in Figure 1-2. Because of its array-like structure, this architecture is naturally suited to parallel operation.

An alternative proposal seeks to connect a large number of trapped ions by forming a “quantum network”, composed of quantum nodes for processing and storing quantum states connected by quantum channels for distributing quantum information [CZKM96, MRR<sup>+</sup>14]. This is illustrated conceptually in Figure 1-3. In such an architecture, the quantum state of one ion qubit can be converted to the quantum state of a photon that will carry this information through space to another trapped ion qubit. Coupling distant ions in a quantum network can be accomplished through




**Figure 1-2:** Diagram of the quantum charge-coupled device (QCCD). Ions are stored in the memory region and moved to the interaction region for logic operations. Thin arrows show transport and confinement along the local trap axis. Reproduced from Ref. [KMW02].

strong coupling of single photons and ions in the setting of cavity quantum electrodynamics (QED) [Kim08] or through a probabilistic ion photon mapping (not requiring strong coupling) [DBMM04]. This architecture has the advantage of being able to transmit quantum information between distant computers as well as between nodes in a single machine.

### 1.2.3 Challenges to scaling-up

Creating a scalable ion system for quantum information processing is a formidable technical challenge. There are many intricacies to address: incorporating devices for qubit manipulation and readout; obtaining precise control over errors and noise; and



**Figure 1-3:** A notional quantum network composed of quantum nodes for processing and storing quantum states and quantum channels for distributing quantum information. Reproduced from Ref. [Kim08].

efficiently mapping quantum algorithms onto practical hardware.

One important aspect of a cohesive system is optical read-out of many trapped ion qubits. The efficiency of fluorescence collection from an ion sets the rate of quantum state detection of the qubit. Typically, fluorescence is captured using a bulk objective close to the ion and a photomultiplier or charge-coupled detector located outside the vacuum chamber, which results in collection from a small solid angle ( $\sim 5\%$ ) and is not scalable to many ions. Read-out of dense arrays of trapped ions will require efficient light collection from many ions in parallel. For a distributed architecture, with remote atomic nodes connected by photons, light collection is also a critical factor determining the efficiency of remote entanglement generation [LHM<sup>+</sup>09].

A second challenge to scaling-up is the presence of increasing electric field noise as ions are brought closer to the electrode surface, which limits gate fidelity and coherence time [TKK<sup>+</sup>00]. Electric field noise couples to the trapped motion of the ion, resulting in heating of the motional state. An efficient method of performing gates on multiple ions in the same trap is to use their shared motional state as a



“quantum bus”, but if the motional state is altered by noise during the gate, this limits the fidelity. The physical origin of this noise, and effective ways to mitigate it, are yet to be understood.

Another line of inquiry asks how the performance and scale of technologies employed by conventional “classical” computers can be leveraged for quantum information processing with trapped ions. Repurposing existing computer hardware for ion systems requires trap designs to incorporate new materials and size constraints. Additionally, the computational power of a trapped ion system is dependent on how quantum algorithms are actually implemented within its hardware. Designing optimal trap geometries and orchestrating how ions will be moved among different zones during computation poses a significant challenge.

## 1.3 This work

This thesis explores how we can engineer better ion traps towards a scalable quantum computer architecture. Our work is concerned with creating a complete, cohesive ion system through addressing several important challenges:

1. Achieving efficient fluorescence collection from many ions in parallel;
2. Understanding the mechanisms responsible for motional heating in ion traps, so that excessive heating can be mitigated;
3. Harnessing the capabilities of the CMOS platform for quantum information processing with trapped ions;
4. Establishing techniques to efficiently map quantum algorithms onto a planar trap array, like the QCCD;
5. Analyzing and proposing an integrated trap module capable of serving as a building block for a larger quantum system.

Our results are described in Chapters 4 through 8, and are briefly summarized below:

- **Chapter 4**

We demonstrate a novel ion trap with transparent electrodes, which allows fluorescence to be transmitted through the trap. As a proof-of-concept prototype, we sandwich a commercial photodiode with a transparent trap, and detect the fluorescence emitted by trapped ions through the electrodes. This work illustrates one possibility for efficient, parallelized fluorescence collection.

- **Chapter 5**

To better understand the causes of motional heating, we create a trap with a qualitatively different surface. We implement an ion trap passivated with a monolayer of graphene, which we hypothesize will result in lower motional heating because graphene will protect the trap electrodes from oxidation and other contamination. Surprisingly, the heating rate we measure in the graphene-coated trap is significantly higher than would be expected in a comparable, uncoated trap.

- **Chapter 6**

We design an ion trap suitable for fabrication using a standard CMOS foundry process. Experimental testing reveals that the trap functions favorably, with a measured heating rate similar to that seen in other surface-electrode traps of similar size.

- **Chapter 7**

We show that many quantum algorithms (such as Shor's factoring algorithm) can be realized efficiently using the quantum Fourier transform (QFT) as an algorithmic primitive. Further, we develop a protocol for implementing the QFT in a planar trap array (made up of multiple trap zones) that is SIMD

(Single Instruction Multiple Data). The SIMD algorithm greatly simplifies the hardware required for ion movement.

- **Chapter 8**

Lastly, we combine the results of Chapters 6 and 7 to envision a fully-integrated trap module capable of implementing the SIMD QFT introduced in Chapter 7. We examine the requirements for such a module, and explore how CMOS-compatible devices could be integrated to create a scalable building block.

These investigations focus on microfabricated, surface-electrode ion traps experimentally verified in a bath cryostat at cryogenic temperatures (4 K and 77 K), where low pressure is achievable within a short time frame, and heating of the ion's motional state is suppressed [LGA<sup>+</sup>08].

### 1.3.1 Publications

Some of the work presented in this thesis has been published in the following articles:

- **A. M. Eltony**, D. Gangloff, M. Shi, A. Bylinskii, V. Vuletić, and I. L. Chuang. Technologies for trapped-ion quantum information systems. *To appear in a Springer Special Issue on Trapped Ion Quantum Information Processing*.
- **A. M. Eltony**, H. G. Park, S. X. Wang, J. Kong, and I. L. Chuang. Motional Heating in a Graphene-Coated Ion Trap. *Nano Letters* 14, 5712 - 5716 (2014).
- K. K. Mehta, **A. M. Eltony**, C. D. Bruzewicz, I. L. Chuang, R. J. Ram, J. M. Sage, and J. Chiaverini. Ion Traps Fabricated in a CMOS Foundry. *Applied Physics Letters* 105, 044103 (2014).
- **A. M. Eltony**, S. X. Wang, G. M. Akselrod, P. F. Herskind, and I. L. Chuang. Transparent ion trap with integrated photodetector. *Applied Physics Letters* 102, 054106 (2013).

### 1.3.2 Contributions of co-workers

The work presented in this thesis is stronger for having been undertaken collaboratively. The contributions my co-workers made to this thesis are summarized below.

The cryogenic ion trapping experiment has developed through several generations of students, starting with the original builders: Waseem Bakr and Paul Antohi, followed by Jaroslaw Labaziewicz and Shannon Wang. Likewise, the laser systems originally constructed by Jaroslaw Labaziewicz and Ruth Shewmon have continued to evolve through work by Jeffrey Russom, Peter Herskind, and Michael Gutierrez. Many co-workers have contributed to creating the ion traps tested: Shannon Wang and Gleb Akselrod fabricated the indium-tin oxide (ITO) ion traps, David Meyer and Molu Shi fabricated niobium and gold traps, and Yufei Ge fabricated the copper traps used for graphene synthesis.

Shannon Wang suggested developing an ion trap with indium-tin oxide electrodes, and Peter Herskind conceived of the compact entanglement unit introduced in Chapter 4. Hyesung Park developed the graphene growth process described in Chapter 5, and synthesized graphene on each of the traps tested. Karan Mehta adapted my trap layout to conform to CMOS design rules and created the final mask presented in Chapter 6. Colin Bruzewicz, Jeremy Sage, and John Chiaverini characterized the resulting ion traps in a cryogenic experiment they have developed at Lincoln Laboratories, as discussed in Chapter 6. Rich Rines created the shuttling patterns, and contributed significantly to the QFT algorithm developed in Chapter 7. Discussions with Colin Bruzewicz, Jeremy Sage, John Chiaverini, Rajeev Ram, and especially Karan Mehta provided inspiration for the outlook presented in Chapter 8.

Finally, I gratefully acknowledge support from the MQCO Program with funding from IARPA, and the National Science and Engineering Research Council of Canada's Postgraduate Scholarship program.

# Chapter 2

## Introduction to ion trapping


This chapter introduces the fundamental mechanisms and interactions required to turn an atomic ion into a qubit. First, the basics of how an ion can be confined in space using electromagnetic fields is introduced in Section 2.1. Next, the internal structure of the  $^{88}\text{Sr}^+$  trapped ion qubit is described in Section 2.2. Following this, we explore control of the ion's internal degrees of freedom using lasers in Section 2.3. Finally, we describe laser cooling techniques used to reduce the kinetic energy of a trapped ion and place it in the ground state of its external degrees of freedom (its trapped motion) in Section 2.4.

### 2.1 Spatial confinement

Owing to its net charge, a single atomic ion can be isolated and held in place using relatively simple technology: the Paul trap, now in use for many decades. Driven by applications in quantum information science, there has been a push to miniaturize this technology and adapt it to become compatible with modern microfabrication processes.

### 2.1.1 Paul traps

One of the advantages of trapped-ion qubits is the ease with which they can be kept localized in space. Owing to their net charge, atomic ions can be confined for hours using only electric fields, though a static electric field will not do. Because it must satisfy Laplace's equation,  $\nabla^2 V = 0$ , the curvature of a static potential  $V$  will have an anti-confining component, meaning that a static electric field cannot form a stable ion trap. Intuitively though, this can be circumvented by rapidly switching between confinement in one direction and another, such that viewed on a slower timescale, the particle is stably trapped. This is the idea at work in a Paul trap (named for Wolfgang Paul, who was awarded the Nobel Prize in Physics for its invention [PH53, Pau90]), which employs an oscillating electric field to form a dynamically stable trap.



**Figure 2-1:** (a) Electrode geometry of a generic 4-rod Paul trap (with the ion position indicated in blue), as discussed in the text. The voltages applied to diagonally-opposed rods are the same. (b) A cross-sectional view of the 4-rod Paul trap showing electric field lines and the coordinates used in the text.

A Paul trap confines charged particles at the center of a radio-frequency (RF) electric quadrupole potential. A simple trap geometry, known as a four-rod Paul trap, is pictured in Figure 2-1. An RF potential is applied between diagonally opposite rods, which are fixed in a quadrupole configuration, to provide confinement along the central axis of the rods. Static electric potentials are applied to plug the ends of the rods, creating a trap in the center. Traps of this kind, which are constructed from precisely-machined metal rods, have been widely used for experiments in areas

ranging from mass spectrometry to quantum information science.

Typically, the amplitude of the ion motion is small relative to the spacing between electrodes ( $R$ ), so only the lowest-order (harmonic) terms in an expansion of the electrode potentials need be considered. Near the central axis of the trap (along the  $z$ -axis in Figure 2-1), the RF potential is well approximated by [WMI<sup>+</sup>97]

$$V_D(x, y, z) = \frac{1}{2} \kappa_D V_{RF} \cos(\Omega_{RF} t) \left( 1 + \frac{x^2 - y^2}{R^2} \right), \quad (2.1)$$

where  $\kappa_D$  is a geometric factor (of order 1). The DC potential is approximately

$$V_S(x, y, z) = \kappa_S V_{DC} \left( z^2 - \frac{x^2 + y^2}{2} \right), \quad (2.2)$$

where again  $\kappa_S$  is a geometric factor (of order 1).

In the  $z$ -direction (often called the “axial direction”), a static harmonic well is formed, resulting in an oscillation frequency for a single ion of

$$\omega_z = \sqrt{\frac{2\kappa_S Q V_{DC}}{M}}, \quad (2.3)$$

(where  $M$  is the ion’s mass and  $Q$  is the ion’s charge). There is no force along  $\hat{z}$ , so all vectors can be restricted to the  $x$ - $y$  plane. In the  $x$  and  $y$  directions, the potentials  $V_D$  and  $V_S$  combine to yield the equations of motion:

$$\frac{d^2 r_j}{d\zeta^2} + (a_j + 2q_j \cos(2\zeta)) r_j = 0, \quad (2.4)$$

$$j = \{x, y\},$$

(where  $\vec{r} = [r_x, r_y]$  is the position of the ion in the  $x$ - $y$  plane), which take the form of the canonical Mathieu equation. Here,  $\zeta = \frac{\Omega_{RF}}{2} t$  is the dimensionless time, and the

dimensionless Mathieu parameters  $a_j$  and  $q_j$  are given by

$$a_x = a_y = -\frac{4\kappa_S QV_{DC}}{M\Omega_{RF}^2}, \quad (2.5)$$

$$q_x = -q_y = \frac{2\kappa_D QV_{RF}}{M\Omega_{RF}^2 R^2}. \quad (2.6)$$

For certain regions of the Mathieu parameter space defined by  $a_j$  and  $q_j$ , the ion follows bounded trajectories, so a stable trap is formed [Gho95].

In practice, the stable trajectories chosen most often satisfy  $a_j < q_j \ll 1$ , in which case the solution to Equation 2.4 is well approximated by

$$r_j(t) = r_j^0 \cos(\omega_j t) \left( 1 + \frac{q_j}{2} \cos(\Omega_{RF} t) \right), \quad (2.7)$$

$$j = \{x, y\}.$$

Here, the oscillation frequencies in the x and y directions (often called the “radial directions”) are given by

$$\omega_j = \frac{\Omega_{RF}}{2} \sqrt{a_j + \frac{q_j^2}{2}}, \quad (2.8)$$

and  $r_j^0$  depends on initial conditions. So, we see that the ion movement is made up of a large amplitude, harmonic motion at frequency  $\omega_j$  called the “secular motion”, and an amplitude modulated, fast driven motion at frequency  $\Omega_{RF}$  called the “micromotion”.

In the limit that  $q_i$  is small, we can neglect the micromotion and treat the ion as if it were confined in a three-dimensional harmonic potential (known as a “pseudopotential”) of the form

$$Q\Psi = \frac{1}{2}M\omega_x^2 x^2 + \frac{1}{2}M\omega_y^2 y^2 + \frac{1}{2}M\omega_z^2 z^2. \quad (2.9)$$

The oscillation frequencies  $\omega_x$ ,  $\omega_y$ ,  $\omega_z$  (defined above) are known as the “secular frequencies” because they describe the secular motion of the ion; they are usually on



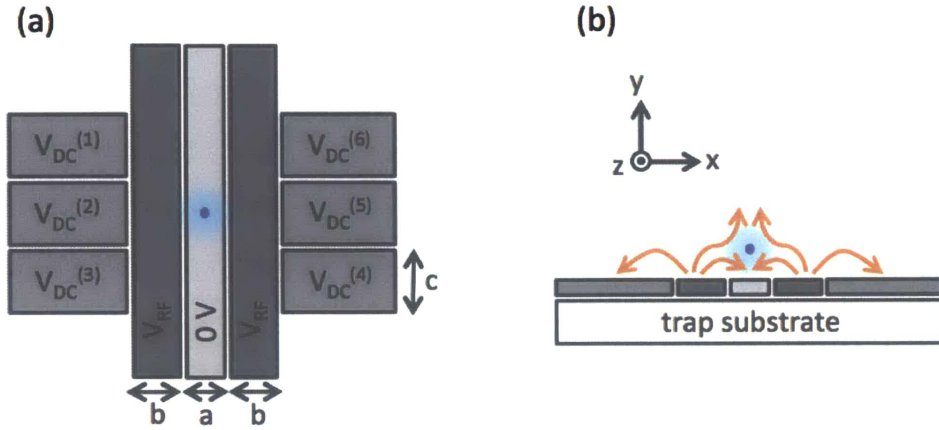
the order of  $\approx 1$  MHz for a trap frequency  $\Omega_{RF}$  of about  $\approx 20$  MHz. The resulting depth of the pseudopotential well in the radial and axial directions (known as the “trap depth”) is typically  $\approx 2$  eV. When sufficient vacuum is obtained (usually  $< 10^{-10}$  Torr), the lifetime of a trapped ion can be on the order of days.

### 2.1.2 Surface-electrode traps

The quadrupole potential used for trapping can also be produced by a combination of two-dimensional RF and DC electrodes [CBB<sup>+</sup>05]. For example: visualize squashing three of the rods into a plane (RF-ground-RF) and pulling the fourth electrode up to form a distant ground above the plane with the ion trapped in between. Traps of this form, known as “surface-electrode” or “planar” ion traps, can often be constructed using standard printed circuit board (PCB) manufacturing, or microfabrication techniques, making them far simpler to build than their precisely-machined, three-dimensional counterparts. Because planar traps are more amenable to being shrunk in size and replicated, they are better suited for scaling up to dense arrays of ions [KMW02]. Planar traps may also better facilitate integration of optical components [VCA<sup>+</sup>10, BEM<sup>+</sup>11, KHC11] or electronic devices [KPM<sup>+</sup>05].

A generic surface-electrode trap is pictured in Figure 2-2. The central three RF electrodes provide confinement in the radial directions (the x and y directions), and the outer DC electrodes add confinement in the axial direction (the z direction). Typically, an ion is trapped on the order of  $100 \mu\text{m}$  above the surface.

The equations of motion describing an ion in a planar trap are similar to those for a 4-rod trap (as discussed in Section 2.1.1). However, determining the electric potential generated in the space above the trap plane due to an arbitrary electrode configuration can be considerably more difficult. This amounts to a boundary-value problem for the Laplace equation in three dimensions with Cauchy boundary conditions. Often, the solution is obtained numerically using finite element methods [SPM<sup>+</sup>10]. But, in



**Figure 2-2:** (a) A generic surface-electrode ion trap electrode geometry. When  $V_{DC}^{(1)} = V_{DC}^{(3)} = V_{DC}^{(4)} = V_{DC}^{(6)}$  is positive, and  $V_{DC}^{(2)} = V_{DC}^{(5)}$  is negative, a symmetric trap is formed with the ion trapped above the surface at the center (indicated in blue). The trap dimensions,  $a$ ,  $b$ ,  $c$  are typically chosen to optimize properties such as the trap depth. (b) A cross-sectional view of the trap showing electric field lines and the coordinates used in the text.

cases where the electrodes are all rectangular (as in Figure 2-2), an analytic solution can be found by assuming a piecewise-constant Dirichlet boundary condition [Hou08].


Because the electrodes are confined to lie in a single plane, the trap depth of a surface trap is smaller than for a three-dimensional trap, usually  $\sim 100$  meV. Another drawback of the surface-electrode trap geometry is restricted laser access; lasers are constrained to lie in a plane parallel to the trap. Various laser beams are required for motional cooling and qubit manipulation (see Section 2.4). A laser beam must have a projection onto the appropriate trap axis to address a particular trap mode. This means that the symmetry of the surface-trap geometry must be explicitly broken using either the electrode shapes themselves or the DC voltages applied, so that all of the principal trap axes have a projection onto the cooling laser beam, and hence all trap axes can be cooled and manipulated.

## 2.2 Trapped ion qubit

As we will see, the electronic or nuclear spin states of an atomic ion confined in a Paul trap form a viable qubit. The trapped-ion qubit can be manipulated using laser or microwave pulses, and read-out by analyzing emitted fluorescence.

### 2.2.1 $^{88}\text{Sr}^+$ atomic ion

Our preferred ion is  $^{88}\text{Sr}^+$ , chosen for its mostly visible-wavelength atomic transitions. All the atomic states of interest are accessible using diode lasers, streamlining the laser systems required for experiments.  $^{88}\text{Sr}^+$  is a hydrogen-like atom, with no nuclear spin; a simplified level structure is shown in Figure 2-3. The strong transitions to the P states (at 422 nm and 408 nm) can be used for efficient laser cooling and state detection. The metastable D states ( $4\text{D}_{3/2}$  and  $4\text{D}_{5/2}$ ) have long lifetimes, so are well-suited to encoding quantum information. We choose to use the  $5\text{S}_{1/2} \leftrightarrow 4\text{D}_{5/2}$  transition as a qubit, with  $5\text{S}_{1/2} = |0\rangle$  and  $4\text{D}_{5/2} = |1\rangle$ .



**Figure 2-3:** Level structure of  $^{88}\text{Sr}^+$ , simplified to show only transitions relevant for cooling and state-detection. Air wavelengths are indicated for each transition, as well as the transition lifetime (shown in brackets).

## 2.2.2 Quantized ion motion

The kinetic energy of a trapped ion can be reduced by laser cooling so that it is comparable to  $\hbar\omega_i$ , at which point the motion of the ion in the confining potential is quantized. If we define (standard) creation/annihilation operators by

$$\hat{a}_j^\dagger = \sqrt{\frac{M\omega_j}{2\hbar}} r_j + i \frac{1}{\sqrt{2M\hbar\omega_j}} p_j \quad (2.10)$$

$$\hat{a}_j = \sqrt{\frac{M\omega_j}{2\hbar}} r_j - i \frac{1}{\sqrt{2M\hbar\omega_j}} p_j, \quad (2.11)$$


then we can write the Hamiltonian of the trapped ion in the standard form of a quantum harmonic oscillator:

$$\hat{H} = \sum_j \hbar\omega_j \left( \hat{a}_j^\dagger \hat{a}_j + \frac{1}{2} \right). \quad (2.12)$$


The complete Hilbert space for the trapped ion is then composed as a tensor product of the atomic (2-level) states and the motional (Fock) states. This is illustrated diagrammatically in Figure 2-4.

## 2.2.3 Quantum state detection

To perform state detection on the qubit transition ( $5S_{1/2} \leftrightarrow 4D_{5/2}$ ), the  $5S_{1/2} \leftrightarrow 5P_{1/2}$  transition (at 422 nm) is continuously driven; if the ion is in the  $5S_{1/2}$  state, this will result in photon scattering (at 422 nm), and fluorescence will be detected, meanwhile if the ion is the  $4D_{5/2}$  state, no photons can be scattered (at 422 nm) and no fluorescence will be detected. This state detection method is called “electron shelving” and is illustrated in Figure 2-5.



**Figure 2-4:** The tensor product of the atomic and motional states of the trapped ion result in a ladder of combined states. A transition which affects only the atomic state is known as a “carrier transition” (shown in black); a transition which increases the number of motional quanta by one is known as a “blue sideband transition” (shown in blue); and a transition which decreases the number of motional quanta by one is known as a “red sideband transition” (shown in red).



**Figure 2-5:** Relevant atomic transitions for quantum state detection using the electron shelving method. Light scattered (or not) on the  $5S_{1/2} \leftrightarrow 5P_{1/2}$  transition indicates the state of the qubit.

## 2.3 Basic laser-ion interactions

The internal (atomic) state of a trapped ion can be controlled through interaction with coherent light sources (lasers).

### 2.3.1 Hamiltonian

The Hamiltonian representing a two-level atom in a harmonic trap, interacting with the traveling wave of a single mode laser, is [VvL03]:

$$\begin{aligned}
 \hat{H} &= \hat{H}_0 + \hat{H}_I, \\
 \hat{H}_0 &= \hbar\omega_0\hat{\sigma}_z + \hbar\omega_j\hat{a}_j^\dagger\hat{a}_j, \\
 \hat{H}_I &= \hbar\Omega(\hat{\sigma}_+ + \hat{\sigma}_-)(e^{i(kr_j - \omega_L t + \phi)} + e^{-i(kr_j - \omega_L t + \phi)}), \tag{2.13}
 \end{aligned}$$

where the state of the ion is contained in  $\hat{H}_0$ , and the laser-ion interaction is contained in  $\hat{H}_I$ . For simplicity, we are only considering one of the three trap modes (equivalent to a one-dimensional harmonic oscillator). The operators  $\hat{\sigma}_z, \hat{\sigma}_+, \hat{\sigma}_-$  are the usual Pauli spin operators, and  $\hat{a}_j, \hat{a}_j^\dagger$  are the annihilation/creation operators for the chosen trap mode (described in Section 2.2.2). As before,  $\omega_j$  is the oscillation frequency in a particular trap direction  $r_j$ ,  $k$  is the laser wavenumber,  $\omega_L$  is the laser frequency, and  $\phi$  is the laser phase. The coupling constant  $\Omega$ , known as the Rabi frequency, fixes the interaction strength.

For an ion in a harmonic trap, the spatial extent of the zero-point wave function (in a chosen direction  $j$ ) is given by

$$\sigma_j = \sqrt{\langle 0 | (r_j)^2 | 0 \rangle} = \sqrt{\frac{\hbar}{2M\omega_j}} \tag{2.14}$$

which is typically  $\approx 10$  nm. The ratio of this length scale to the wavelength of the atomic transition is known as the ‘‘Lamb-Dicke’’ parameter. For a particular trap

mode  $j = x, y, z$ , the Lamb-Dicke parameter is given by

$$\eta_j = k\Phi\sqrt{\frac{\hbar}{2M\omega_j}}, \quad (2.15)$$

where  $k$  is the laser wavenumber, and  $0 \leq \Phi \leq 1$  is a geometric factor which takes into account the angle of the laser to the oscillation axis. The Lamb-Dicke regime is defined by the condition that

$$\eta_j^2(2n+1) \ll 1. \quad (2.16)$$

Written in terms of the Lamb-Dicke parameter, (and making the rotating wave approximation) the interaction term in the Hamiltonian becomes:

$$\hat{H}_I = \hbar\Omega \left( e^{i[\eta_j(\hat{a}_j + \hat{a}_j^\dagger) - \omega_L t + \phi]} \hat{\sigma}_+ + e^{-i[\eta_j(\hat{a}_j + \hat{a}_j^\dagger) - \omega_L t + \phi]} \hat{\sigma}_- \right) \quad (2.17)$$

### 2.3.2 Sideband transitions

Experimentally, it is advantageous to work in the so-called ‘‘Lamb-Dicke regime’’, in which  $\eta_j \ll 1$  (meaning that the atomic wavepacket is confined to a space much smaller than the wavelength of the atomic transition). In this regime, the interaction Hamiltonian can be expanded to obtain the three lowest order terms [Roo00]:

$$\hat{H}_I = \hbar\Omega_{n,n} (\hat{\sigma}_+ + \hat{\sigma}_-) + i\hbar\Omega_{n-1,n} (\hat{a}_j \hat{\sigma}_+ - \hat{a}_j^\dagger \hat{\sigma}_-) + i\hbar\Omega_{n+1,n} (\hat{a}_j^\dagger \hat{\sigma}_+ - \hat{a}_j \hat{\sigma}_-). \quad (2.18)$$

The first term in this interaction Hamiltonian, with coupling strength:  $\Omega_{n,n} \approx \Omega(1 - \eta_j^2 n)$ , corresponds to absorption/emission events that do not change the motional state of the ion, known as ‘‘carrier’’ transitions. The second term, with coupling strength:  $\Omega_{n-1,n} \approx \Omega(\eta_j \sqrt{n})$ , corresponds to absorption events that decrease the motional quantum number,  $n$ , of the ion by one quanta, known as ‘‘red sideband’’ transitions. Likewise, the third term, with coupling strength:  $\Omega_{n+1,n} \approx \Omega(\eta_j \sqrt{n+1})$ ,

corresponds to absorption events that increase the motional quantum number,  $n$ , of the ion by one quanta, known as “blue sideband” transitions. These transitions are shown in Figure 2-4. In the Lamb-Dicke regime, processes that change the motional quantum number  $n$  by more than one are strongly suppressed.


## 2.4 Laser cooling

Laser cooling encompasses those techniques in which laser-ion interactions are exploited to reduce the kinetic energy of trapped atoms. The method of laser cooling chosen depends on the properties of the transition and the trap [Roo00]:

When the trap frequency  $\omega_j$  is much smaller than the decay rate  $\Gamma$  of the atomic transition, the motional sidebands are spaced more closely than the linewidth of the transition. The velocity of the ion (due to the confining potential) changes much more slowly than the time required to absorb or emit a photon. So, the ion acts like a free particle exposed to a laser frequency with a time-dependent Doppler shift. In this regime, the laser can exert a velocity-dependent radiation pressure to cool the ion. This is known as Doppler cooling [HS75]. For an  $^{88}\text{Sr}^+$  ion, Doppler cooling is performed on the  $5\text{S}_{1/2} \leftrightarrow 5\text{P}_{1/2}$  transition (at 422 nm). Because there is some probability ( $\approx 1/20$ ) of decay to the  $4\text{D}_{3/2}$  state during each cooling cycle, this state needs to be continuously repumped as well (using a laser at 1091 nm) to maintain continuous cooling (see Figure 2-6).

Conversely, when the trap frequency  $\omega_j$  is much larger than the decay rate  $\Gamma$  of the atomic transition, the motional sidebands are spaced further apart than the linewidth of the transition. This is known as the “resolved sideband regime”, in which the laser can be tuned to address primarily one chosen sideband (see Figure 2-7). When this is done such that the energy of the emitted photon is greater than that of the absorbed photon, cooling is possible. This is known as sideband cooling [WI79]. For an  $^{88}\text{Sr}^+$  ion, sideband cooling is performed on the qubit transition,  $5\text{S}_{1/2} \leftrightarrow 4\text{D}_{5/2}$  (at 674 nm).





**Figure 2-6:** Relevant atomic transitions for Doppler cooling of  $^{88}\text{Sr}^+$ . During each cooling cycle on the  $5\text{S}_{1/2} \leftrightarrow 5\text{P}_{1/2}$  transition (at 422 nm), there is some probability ( $\approx 1/20$ ) of decay to the  $4\text{D}_{3/2}$  state, so this state is continuously repumped using a laser at 1091 nm.

To speed-up the cooling cycles, the lifetime of the upper level ( $4\text{D}_{5/2}$ ) is artificially shortened by coupling it to the the  $5\text{P}_{3/2}$  state (which has a much higher decay rate) using a laser at 1033 nm.


### 2.4.1 Doppler cooling

For simplicity, we will consider Doppler cooling of a two-level atom (with atomic frequency  $\omega_0$ ) interacting with a traveling wave laser field (of frequency  $\omega_L$ ). Each photon making up the laser field carries momentum:

$$\hbar k = \hbar \frac{\omega_L}{c}, \quad (2.19)$$

and these photons are scattered on the atomic transition at a rate of  $\Gamma \rho_{ee}$ , where  $\Gamma$  is the decay rate of the cooling transition and  $\rho_{ee}$  is the excited state probability. For an atom moving with velocity  $v$ , the excited state probability depends on the velocity via the Doppler shift [VvL03]:

$$\rho_{ee} = \frac{\Omega^2}{\Gamma^2 + 4(\Delta - kv)^2} \quad (2.20)$$



**Figure 2-7:** This cartoon spectrum illustrates what we mean by the “resolved sideband regime”. Two peaks in the spectrum due to sideband transitions (corresponding to a trap mode with secular frequency  $\omega_x$ ) are distinguishable from the peak due to the carrier transition (at frequency  $\omega_0$ ). By adjusting the laser frequency, it is possible to address an individual sideband transition (not drawn to scale).

(assuming low saturation), where  $\Delta = \omega_L - \omega_0$  is the detuning of the laser, and  $\Omega = \Gamma \sqrt{\frac{I}{2I_{sat}}}$  is the Rabi frequency. This results in a radiation pressure force on the atom of

$$F = \hbar k \Gamma \rho_{ee}. \quad (2.21)$$

For small velocities, we can expand the excited state probability to first order in  $v$ :

$$\rho_{ee} \approx \frac{\Omega^2}{\Gamma^2 + 4\Delta^2} \left( 1 + \frac{8\Delta k}{\Gamma^2 + 4\Delta^2} v \right), \quad (2.22)$$

resulting in a force that is (approximately) made up of a constant, radiation pressure component:

$$F_{constant} = \hbar k \Gamma \frac{\Omega^2}{\Gamma^2 + 4\Delta^2}, \quad (2.23)$$

and a viscous damping term (provided the detuning  $\Delta$  is negative):

$$F_{damping} = \hbar k \Gamma \frac{8\Delta k \Omega^2}{(\Gamma^2 + 4\Delta^2)^2} v. \quad (2.24)$$

Cooling is most efficient when we set  $\Delta = -\frac{\Gamma}{2}$ , in which case the damping force is

$$F_{damping} = -\frac{\hbar \Omega^2 k^2}{\Gamma^2} v. \quad (2.25)$$

The lowest temperature reachable by Doppler cooling is constrained by the recoil velocity induced by the photons scattered (photon recoil). For optimal settings, Doppler cooling can reduce the energy of a trapped ion to a limit of [DCT99, Ste86]

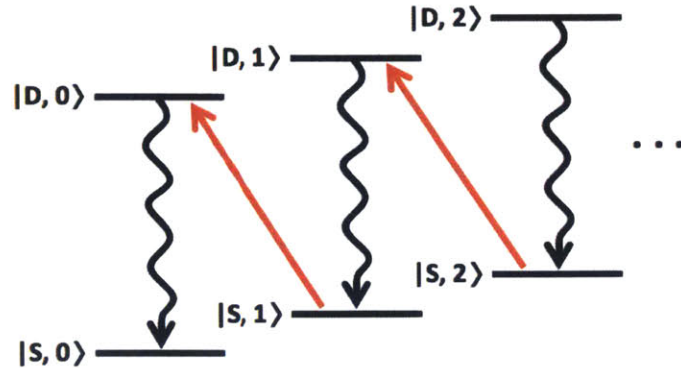
$$\bar{n} \approx \frac{\Gamma}{2\omega_j} \quad (2.26)$$

where  $\bar{n}$  is the mean occupation number in the harmonic trap. The corresponding temperature:  $T_D = \frac{\hbar\Gamma}{2k_B}$  (where  $k_B$  is Boltzmann's constant) is known as the Doppler temperature. Because the regime where this cooling method works is defined by:  $\omega_j \ll \Gamma$ , the mean quantum number is typically at least 10 following Doppler cooling. Due to photon recoil, it is not possible to cool to the motional ground state using only Doppler cooling methods.

## 2.4.2 Sideband cooling

Doppler cooling is typically used to efficiently cool an ion from an initial temperature (on the order of hundreds of quanta) to 10s of quanta. At this point, sideband cooling can be used to further cool the ion to the motional ground state [DBIW89]. To gain intuition about how sideband cooling works, it is useful to visualize the combined motional and atomic states of the ion as a ladder of states, as shown in Figure 2-8.

A single sideband cooling cycle proceeds as follows: first, a red sideband pulse is



**Figure 2-8:** Illustration showing sideband cooling of a trapped ion. Red sideband pulses (shown in red) excite the atom to the upper state, while reducing the number of motional quanta by one. Then, spontaneous emission occurs (shown in black), which does not change the motional state. After repeated cycles, the ion is left in the motional ground state.

applied, resulting in an excitation to the upper atomic state, and reduction by one of the motional quantum number. In the Lamb-Dicke regime (described in Section 2.3.2), spontaneous decays in which the motional state is changed are strongly suppressed, so the following spontaneous emission will most often leave the ion’s motional state unchanged. Hence, every cycle will remove one motional quantum until the ion has reached the bottom of the “ladder” of states (the ground state of the motion) at which point it is no longer affected by the red sideband pulses and cooling to the ground state is completed. In practice, because the lifetime of the upper atomic level is long (by choice), it is usually coupled to another atomic level with a higher decay rate to speed up each of the sideband cooling cycles.

### 2.4.3 Motional state detection

Most laser cooling techniques (including Doppler cooling and sideband cooling) leave the ion in a thermal state. For a thermal distribution, the probability of occupying

the  $n$ th Fock level is:

$$P_n = \frac{\bar{n}^n}{(\bar{n} + 1)^{n+1}}, \quad (2.27)$$

where  $\bar{n}$  is the mean motional quantum number. In the Lamb-Dicke regime (described in Section 2.3.2), the transition probability on the blue sideband is well-approximated by [LMM<sup>+</sup>97]:

$$P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = +\omega_j) = \sum_{n=0}^{\infty} P_n \sin^2 \left( \frac{\Omega_{n+1,n} t}{2} \right), \quad (2.28)$$

where  $\delta = \omega_L - \omega_0$  is the detuning of the qubit laser frequency from the atomic resonance. Rewriting the expression for the transition probability on the blue sideband as

$$P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = +\omega_j) = \frac{1}{2} \left( 1 - \frac{1}{\bar{n} + 1} \sum_{n=0}^{\infty} \left( \frac{\bar{n}}{\bar{n} + 1} \right)^n \cos(\Omega_{n+1,n} t) \right), \quad (2.29)$$

we see that with the qubit laser frequency fixed at  $\delta = +\omega_j$ ,  $\bar{n}$  can be extracted by scanning the pulse length  $t$  and fitting the resulting data in either the time domain or the frequency domain [WMM<sup>+</sup>96]. This is one way of characterizing the motional state of an ion.

Similarly, the transition probability on the red sideband is given by:

$$P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = -\omega_j) = \sum_{n=1}^{\infty} P_n \sin^2 \left( \frac{\Omega_{n-1,n} t}{2} \right). \quad (2.30)$$

Shifting indices, this can be rewritten as:

$$\begin{aligned}
 P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = -\omega_j) &= \sum_{n=0}^{\infty} P_{n+1} \sin^2\left(\frac{\Omega_{n,n+1}t}{2}\right) \\
 &= \frac{\bar{n}}{\bar{n}+1} \sum_{n=0}^{\infty} P_n \sin^2\left(\frac{\Omega_{n+1,n}t}{2}\right) \\
 &= \frac{\bar{n}}{\bar{n}+1} P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = +\omega_j). \tag{2.31}
 \end{aligned}$$

Hence, for any (fixed) Rabi frequency  $\Omega$  and pulse length  $t$ , the average number of quanta in a thermal distribution can be obtained through:

$$\begin{aligned}
 \bar{n} &= \frac{1}{r-1}, \tag{2.32} \\
 r &= \frac{P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = +\omega_j)}{P(|\downarrow\rangle \rightarrow |\uparrow\rangle, \delta = -\omega_j)}.
 \end{aligned}$$

So, by comparing the transition strengths of the blue and red sidebands, we can extract the mean motional quantum number [DBIW89]. This is a second, often simpler, way to determine the motional state of an ion, but it requires the sidebands to be resolvable.

# Chapter 3

## Experimental apparatus

This chapter describes the experimental system in which this thesis work was undertaken. Because this apparatus has previously been documented in Refs. [Bak06, LGA<sup>+</sup>08, ASA<sup>+</sup>09, Wan12], this chapter will focus on changes made during this work. First, the planar trap geometries employed are introduced, with a focus on a new geometry tested, in Section 3.1. Next, the experimental cryostat and associated hardware are described in Section 3.2. Changes to the optics used for laser delivery to the ion and for fluorescence collection are discussed in Section 3.3. Following this, the electronics and software systems required for experimental control are described in Section 3.4. Finally, some baseline measurements for the system are presented in Section 3.5.

### 3.1 Planar ion traps

This work focuses exclusively on microfabricated, surface-electrode ion traps. Two trap geometries have been explored: the (so-called) “Tower of London” trap geometry and the “Quanta Classic” trap geometry.

### 3.1.1 The “Tower of London” trap geometry

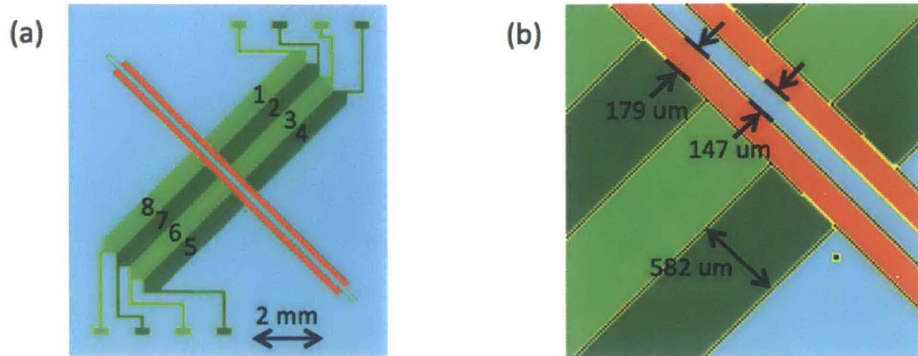
The Tower of London trap geometry is pictured in Figure 3-1, along with the dimensions required to set an ion height of about  $150\ \mu\text{m}$  above the surface. The two RF electrodes are usually directly connected together. This electrode geometry creates a linear Paul trap with the possibility for shuttling ions over a short distance ( $\approx 1\ \text{mm}$ ). The relative sizes of the electrodes are chosen such that the ratio of the RF electrode width to the ground electrode width is  $\approx 1.19$ , which maximizes the trap depth. The absolute electrode widths are chosen to set a trap height of  $\approx 150\ \mu\text{m}$ , which is convenient for laser access to the ion. The width of the DC electrodes is chosen to be  $\approx 4$  times the ground electrode width, which is optimal for shuttling ions (this allows for minimal change in the potential curvature with changing ion position). The length of the RF electrodes is chosen to be as long as possible (to minimize edge effects that would create micromotion in the axial direction), while not so long that the trap capacitance becomes very high (making it difficult to drive RF using a resonator). The DC electrodes break out into wires to match the ceramic pin grid array (CPGA) connections (to simplify wirebonding and avoid wirebonds intersecting laser beams).

Some typical trap voltages and the resulting trap parameters can be found in Table 3.1. The pseudopotential (computed using the model from Ref. [Hou08]) generated for these trap voltages is shown in Figure 3-2.

### 3.1.2 The “Quanta Classic” trap geometry

The Quanta Classic trap geometry is pictured in Figure 3-3, along with the dimensions required to set an ion height of about  $100\ \mu\text{m}$  above the surface. This electrode geometry creates a linear Paul trap with a single trapping zone. The notch in the central electrode defines a point where the RF field is zero, such that there exists a single, well-defined trap site. The RF electrode width is also asymmetric in the trapping region, which serves to create a tilt of the principal trap axes such that the





**Figure 3-1:** (a) Schematic showing the trap electrodes in the Tower of London geometry. The RF electrodes are shown in red, the ground electrodes are shown in blue, and the eight DC electrodes are shown in green. (b) The central trapping region of the Tower of London geometry with key dimensions indicated (for an ion height of about  $150 \mu\text{m}$ ). Dimensions are shown net of gaps between electrodes (which are  $15 \mu\text{m}$ ). The small marking visible in the outer ground electrode (in the bottom right of this image) is an alignment marking, used to align imaging optics to the trap during installation in the vacuum chamber.

cooling lasers (confined to lie in a plane parallel to the trap) have a projection onto all the axes and hence can cool all three degrees of motion of the ion. For more details about this trap geometry, and ideal trap voltages, see the theses [Lab08] and [Wan12].


### 3.1.3 Packaging traps

Details of trap fabrication will be discussed in Chapters 4, 5, and 6. Finished ion trap chips are packaged in a 100-pin ceramic pin grid array carrier (CPGA, Kyocera Corp. KD-P90171). Connections to the trap electrodes are made with aluminum wire bonds. To reduce RF pick-up and noise,  $1 \text{ nF}$  bypass capacitors (100 V, Skyworks SC99906068) are connected to each DC electrode and grounded on the CPGA using electrically-conductive, silver epoxy (EPO-TEK H20E). The amorphous ceramic material of the CPGA does not conduct heat well, so for good thermal contact the trap must be explicitly heat-sunk to the 4 K baseplate of the cryostat. To accomplish this, a thermally conductive trap substrate was chosen: sapphire, which has excellent ther-

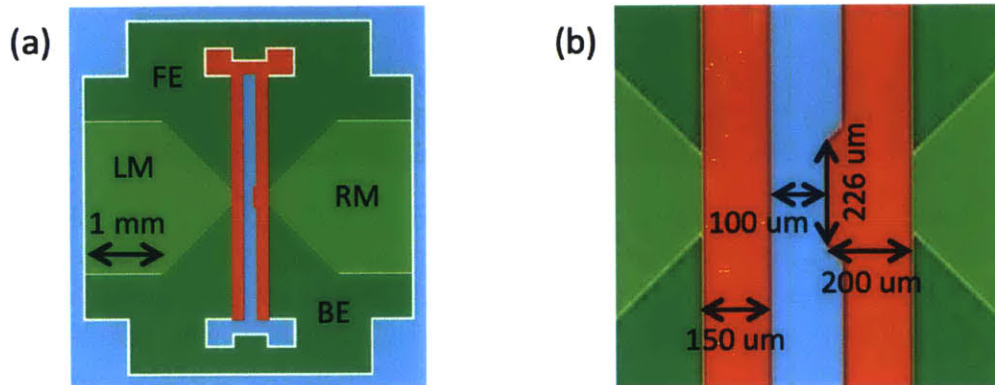
Electrode:	Voltage applied:
DC1	0 V
DC2	+19.326 V
DC3	-9.862 V
DC4	+19.326 V
DC5	-9.778 V
DC6	+0.823 V
DC7	-9.778 V
DC8	0 V
RF	150 V (amplitude)

**Table 3.1:** Computed voltages to create a trap centered between DC electrodes 3 and 6 with a tilt of the principal axes of  $15^\circ$ . For an RF frequency of 28 MHz, this corresponds to secular frequencies:  $\omega_x = 1.2$  MHz,  $\omega_y = 1.5$  MHz,  $\omega_z = 0.43$  MHz, and a trap depth of 56 meV.


mal conductivity at low temperature (orders of magnitude better than quartz, and comparable to copper). A copper spacer beneath the trap can then make a strong thermal contact to the substrate and provide a thermal bridge to the baseplate via a copper tab. A packaged, and heat sunk trap can be seen in Figure 3-4.



**Figure 3-2:** Computed pseudopotential in the  $\hat{x} - \hat{y}$  plane at the trap center. An ion would ideally be trapped at  $(x, y, z) = (0, 150, 0) \mu\text{m}$ .



**Figure 3-3:** (a) Schematic showing the trap electrodes in the Quanta Classic geometry. The RF electrode is shown in red, the ground electrodes are shown in blue, and the DC electrodes are shown in green. The labels LM, RM, BE, FE are the names given to the four DC electrodes. (b) The central trapping region of the Quanta Classic geometry with key dimensions indicated (for an ion height of about  $100 \mu\text{m}$ ). Dimensions are shown net of gaps between electrodes (which are  $10 \mu\text{m}$ ).



**Figure 3-4:** A surface-electrode trap packaged in a CPGA and installed in the cryostat. A copper tab connected to a copper spacer sitting below the trap can be seen, which provides a thermal contact to the 4 K baseplate. Bypass capacitors are also visible which are connected to each of the 8 DC electrodes with wirebonds.

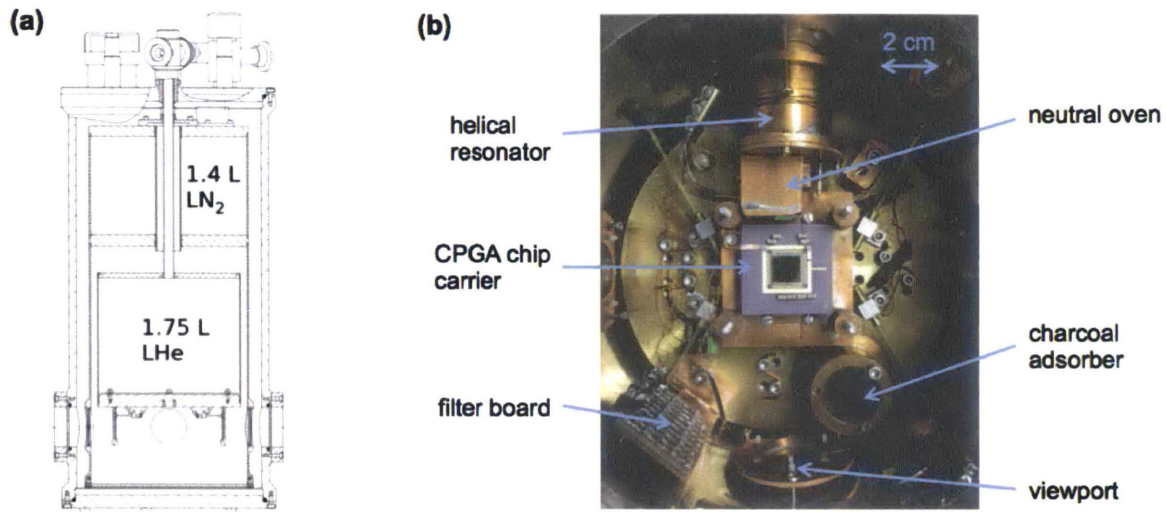
## 3.2 Experimental chamber

Traps are tested at cryogenic temperatures (4 K and 77 K), where low pressure is achievable within a short time frame, allowing for rapid trap characterization. Heating of the ion's motional state is also suppressed at cryogenic temperatures [LGA<sup>+</sup>08].

### 3.2.1 Bath cryostat

The simplest system for trapping at cryogenic temperatures is a bath cryostat, in which a vacuum enclosure is thermally anchored to one or more insulated tanks of cryogenics. We use a commercial liquid helium bath cryostat (QMC Instruments, Model TK 1813), depicted in Figure 3-5. The trap and other required components are mounted to the 4 K baseplate. There are three viewports on the chamber, allowing laser access to the trapping region, and one 55-pin electrical feedthrough. The base pressure at room temperature is only  $\approx 10^{-6}$  Torr, but cryosorption, aided by two activated charcoal adsorbents (mounted on the helium baseplate and at the top of the nitrogen shield to increase the effective surface area), results in sufficiently low pressure (estimated to be  $\approx 10^{-12}$  Torr [Bak06] in this system) for stable trapping. A filter board containing low-pass filters (2nd order, 4 kHz corner frequency) for each of the DC electrodes is heat sunk to the 4 K baseplate. To reduce the heat load on the LHe tank, the neutral oven (used to produce a beam of neutral Sr atoms at the trap site for subsequent photoionization and loading) and the helical resonator (used to step-up the RF trap voltage) are thermally anchored to the 77 K shield. Further details about the construction of the cryostat can be found in Ref. [ASA<sup>+</sup>09] and Ref. [Bak06].

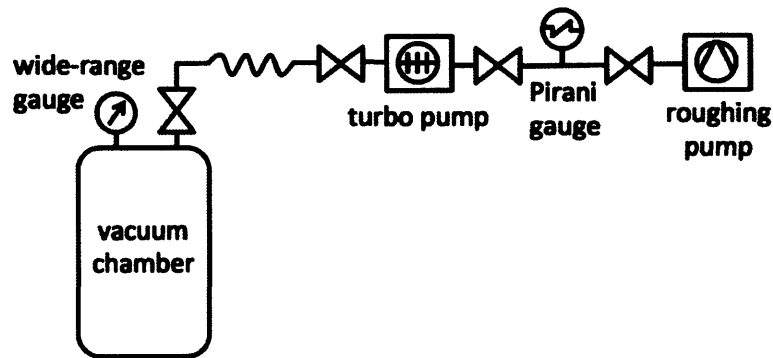




**Figure 3-5:** (a) Schematic of the bath cryostat, with liquid helium (LHe) and liquid nitrogen (LN<sub>2</sub>) tanks labeled. The 77 K shield attached to the LN<sub>2</sub> tank encloses the LHe tank and the 4 K baseplate. The working area (where the trap is mounted) is built around the 4 K baseplate, directly below the LHe tank. (b) Photograph of the 4 K baseplate with major components labeled.

### 3.2.2 Vacuum assembly

The vacuum assembly for the cryostat has been improved from Ref. [Wan12]. The turbo pump is now connected via the pumping line connection on the cryostat, which has a gate valve to seal-off the chamber once base pressure is reached. The spare vacuum port (previously connected to the turbo pump) is now being used for a wide-range gauge (Edwards, WRG-S-NW25). This allows us to measure the pressure at the chamber, as opposed to a distance away at the turbo pump (as before). The switch to using a wide-range gauge (which has a Pirani gauge for the upper pressure range, and an inverted magnetron gauge for the lower range) instead of an ion gauge (as before) also allows us to monitor the pressure during the entire process of pumping down the chamber. Multiple gate valves have been added to the vacuum lines, as can be seen in Figure 3-6.



**Figure 3-6:** Diagram showing the vacuum apparatus for the cryostat (all the indicated valves are gate valves).

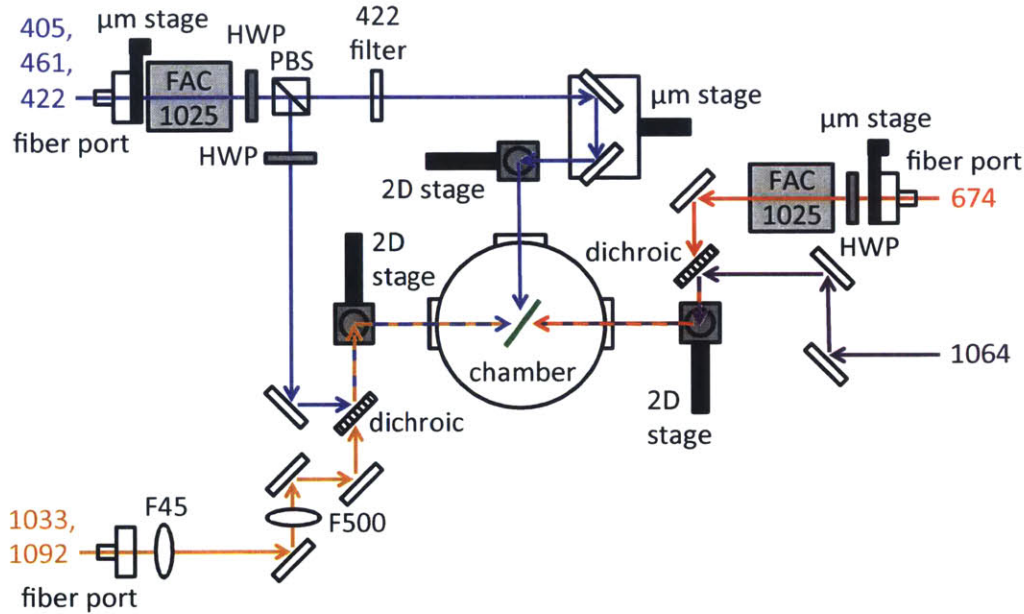
### 3.3 Light delivery and collection

There are three viewports on the side of the cryostat used to deliver laser beams to the trap, and one large window below the vacuum chamber (on the bottom of the cryostat) used for ion imaging. Along with the lasers used for cooling and qubit manipulation, two additional lasers at 461 nm and 405 nm are input to ionize neutral strontium atoms.

#### 3.3.1 Laser delivery optics

The optics for light delivery to the cryostat have been completely redesigned from Ref. [Wan12] in order to simplify loading of the trap and improve the stability of beam alignment. The Doppler cooling light (at 422 nm) is now co-propagating with the photoionization light (at 461 nm and 405 nm) as well as the repumper light (at 1033 nm and 1092 nm), which greatly simplifies alignment of these laser beams to the ion, and hence initial loading of the trap. The three blue beams (at 422 nm, 405 nm, 461 nm) are combined into a single fiber on a separate optical breadboard, also housing the double-pass AOM used to frequency-shift and shutter the Doppler cooling beam. As shown in Figure 3-7, all of the beams enter the cryostat from

periscopes adjustable in two-dimensions using motor-controlled stages. This allows for improved repeatability and precision of laser alignment to the ion (the resolution of the motor-controlled stages is  $\approx 5 \mu\text{m}$ ). Physically, the new optics are set-up in a modular fashion on separate optical breadboards which are bolted to the experiment table. All of the laser beam parameters are collected in Table 3.2.



**Figure 3-7:** Diagram showing optics for laser beam delivery to the cryostat. The orientation of the trap axis within the chamber is indicated by a green line. A single Thorlabs achromatic doublet pair (FAC1025-A) is used to focus the two Doppler cooling beams (at 422 nm) to a waist at the position of the ion. A half-wave plate (HWP) is used to adjust the power ratio between the two beam paths, and a single-axis translation stage in the 422-only path (the photoionization light is filtered out) is used to assure the path lengths are equal. Additional half-wave plates in the other 422 nm beam path and the 674 nm beam path assure that the correct polarization is set. Another achromat (FAC1025-A) is used to focus the 674 nm beam to a waist at the trap. The 1033/1092 nm beam is reduced in size and collimated using two lenses with focal lengths of 45 and 500 mm, respectively, while the 1064 nm beam is unaltered from its free-space, collimated waist size. Each beam enters the cryostat via a periscope adjustable using two motor-controlled stages.

Although not used in this work, an Nd:YAG laser is also installed on the optics



Laser purpose:	Wavelength:	Beam waist:	Power:
Doppler cooling	422 nm	30 $\mu\text{m}$	10 $\mu\text{W}$
Photoionization	461 nm	35 $\mu\text{m}$	100 $\mu\text{W}$
Photoionization	405 nm	40 $\mu\text{m}$	100 $\mu\text{W}$
Repumper (Doppler cooling)	1092 nm	200 $\mu\text{m}$	50 $\mu\text{W}$
Repumper (sideband cooling)	1033 nm	200 $\mu\text{m}$	50 $\mu\text{W}$
Sideband cooling (qubit)	674 nm	40 $\mu\text{m}$	1 mW
Laser desorption (pulsed)	1064 nm	3 mm	variable

**Table 3.2:** Summary of laser beam parameters at the cryostat. The beam waist is measured at the position of the ion. The power is a typical experimental value (which is variable), measured at the input of the light-delivery optics on the experimental table (just outside of the vacuum chamber).

table. This laser is intended for laser desorption cleaning of the trap surface. Laser pulses at 1064 nm are emitted into free-space, co-propagating with the 674 beam so that the 1064 nm beam position on the trap can be imaged (the camera is not sensitive to infrared light).

### 3.3.2 Imaging

The overall design of the imaging optics has not substantially changed from Ref. [LGA<sup>+</sup>08]. We simply modified the way the PMT was mounted to be more stable, and to decouple it from the pinhole used to spatially filter light from the ion. Now, a pinhole with a variable diameter, and an adjustable distance from the ion has been installed. Another improvement to light collection resulted from adding alignment markings on the trap so that the lenses inside the cryostat can be better aligned to the trap zone before closing and pumping down. The chamber and imaging optics are now enclosed in black acrylic and cloth to reduce background counts due to stray light.

## 3.4 Control electronics

The experiment relies on a diverse set of control electronics and software systems. Custom-built electronics based on FPGAs or microcontrollers are complemented by commercial electronics (function generators, power supplies, and so on) to run the experimental apparatus and perform data acquisition. These elements are controlled by a master computer via GPIB and network interfaces.

### 3.4.1 DC voltages: SuperDAC

Another major upgrade since Ref. [Wan12] has been the switch to DAC-driven DC electrodes on trap. This has allowed for better voltage resolution (which translates into better resolution with which the ion can be positioned in space) and opened up the possibility for shuttling ions. Previously, power supplies with a resolution of 25 mV set the DC trap voltages, corresponding to an ion positioning accuracy of  $\approx 0.4 \mu\text{m}$ ; a DAC system (with resolution better than 100  $\mu\text{V}$ ) allows for ion placement with an accuracy of  $\approx 1 \text{ nm}$ .

A new hardware system, dubbed “SuperDAC” has been designed for this purpose (schematics in Appendix A). SuperDAC is comprised of three printed circuit boards (PCBs):

1. DAC board: used to set DC voltages in the range of -10 V to +10 V with a 20-bit resolution; provides 4 buffered outputs per board; based on Analog Devices AD5791 DAC (more details in Appendix A.1);
2. Amplifier board: used as a gain stage following the DACs to increase the output voltage range to -20 V to +20 V (or -75 V to +75 V); provides 4 gain stages per board; based on Analog Devices OP42 op amp (more details in Appendix A.2);
3. ADC board: used to measure the SuperDAC outputs with between 20- and 24-bit resolution (depending on the speed) to allow for precise calibration (in


software); provides 4 inputs for measurement; based on the Texas Instruments ADS1248 ADC (more details in Appendix A.3).

These three boards can function completely stand alone, or can be snapped together (using connectors on the front and back sides of each board) to make a 3-layer stack with the DAC board on the bottom, the ADC board in the middle, and the amplifier board on the top (since it may get warm). Owing to this modular design, the different boards will likely find many other uses in the laboratory. A complete three-board unit can drive four DC trap electrodes (and simultaneously measure the voltages applied). Digital control can be achieved using either an FPGA or a microcontroller. The SuperDAC PCBs provide isolation of all of the digital inputs/outputs using digital isolators (Analog Devices, ADUM1402 and ADUM1300).

To drive a “Tower of London” style ion trap with 8 DC electrodes, a microcontroller-based (Arduino Mega2560) version of SuperDAC with 2 DAC boards, 2 ADC boards, 2 amplifier boards, and 8 voltage outputs has been constructed. Figure 3-8 shows the completed unit. The voltages can be set and monitored via the laboratory data acquisition software, or through a stand-alone user interface, which is pictured in Figure 3-9.

The SuperDAC unit was calibrated using a multimeter (Agilent 34401A), and then characterized. First, we investigated how repeatable voltage setting of SuperDAC is, and how repeatability changes over time after SuperDAC has been calibrated. We reset the voltage 50 times and measured the resulting output and ADC reading for 20 different voltage settings (in the range of -10 V to +10 V). We took this data immediately following calibration of SuperDAC and also after 8 hours and 24 hours. We used an Agilent 34401A,  $6\frac{1}{2}$  digit multimeter to measure the output voltage and calculate the errors; the accuracy of the multimeter is specified to be:  $\pm (0.002\%$  of reading) + (0.006% of range), which amounts to about  $\pm 100 \mu\text{V}$  for a reading of 5 V.

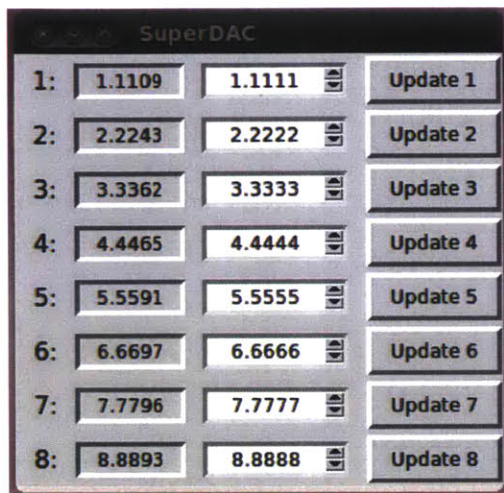
Based on the specifications of the DACs (Analog Devices AD5791), we expect an



**Figure 3-8:** Photographs of a SuperDAC unit with 8 outputs. (a) The front of the box, showing the BNC connections for power ( $\pm 25$  V), the BNC connections for the outputs (1 to 8) and the USB connection for computer control. (b) The inside of the box, showing the top of the two 3-PCB stacks (the amplifier board is on top), the regulator board (which is used to convert the  $\pm 25$  V input power to the various voltages required by the other boards, and the Arduino MEGA microcontroller.

output voltage accuracy of about  $\pm 20 \mu\text{V}$ . The ADCs (Texas Instruments ADS1248) were operated at a rate of 5 samples per second (5 SPS) during characterization, for which the specified readout accuracy is about  $\pm 10 \mu\text{V}$ . The results are tabulated in Table 3.3 and Table 3.4. We are likely limited by the resolution of the multimeter for these measurements. Nonetheless, the average root-mean-square output errors ( $\approx \pm 30 \mu\text{V}$ ) are very close to what we would expect based on the specifications of the DAC, and there is almost no change over 24 hours, which suggests that the calibration is good for more than a day. The average root-mean-square readout errors ( $\approx \pm 50 \mu\text{V}$ ) are a bit higher than predicted from the specifications of the ADC, and after 24 hours the root-mean-square error has increased by about  $15 \mu\text{V}$ , so (depending on the application) it may be necessary to calibrate the ADC board daily.

Next, we investigated how stable the output voltages are over time (when not reset). To do this, we measured the output voltage once every second (again using the Agilent 34401A multimeter) for 12 hours. We used a small voltage set-point ( $50 \text{ mV}$ ) to obtain better accuracy from the multimeter; in this range, the voltage



**Figure 3-9:** Screenshot showing the user-interface for SuperDAC. For each channel, the left box shows the ADC reading, and the right box allows the user to set the voltage (the hardware had not been calibrated yet, which is why the ADC readings are not as close as expected).

accuracy is specified to be  $\pm 5 \mu\text{V}$ . The results are shown in Table 3.5. As before, the measured values are possibly limited by the accuracy of the multimeter. It looks like the voltage stability of the outputs over time is extremely good; there is almost no change measured over a period of 12 hours.

The noise spectrum of one of the SuperDAC outputs is shown in Figure 3-10. AC noise near the secular frequencies is problematic because it can cause heating of the ion's motional state. For SuperDAC, we find the noise level is below -80 dBm for frequencies less than 3 MHz, which is sufficient (the output voltages are low-pass filtered in the cryostat before they are applied to the trap).

### 3.4.2 Trap RF voltage

The RF electronics are as depicted in Ref. [Wan12]. A helical resonator mounted on the 77 K shield is used to step up the RF voltage by factor of  $\approx 20$ , at a resonant frequency of about 30 MHz. The input signal to the resonator is provided by a

Output:	Voltage error after calibration:	Voltage error after 8 hours:	Voltage error after 24 hours:
1	19 $\mu\text{V}$	19 $\mu\text{V}$	29 $\mu\text{V}$
2	18 $\mu\text{V}$	20 $\mu\text{V}$	20 $\mu\text{V}$
3	21 $\mu\text{V}$	21 $\mu\text{V}$	21 $\mu\text{V}$
4	33 $\mu\text{V}$	32 $\mu\text{V}$	31 $\mu\text{V}$
5	34 $\mu\text{V}$	33 $\mu\text{V}$	34 $\mu\text{V}$
6	67 $\mu\text{V}$	66 $\mu\text{V}$	68 $\mu\text{V}$
7	24 $\mu\text{V}$	24 $\mu\text{V}$	27 $\mu\text{V}$
8	22 $\mu\text{V}$	21 $\mu\text{V}$	21 $\mu\text{V}$
Average	30 $\mu\text{V}$	30 $\mu\text{V}$	31 $\mu\text{V}$

**Table 3.3:** Root-mean-square voltage error for the SuperDAC outputs, averaged over the range -10 V to +10 V, with 50 trials per point. Comparison with a multimeter (Agilent, 34401A) is used to compute the errors.

frequency generator (Agilent, 33250A) and a power amplifier (Mini-Circuits, T1A-1000-1R8). A bi-directional coupler outside the chamber is used to measure the reflected power from the resonator and trap, allowing us to determine the resonant frequency. Typical RF voltages at the trap are between 100 and 300 V (amplitude).

### 3.4.3 Data acquisition software


We have built on the data-acquisition software developed in Ref. [LGA<sup>+</sup>08] to incorporate virtually all of the equipment required to trap ions and take data. With the addition of servo-controlled shutters, and motor-controlled optical stages, the experiment can now be fully controlled from a single computer console. Compared to several years ago when running the experiment entailed dashing up and down ladders to turn knobs and jogging to a computer near the lasers to adjust frequencies, this has made a big difference to experimenter stamina.

Output:	Readout error after calibration:	Readout error after 8 hours:	Readout error after 24 hours:
1	46 $\mu\text{V}$	48 $\mu\text{V}$	73 $\mu\text{V}$
2	56 $\mu\text{V}$	56 $\mu\text{V}$	74 $\mu\text{V}$
3	48 $\mu\text{V}$	51 $\mu\text{V}$	70 $\mu\text{V}$
4	57 $\mu\text{V}$	58 $\mu\text{V}$	67 $\mu\text{V}$
5	45 $\mu\text{V}$	51 $\mu\text{V}$	84 $\mu\text{V}$
6	66 $\mu\text{V}$	60 $\mu\text{V}$	83 $\mu\text{V}$
7	46 $\mu\text{V}$	48 $\mu\text{V}$	47 $\mu\text{V}$
8	56 $\mu\text{V}$	56 $\mu\text{V}$	56 $\mu\text{V}$
Average	53 $\mu\text{V}$	54 $\mu\text{V}$	69 $\mu\text{V}$

**Table 3.4:** Root-mean-square readout error for the SuperDAC ADCs, averaged over the range -10 V to +10 V, with 50 trials per point. Comparison with a multimeter (Agilent, 34401A) is used to compute the errors.

Output:	Fluctuation over 2 minutes:	Fluctuation over 1 hour:	Fluctuation over 12 hours:
1	3 $\mu\text{V}$	3 $\mu\text{V}$	4 $\mu\text{V}$
2	2 $\mu\text{V}$	3 $\mu\text{V}$	6 $\mu\text{V}$
3	2 $\mu\text{V}$	2 $\mu\text{V}$	2 $\mu\text{V}$
4	3 $\mu\text{V}$	2 $\mu\text{V}$	2 $\mu\text{V}$
5	3 $\mu\text{V}$	3 $\mu\text{V}$	3 $\mu\text{V}$
6	2 $\mu\text{V}$	5 $\mu\text{V}$	3 $\mu\text{V}$
7	2 $\mu\text{V}$	3 $\mu\text{V}$	3 $\mu\text{V}$
8	2 $\mu\text{V}$	2 $\mu\text{V}$	3 $\mu\text{V}$
Average	2 $\mu\text{V}$	3 $\mu\text{V}$	3 $\mu\text{V}$

**Table 3.5:** Fluctuations (standard deviation) in the SuperDAC output voltages (measured once per second) for different time intervals. A multimeter (Agilent, 34401A) is used for measurement.



**Figure 3-10:** Noise spectrum of SuperDAC unit (channel 1); the noise floor is shown in blue, and the signal is shown in red. The resolution bandwidth was set to 3 kHz and the input impedance was set to 50  $\Omega$ .



## 3.5 Baseline measurements


With a single ion, the first, essential elements of any measurement are placing the ion at the RF null (micromotion compensation), optimizing laser cooling (Doppler cooling and/or sideband cooling) and assuring accurate qubit rotations. This section presents typical experimental scans, acquired with a single ion in a Tower-of-London trap, in order to provide a baseline for measurements taken in this experiment.

### 3.5.1 Micromotion amplitude

Micromotion amplitude is measured by detecting photon counts from the ion that are in phase with the RF drive [BMB<sup>+</sup>98]; it is normalized to the total number of photon counts. Figure 3-11 shows the micromotion amplitude as a function of the compensation voltage applied to the DC electrodes. The compensation voltage is applied differentially to the electrodes so that changing the compensation voltage results in translation of the ion in a chosen direction (in this case, in the  $\hat{x}$ -direction). This data was obtained while running a script to minimize the micromotion amplitude. Usually, the micromotion amplitude is  $\approx 0.02$  following compensation.

### 3.5.2 Blueline spectrum

The spectrum of the transition used for Doppler cooling ( $5S_{1/2} \leftrightarrow 5P_{1/2}$ ) is measured because it is a useful diagnostic for Doppler cooling efficiency. For optimal cooling, the laser power should be set as close to the saturation intensity as possible, and the frequency detuning from resonance should be set to half the linewidth [Roo00]. Because the transition frequency is blue (422 nm) this type of data is referred to as a “blueline” scan; a typical scan is shown in Figure 3-12. The measured linewidth usually exceeds the natural linewidth of 22 MHz (full width at half maximum). After adjusting the power level, the measured linewidth (full width at half maximum) is



**Figure 3-11:** A typical scan obtained while compensating micromotion in the  $\hat{x}$ -direction. The micromotion amplitude is normalized to the total number of photon counts.

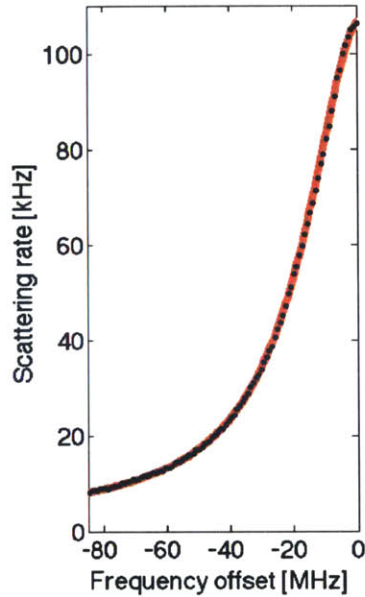
typically about 40 MHz and the maximum photon count rate is close to 100 kHz.

### 3.5.3 Sideband spectrum

The transition probability (quantified experimentally by measuring the shelving rate into the  $4D_{5/2}$  state) for each of the three first-order red sidebands and the three first-order blue sidebands (corresponding to the three trap modes) constitutes a sideband spectrum. From this, the average motional state occupation,  $\bar{n}$ , can be inferred for each trap mode. A typical sideband scan is shown in Figure 3-13. Following sideband cooling, we can reach temperatures with  $\bar{n} \leq 0.1$  motional quanta.


### 3.5.4 Rabi oscillation

With zero detuning from resonance, we drive carrier transitions on the qubit line,  $5S_{1/2} \leftrightarrow 4D_{5/2}$  (at 674 nm). Depending on the laser pulse length applied, we may




**Figure 3-12:** Spectrum of the  $5S_{1/2} \leftrightarrow 5P_{1/2}$  transition (at 422 nm). When the frequency offset from the atomic resonance becomes positive, the ion is heated by the laser. The data points (shown in black) are fitted to a Lorentzian (shown in red), which has a linewidth (full width at half maximum) of 42 MHz.

leave the ion in the  $4D_{5/2}$  state, or in a superposition of the  $5S_{1/2}$  and  $4D_{5/2}$  states. The resulting oscillation in the measured shelving rate as the pulse time is varied (known as Rabi oscillation) can be seen in Figure 3-14. The time for a population inversion (for the population to be completely transferred to the  $4D_{5/2}$  state) is known as the “Pi time”. For typical laser intensities (approximately  $2 \times 10^5 \text{ W/m}^2$  at 674 nm), the Pi time is 5 - 6  $\mu\text{s}$ . The contrast (representing how fully population transfer occurs) is typically close to 99%.



**Figure 3-13:** Sideband spectrum of the qubit transition,  $5S_{1/2} \leftrightarrow 4D_{5/2}$  (at 674 nm) following sideband cooling of all three modes close to the ground state. The lowest frequency sidebands correspond to the axial mode, and the two higher frequency sidebands correspond to the radial modes. Because the tilt of the trap axes is small, there is only a small projection of the laser beam onto the highest frequency radial mode, which is why this line is much narrower.



**Figure 3-14:** Rabi oscillations on the qubit (carrier) transition,  $5S_{1/2} \leftrightarrow 4D_{5/2}$  (at 674 nm), following sideband cooling. Each data point (shown in black) represents 100 experiments, and is fitted to a sinusoid (shown in red), which has a  $\pi$  time of  $5.3 \mu\text{s}$  and a contrast of 99%.

# Chapter 4


## Transparent ion traps

Fluorescence collection sets the efficiency of state detection and the rate of entanglement generation between remote trapped ion qubits. Despite efforts to improve light collection using various optical elements, solid angle capture is limited to  $\approx 10\%$  for implementations that are scalable to many ions. In this chapter, we present a new approach to light capture, which is based on fluorescence detection through a transparent trap using an integrated photodetector. We begin by summarizing prior art, then we introduce our idea, which combines collection efficiency approaching 50% with scalability. We describe fabrication of transparent surface traps with indium tin oxide and experiments to verify stable trapping of single ions. As a proof-of-principle demonstration, the fluorescence from a cloud of ions is detected using a photodiode sandwiched with a transparent trap. Finally, we propose a highly-efficient and compact “entanglement unit” based on this design.

### 4.1 Fluorescence collection from ions

Efficient fluorescence collection is essential for fast, high-fidelity state detection and ion-photon entanglement [LHM<sup>+</sup>09], but is largely still implemented with bulk optics

and conventional photomultipliers (PMTs) or image-intensified charge coupled detectors (CCDs). In the typical apparatus, a high numerical aperture objective is placed near to the ion within the vacuum chamber, and relay optics are used to direct the light to a detector outside the chamber, as shown in Figure 4-1. The resulting solid angle capture is less than 5% of the fluorescence emitted from the ion. Scaling-up to dense arrays of trapped ions will require more efficient fluorescence collection from many ions in parallel.



**Figure 4-1:** Diagram outlining the conventional approach to fluorescence collection, which places a high numerical aperture objective near to the ion, and detects light emitted with a photomultiplier or charge-coupled detector located outside of the vacuum chamber. Dimensions shown are typical for experiments with surface traps.


There are various proposals to enhance fluorescence capture within a scalable architecture. An array of microfabricated phase Fresnel lenses in place of a bulk lens can provide more efficient light collection from multiple ions [SNJ<sup>+</sup>11, JSN<sup>+</sup>11]. Or, a micro mirror embedded into a planar trap can improve solid angle collection from an ion trapped above [HWS<sup>+</sup>11, NKM<sup>+</sup>10] to as much as  $\approx 10\%$  [MVL<sup>+</sup>11]. Greater solid angle capture is possible, at the cost of scalability, by placing the trapping site at the focus of a spherical [SCK<sup>+</sup>11] or parabolic [MGF<sup>+</sup>12] mirror. Similarly, the ion can be trapped within a high-finesse optical cavity for enhanced light collection into a single mode [MKB<sup>+</sup>02, KLH<sup>+</sup>04]. Integration of a planar ion trap with an optical fiber has been demonstrated [KHC11, BEM<sup>+</sup>11], although the solid angle collection



is low:  $\approx 3.5\%$  [VCA<sup>+</sup>10].

## 4.2 New idea: transparent traps

Because the emission pattern of a trapped-ion qubit is isotropic under most conditions, efficient fluorescence collection means capturing the largest solid angle possible. Yet in a surface trap, the  $2\pi$  solid angle obstructed by the trap is typically inaccessible (with the exception of trap designs incorporating a micromirror embedded into the trap beneath the ion [HWS<sup>+</sup>11, MVL<sup>+</sup>11, NKM<sup>+</sup>10]). This can be circumvented by collecting fluorescence *through* a planar trap via the underutilized  $2\pi$  solid angle below the ion. A surface electrode trap made of transparent materials would in principle allow for collection efficiency approaching 50%. This concept is illustrated graphically in Figure 4-2. Combining such a transparent trap with an array of detectors beneath opens up the possibility of massively parallel light collection.



**Figure 4-2:** Conceptual drawing showing how an ion trap with transparent electrodes would allow fluorescence typically obstructed by the trap to be transmitted to a photodetector below. For the (typical) dimensions shown, the solid angle capture possible approaches 50%.

Indium tin oxide (ITO), an electrical and optical conductor commonly used for applications such as touch screens and LCD displays, is a natural choice for transparent electrodes. But trapping ions with ITO electrodes poses significant difficulty because ITO has a resistivity about 1000 times higher than metals typically used for trap electrodes, which makes it difficult to apply sufficient RF voltage on trap. Because ITO is

an oxide, it may also be susceptible to laser-induced charging of its surface, which results in strong and difficult-to-control forces on the ions, inducing micromotion, large displacements, or even making the ions untrappable [HBHB10, WHL<sup>+</sup>11, SLM<sup>+</sup>12]. Motivated by these challenges, we have created a transparent surface trap fabricated with ITO, and collected ion fluorescence through it. We observe stable trapping in a first ITO trap. We then demonstrate a proof-of-principle prototype for scalable fluorescence detection in a second ITO trap by collecting light from a trapped ion cloud using a standard photodiode mounted below.

### 4.3 ITO trap fabrication


The Quanta Classic trap geometry is chosen (as described in Section 3.1.2), with trap frequencies in the range of 0.8-1.3 MHz for an RF frequency of 35 MHz, and a trap depth of about 300 meV. Trap fabrication begins with optical lithography using NR9-3000PY photoresist on quartz substrate. Next, ITO is deposited by RF sputtering with argon gas at a rate of 0.5 Å/s. Finally, the resist pattern is transferred to the ITO via lift-off with RD6. The resulting optical transmission of the ITO samples (including the polished quartz substrate), measured over a 4 mm<sup>2</sup> area with a 422 nm light source at room temperature, averages to  $\approx 60\%$ , which is about 10% lower than expected for commercial films [KPH<sup>+</sup>99]. The measured resistivity of ITO varies from  $1 \times 10^{-5} \Omega\text{m}$  to  $2 \times 10^{-5} \Omega\text{m}$  for different sputtering runs. To improve the conductivity of the RF electrodes, an additional lithography step is performed to deposit a thin gold layer on the RF electrodes only.

The first ITO trap, ITO-4K (see Figure 4-3(a)), has 400 nm of ITO for all trap electrodes, and an additional 50 nm of gold on the RF electrodes only, bringing the resistivity down to  $\approx 2 \times 10^{-8} \Omega\text{m}$ , which is comparable to a fully-metal trap. The second ITO trap, ITO-PD (see Figure 4-3(b)), also has 400 nm of ITO for all trap electrodes, but only 5 nm of gold on the RF electrodes for improved transparency



( $\approx 60\%$  for 5 nm of Au [SNTK03]) with a resistivity of  $\approx 3 \times 10^{-8} \Omega\text{m}$ .

This second trap is sandwiched with a photodiode. For a proof-of-concept demonstration, a commercially available PIN photodiode (Advanced Photonix PDB-C613-2) is used. This photodiode has an active area of  $86.4 \text{ mm}^2$ , corresponding to light collection from approximately 50% of the solid angle. The photodiode efficiency drops significantly at cryogenic temperatures due to carrier freeze-out. Measurements at 422 nm reveal that the photodiode responsivity changes little from  $\approx 0.1 \text{ A/W}$  as it is cooled from room temperature to 77 K, but plummets to  $\approx 0.01 \text{ A/W}$  at 4 K. For this reason, measurements with the photodiode are performed at 77 K.



**Figure 4-3:** (a) ITO-4K trap mounted in a CPGA; 50 nm of Au is visible on the RF electrodes and contact pads for wire bonding. (b) ITO-PD trap mounted on photodiode in a CGPA; with only 5 nm of Au on the RF electrodes, the photodiode is visible through the trap. (c) Diagram of trap geometry showing RF electrodes (red), ground electrodes (blue), and DC electrodes (green).

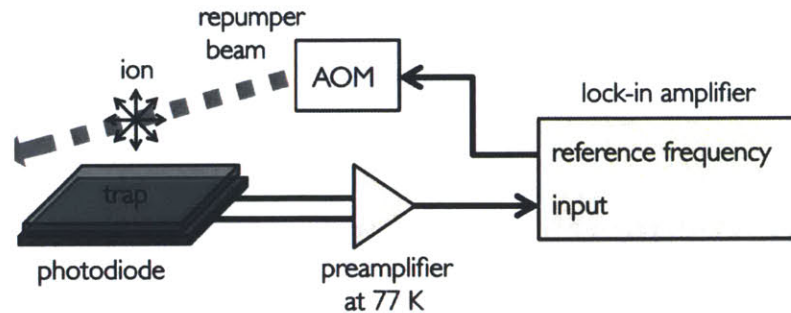
## 4.4 Trap verification and proof-of-concept prototype

Two ITO traps (ITO-4K and ITO-PD) have been tested at cryogenic temperatures (4 K and 77 K), where low pressure is achievable within a short time frame, and

heating of the ion's motional state is suppressed [LGA<sup>+</sup>08].  $^{88}\text{Sr}^+$  ions are trapped 100  $\mu\text{m}$  above the trap surface, which is heat sunk to the 4 K stage of a bath cryostat. The fluorescence emitted on the Doppler cooling transition ( $5\text{S}_{1/2} \leftrightarrow 5\text{P}_{1/2}$ ) at 422 nm is collected for state detection.

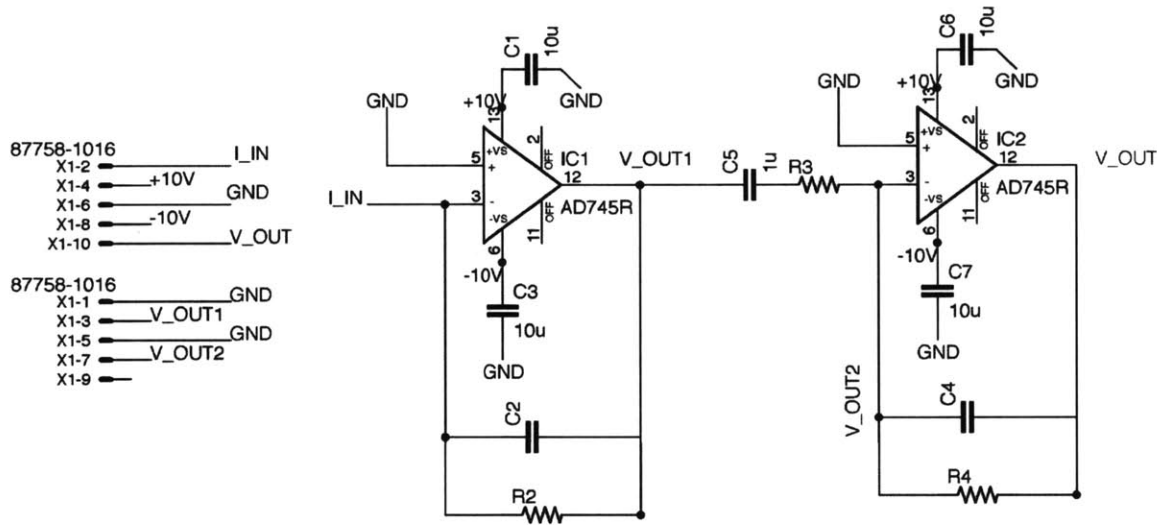
For ITO-4K, the estimated trap-surface temperature was 6 K. Single  $^{88}\text{Sr}^+$  ions were stably trapped with a lifetime of several hours, comparable to metal traps. No significant change in the micromotion amplitude was measured after crashing either a 405 nm or a 461 nm laser beam with intensity  $\approx 5 \text{ mW}/\text{mm}^2$  into the center of the trap[WHL<sup>+</sup>11] for 10 minutes, indicating that charging is not a major problem for ITO traps operated at cryogenic temperatures.

ITO-PD has an estimated trap-surface temperature of 77 K. Because the photodiode used has no internal gain mechanism, the resulting picoAmp-scale photocurrent is difficult to distinguish from electrical noise, making lock-in detection essential. The intensity of the laser addressing the  $4\text{D}_{3/2} \leftrightarrow 5\text{P}_{3/2}$  transition is chopped at 300 Hz, resulting in modulation of the ion fluorescence as population is successively trapped then pumped from the metastable  $4\text{D}_{3/2}$  state during Doppler cooling. Inside the cryostat, a custom preamplifier circuit mounted to the 77 K shield amplifies the signal with low added noise before it is input to a lock-in amplifier (Stanford Research Systems SR530) outside the chamber, as shown in Figure 4-4.



**Figure 4-4:** Apparatus for detection of ion fluorescence through a transparent trap using a photodiode mounted below. The acousto-optic modulator (AOM) is used to modulate the repumper for lock-in detection.


The preamplifier circuit consists of a transimpedance amplifier followed by a voltage gain stage, resulting in a net gain of  $\approx 2 \times 10^5$  V/A. The compact design is based on a low noise op amp (Analog Devices, AD745R); as shown in Figure 4-5 (the schematic) and Figure 4-6 (the board layout). The circuit was characterized for operation at 77 K where Johnson noise is suppressed. Figure 4-7 shows how the circuit is mounted to the 77 K shield using a copper spacer which fits in between the top of the shield and the plate sealing the inner vacuum chamber. The expected signal for  $\approx 50$  ions in this setup is estimated in Table 4.1 and compared with our conventional bulk optics and photomultiplier (PMT) setup.



**Figure 4-5:** Schematic for the preamplifier circuit.

The pressure in the cryostat when operated at 77 K is insufficient ( $\approx 1 \times 10^{-7}$  Torr) for stable trapping, so a cloud of ions was continually loaded. Figure 4-8(a) plots the photodiode voltage during initial loading of an ion cloud against the photon counts for light collected at the same time using bulk optics and a photomultiplier, indicating that the photodiode response is proportional to the fluorescence rate. Variation in the photodiode signal is likely due to scatter from the modulated repumper beam. Figure 4-8(b) compares the photodiode voltage before and after an ion cloud was loaded,






**Figure 4-6:** (a) The PCB artwork for the preamplifier circuit. The top layer routing is shown in red, the bottom layer routing is shown in blue, and the drill locations are shown in green. (b) The populated preamplifier circuit board before installation in the cryostat.

showing that the ion cloud fluorescence is distinguishable from the background. Over a measurement interval of a few minutes, the detected signal from the ion cloud averages to  $175 \pm 49$  mV, which is consistent with the signal predicted above for  $\approx 50$  ions.

These experiments indicate that significant improvements in quantum state detection are possible using transparent traps with integrated detectors. Assuming a noiseless amplifier, our photodiode signal could be used to distinguish between quantum states with greater than 99% fidelity with a 1 ms integration time. More generally, nearly 50% solid-angle collection is possible by using a photodiode with a large active area or focusing fluorescence onto the photodiode using additional optics below the trap. Then, the only losses before the detector occur in the ITO film and in the substrate. For commercially available ITO films these losses could be as low as 10% [KPH<sup>+</sup>99]. Replacing the photodiode by a device with an internal gain mechanism such as the Visible Light Photon Counter (VLPC), which has a quantum efficiency of




**Figure 4-7:** (a) Photograph showing how the preamplifier circuit is mounted to the 77 K radiation shield. The top of the chip packages make thermal contact to the copper mount. (b) Close-up view showing how the preamplifier circuit board is attached.

	Light collection efficiency	Power at the detector	Detector quantum efficiency	Photodiode current	Lock-in amplifier output
ITO-PD	30%	60 pW	30%	6 pA	120 mV
PMT	5%	10 pW	20%	n.a.	n.a.


**Table 4.1:** Comparison of photodiode and photomultiplier (PMT) collection efficiencies, with expected signal values for 50 ions, assuming a scattering rate of  $\approx 10^7$  photons/s per ion at 422 nm, resulting in  $\approx 200$  pW of total fluorescence into  $4\pi$  solid angle.

88% at 694 nm and 4 K [MKH09a], would allow a total detection efficiency of nearly 40%, compared to the typical 1-5% possible with a conventional photomultiplier and bulk optics. For our system, this would mean reducing the time required for quantum state detection with 99% fidelity from the current 200  $\mu$ s to only 5  $\mu$ s.



**Figure 4-8:** (a) Photodiode voltage and photomultiplier count rate, both background subtracted, during loading of an ion cloud. Each point is averaged over 30 seconds. (b) Histogram of photodiode voltages over a period of several minutes without ions (red), and after loading an ion cloud (blue).

## 4.5 Conclusion



**Figure 4-9:** Diagram of proposed compact entanglement unit (see text).

Our measurements establish that ITO is a viable material to use for microfabricated traps, and provide a first demonstration of light collection from ions through a trap with an integrated photodetector. The ability to form transparent traps opens up many possibilities to integrate ions with devices to transfer and detect light efficiently in a scalable architecture. One particularly interesting application is a compact entan-

glement unit, as shown in Figure 4-9. Here, transparent traps mounted on adjacent faces of a beam splitter house two (or more) ions to be entangled. Diffractive optics below the traps overlap the images of the two ions on detectors at the opposite faces of the cube, allowing for heralded entanglement generation between the ions [LHM<sup>+</sup>09]. For current state-of-the-art experiments, relying on bulk optics and light transmission in fibers, the total coupling efficiency for photons is only  $\approx 0.004$  [MOH<sup>+</sup>09], resulting in an entanglement generation rate of only  $\approx 2 \times 10^{-3} \text{ s}^{-1}$  (when the experiment is repeated at 100 kHz, assuming a branching ratio of 0.005, and a photodetector quantum efficiency of 15%). For the proposed compact entanglement unit (neglecting losses due to reflection at interfaces, and absorption in materials other than ITO), the coupling efficiency is  $\approx 0.45$ , resulting in a probability of entanglement generation of  $\approx 30 \text{ s}^{-1}$ , which is  $\approx 10^3$  times higher than the best rates achieved to date.





# Chapter 5

## Graphene-coated ion traps

Electric field noise originating from metal surfaces is a hindrance for a variety of micro-engineered systems, including for ions in microtraps, but is not well understood at the microscopic level. For trapped ions, it is manifested as motional-state heating inexplicable by thermal noise of electrodes alone, but likely surface-dependent. In this chapter, we investigate the role of surface properties in motional heating by creating an ion trap with a unique exterior. We describe the fabrication of an ion trap coated with graphene, a material free of surface charge and dangling bonds, and our experiments to characterize the trap. Surprisingly, we measure an average heating rate of  $1020 \pm 30$  quanta/s, which is about 100 times higher than is typical for an uncoated trap operated under similar conditions. Lastly, we discuss potential mechanisms to explain this result.

### 5.1 Motional heating of trapped ions

Miniaturization of traps is a natural first step towards realizing trapped-ion quantum computation at scale. But trap miniaturization comes at a cost: ions are exposed to worsening electric field noise with proximity to the trap surface (typically metal).

Noise near the motional frequency of the trapped ion couples to its charge, resulting in motional-state heating and gate error [TKK<sup>+</sup>00]. This is problematic because the shared motional state of ions in a trap forms a convenient “bus” used for nearly all multi-qubit gates; if the motional state is altered by noise during a gate operation, this will limit the fidelity. Perplexingly, the heating rates observed in experiments are larger than expected due to thermal (Johnson) noise from resistance in the trap electrodes or external circuitry [TKK<sup>+</sup>00, DK02].

Electric field noise is an obstacle shared by other micro- and nanofabricated devices incorporating metal surfaces and interfaces, including field-effect transistors [SDA<sup>+</sup>12], superconducting circuits [PAY<sup>+</sup>09], and microtraps for ultracold atoms [LTCV04]. It also poses a challenge to precision measurements ranging from studies of non-contact friction [SMS<sup>+</sup>01] and measurements of the Casimir-Polder force [KSDL10] to gravitational wave experiments [PSG08]. Because of their charge, and the exactness with which their motional state can be controlled and measured, trapped ions are exquisite probes for studying electric field noise near surfaces.

The origin of the excessive noise seen in ion traps still eludes researchers, but the heating rate has been observed to diminish by about 100 times when trap electrodes are cooled from room temperature (295 K) to 4 K, suggesting that it is thermally activated [DOS<sup>+</sup>06, LGA<sup>+</sup>08, LGL<sup>+</sup>08]. Termed “anomalous” [DK02], the heating rates observed are still orders of magnitude higher than expected from Johnson (thermal) noise in trap electrodes and connecting circuits [TKK<sup>+</sup>00]. The observations are consistent with an unfortunate  $d^{-4}$  scaling law (where  $d$  is the distance from the ion to the nearest electrode), as would be expected for a random distribution of fluctuating charges or dipolar patch potentials on the electrode surfaces [DHL<sup>+</sup>04].

Experimental evidence indicates that surface effects play a significant role in motional heating. An increased heating rate was measured near the loading zone of a trap [DNM<sup>+</sup>11], and no observable change in heating rate was measured as the bulk

of the trap electrodes transition to the superconducting state [WGL<sup>+</sup>10]. Further, the heating rate was recently shown to be insensitive to the trap material over a large temperature range [CS14]. Theoretically, a correlation length associated with surface disorder has been shown to characterize electric field noise generated near the surface [DCG09, LHC11]. A possible microscopic model points to fluctuations of the electric dipoles of adsorbed molecules on the trap surface as the source of excessive noise [DNM<sup>+</sup>11, SNRWS11]. This model is physically feasible; after exposure to the atmosphere, the trap surface typically develops at least several monolayers of adsorbates, and depending on the electrode material, a native oxide layer may also form.

## 5.2 Our approach: engineered surface

With the observations about motional heating described in Section 5.1 in mind, we became interested in creating a trap with distinctly dissimilar surface properties from metal. In particular, we were intrigued by the surface attributes of graphene—a material which exhibits an exterior free of dangling bonds and surface charge. The surface of graphene prepared by standard peel-off method is so pristine that atomic layer deposition (ALD) of metal oxide onto graphene fails because there are no dangling bonds or surface groups for ALD precursors to react with [WTD08]. Moreover, large graphene films grown on metal by chemical vapor deposition (CVD) have been shown to protect the underlying metal surface from oxidation and reaction with air, even following heating [SACS10, CBL<sup>+</sup>11]. These properties make graphene a compelling ion trap material to explore.

In this work we investigate motional heating in a surface trap with graphene-coated copper electrodes. Graphene-covered metal is chosen instead of wholly graphene in order to enhance conductivity of the RF electrodes. We hypothesize that such a trap will induce lower motional heating, because graphene will act as a field-transparent coating free of dangling bonds or surface charges, and may also help

by reducing contamination of the metal electrodes with potentially noise-generating oxides and surface adsorbates. We develop a fabrication process to synthesize monolayer graphene on a surface-electrode ion trap, and measure the resulting heating rate of a single ion confined above the trap surface.

### 5.2.1 Graphene as a protective coating for metals

Among its many unique properties [GN07], graphene has recently been shown to act as an extremely thin protective coating for metals [CBL<sup>+</sup>11]. Graphene is a single atomic layer of carbon atoms forming a tightly-packed, 2D honeycomb lattice. With complete internal  $sp^2$  bonding in each sheet, graphene has exceptional thermal and chemical stability [dHBW<sup>+</sup>07]. These surfaces of  $sp^2$  carbon allotropes form a natural diffusion barrier, allowing graphene to act as a thin—only 0.34 nm per layer—separation barrier to reactants. Early studies showed that even with defects, a closed graphene membrane (“balloon”) is impermeable to standard gases (helium, argon and air) [BVA<sup>+</sup>08]. A graphene coating has also been shown to inhibit reaction of an underlying metal with oxygen. For an intercalation layer of thin Fe film underneath graphene on Ni, the X-ray photoelectron spectroscopy (XPS) of core levels (and ultraviolet photoelectron spectroscopy (UPS) of the graphene-derived  $\pi$  states in the valence band) were found to be unchanged following exposure to large amounts of oxygen [DFRL08]. Similarly, in studies of emission of spin-polarized secondary electrons from Ni with a graphene coating, the magnitude of spin polarization was found to be insensitive to exposure to oxygen [DFL08]. Monolayer graphene, grown on Ru films, was found to uniformly cover curved surfaces and to protect the underlying metal surface against reaction with ambient gases [SACS10]. Additionally, oxygen adsorbates are only weakly bound to the graphene exterior, compared to uncoated metal.


## 5.3 Trap design

The ion trap geometry chosen for this investigation—the Quanta Classic (as described in Section 3.1.2)—has been well-characterized in terms of single-ion heating rates [LGA<sup>+</sup>08]. It is a five-electrode design in which a single ion can be confined 100  $\mu\text{m}$  above the trap surface (see Figure 5-1(a)). For this design, the trap frequencies (the frequencies of the harmonic trap, or “secular frequencies”) are in the range of 0.8 to 1.3 MHz for an RF frequency of 35 MHz, and the trap depth is about 300 meV. We choose copper as our electrode material for compatibility with MEMS and CMOS fabrication processes, and because it is most amenable to monolayer graphene growth.

The planar dimensions of the trap are 10 mm  $\times$  10 mm, which is typical for a surface trap, but large enough to preclude using graphene prepared by standard peel-off method to coat the electrodes. Instead, we employ CVD synthesis techniques on metal substrate to create large-scale graphene films. CVD synthesis methods utilize the temperature-dependent solubility of carbon in transition metals to control carbon precipitation on the surface of metal films, producing one or more graphene layers. Single-layer growth has been demonstrated on Cu [LCA<sup>+</sup>09], Pt [FKO05] and Ir [CDBM08], and multiple-layer growth on Ni [CZL<sup>+</sup>09, RJH<sup>+</sup>09, RTJ<sup>+</sup>09] and Ru [SFS08]. Remarkably, the resulting films remain uniform across multiple grain boundaries and surface steps of the underlying seed metal [LCA<sup>+</sup>09, SACS10]. We choose to grow monolayer graphene in order to avoid potential contamination build-up between multiple layers.

A graphene sheet can be synthesized on copper foil and transferred onto another substrate using PMMA film [LZC<sup>+</sup>09], and can also be patterned using photolithography [LRVGP09]. However, contamination and structural damage are liable to occur during these processes. There is also a risk that graphene will cause shorting between RF electrodes due to misalignment of the transferred film, or rough edges left by the lithography process. To avoid these concerns, we choose to synthesize graphene

directly onto the trap electrodes. Typically, CVD synthesis of graphene proceeds using Cu foils several 10s of microns thick [KZJ<sup>+</sup>09], but growth on evaporated Cu films only 500 nm thick has been demonstrated [LRVGP09] with adjustments to the growth conditions. We build on this work to synthesize graphene directly onto evaporated copper trap electrodes one micron thick. By limiting the number of processing steps, this approach reduces the risk of contamination or damage to the electrodes and graphene coating.



**Figure 5-1:** (a) Diagram of trap geometry showing RF electrodes (red), ground electrodes (blue), and DC electrodes (green). A single ion is confined  $100\ \mu\text{m}$  above the center of the ground electrode adjacent to the notch in the RF electrode. (b) Photograph showing one of the graphene-coated ion traps tested (MLG-1a) mounted in a CPGA (all traps tested are from a single fabrication batch and therefore look very similar).

## 5.4 Trap fabrication

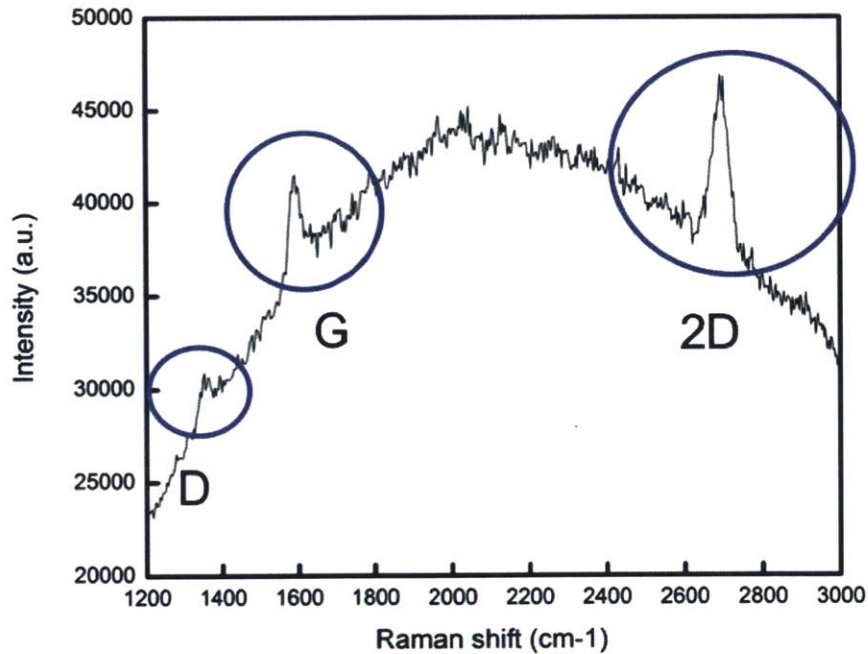
Fabrication begins with preparation of graphene growth substrates, i.e. the ion traps. Copper traps are fabricated on quartz substrate by e-beam evaporation of a 20 nm titanium adhesion layer, followed by evaporation of  $1\ \mu\text{m}$  of copper. Electrodes are patterned by coating the wafer with NR9-3000 photoresist and performing optical

lithography. Our apparatus for graphene synthesis is a low-pressure CVD chamber, similar to previous reports [RJH<sup>+</sup>09, RTJ<sup>+</sup>09]. CVD synthesis begins with annealing of the sample in hydrogen gas ( $H_2$ ) to smooth the metal surface and activate grain growth. Then, carbon is introduced into the electrodes—which act as the seed metal—by decomposing diluted methane gas ( $CH_4$ ). Finally, controlled cooling of the metal-carbon solid solution in a hydrogen environment promotes graphene precipitation. Because there is no seed metal in the gaps between the electrodes, no graphene is produced there.

With thin metal films, there is a propensity for dewetting to occur during graphene synthesis. We found that the parameters for graphene growth needed to be carefully optimized (in particular, the times for each step) in order to avoid formation of holes in the copper electrodes. The optimized schedule for our growth process is detailed in Table 5.1. To confirm the presence of monolayer graphene following CVD synthesis, Raman spectra are taken of the samples. Although there is significant Cu background, the G, D and 2D peaks of graphene are clearly visible. A small G/2D ratio, and a small D peak are observed, which suggest that the graphene is a single layer thick. The Raman spectra of a trap can be seen in Figure 5-2. Surface coverage was inferred from optical image intensity measured across graphene films transferred from copper samples onto 300  $\mu\text{m}$ -thick  $\text{SiO}_2$  substrate [LCA<sup>+</sup>09]. We estimate that there is  $\geq 90\%$  monolayer graphene coverage.

Growth step	Gases in chamber	Flow rate [sccm]	Temperature [°C]	Time [minutes]
1	H <sub>2</sub>	10	warm-up	10
2	H <sub>2</sub>	10	1000	10
3	CH <sub>4</sub>	20	1000	10
	H <sub>2</sub>	10		
4	H <sub>2</sub>	10	cool-down	20

**Table 5.1:** Growth parameters chosen for graphene synthesis on evaporated copper traps. During warm-up (and cool-down) the temperature of the chamber is being ramped up from room temperature to 1000 °C (and vice versa).



**Figure 5-2:** Raman spectra of a copper trap following graphene synthesis (with an integration time of 30 seconds). The broad hump is the Cu background and the indicated peaks correspond to graphene. The small G/2D ratio and small D peak suggest that the graphene is primarily monolayer.



## 5.5 Trap testing and heating rate measurements

Graphene-coated ion traps are packaged in a ceramic pin grid array (CPGA) carrier, which is then mounted to the vacuum chamber of a bath cryostat (see Section 3.2). Operation at cryogenic temperatures (77 K or 4 K) allows for vacuum pressure below  $10^{-10}$  Torr within a short time frame and avoids the need to outgas the vacuum chamber after installing the new trap. There is also a thermal suppression of the heating rate (relative to operation at room temperature) [DOS<sup>+</sup>06, LGA<sup>+</sup>08, LGL<sup>+</sup>08], which makes it possible to measure motional heating with high accuracy (by comparing sideband strengths).


During testing, a single  $^{88}\text{Sr}^+$  ion is loaded into the trap via photoionization of a thermal vapor and confined 100  $\mu\text{m}$  above the surface. Doppler cooling on the dipole-allowed  $5\text{S}_{1/2} \rightarrow 5\text{P}_{1/2}$  optical transition at 422 nm reduces the ion temperature to below 1 mK. A narrow diode laser locked to a stable external cavity permits access to the quadrupole-allowed  $5\text{S}_{1/2} \rightarrow 4\text{D}_{5/2}$  optical transition at 674 nm. This transition is used for sideband cooling [WI79, DBIW89] the lowest frequency mode of the ion at 0.86 MHz to the motional ground state with fidelity exceeding 95% [LGA<sup>+</sup>08].

Fluorescence emitted on the Doppler cooling transition at 422 nm is collected to distinguish between the  $5\text{S}_{1/2}$  and  $4\text{D}_{5/2}$  states (the ion emits no fluorescence when it is shelved in the metastable  $4\text{D}_{5/2}$  state). For the lowest frequency trap mode, the transition strengths on the red and blue motional sidebands are determined by measuring the probability of shelving in the  $4\text{D}_{5/2}$  state. The mean number of motional quanta  $\bar{n}$  can then be read out by comparing the transition strengths on the sidebands [TKK<sup>+</sup>00]. To measure the heating rate  $\dot{\bar{n}}$ , a delay time is added before  $\bar{n}$  read-out, which is then varied to monitor the average number of quanta versus time. For a bare metal trap (gold or copper of the same design chosen for this work) operated under comparable conditions, at 4 K, we typically measure a heating rate of about  $\dot{\bar{n}} = 10$  quanta/s. [LGA<sup>+</sup>08, LGL<sup>+</sup>08]. We have not observed any significant

difference in heating rate for metal traps that have been annealed at high temperature, versus other metal traps [WGL<sup>+</sup>10].

## 5.6 Results

Three graphene-coated ion traps were tested: MLG-a, MLG-b, and MLG-c. A packaged trap (MLG-c) is shown in Figure 5-1(b). The first trap tested, MLG-a, could confine ion clouds at 77 K, but was not tested at 4 K due to shorting between DC electrodes. The shorting was most likely caused by excessive Sr discharge coating the trap when the oven was outgassed (a large amount of Sr was observed on the CPGA when the trap was removed from the cryostat). The following two traps, MLG-b and MLG-c, had no shorting problems and both allowed for stable trapping and manipulation of a single  $^{88}\text{Sr}^+$  ion (see Figure 5-3(a)). The ion lifetime with cooling lasers on was several hours, similar to traps made of uncoated metal [LGA<sup>+</sup>08]. For MLG-b, the heating rate averaged over 20 measurements was  $\dot{n} = 1020 \pm 30$  quanta/s. The data are shown in Figure 5-3. For MLG-c, after trying many different cooling and voltage settings, sideband cooling could not be optimized to allow for ground state cooling. It is likely that the heating rate was prohibitively high for sideband cooling to outpace heating.



**Figure 5-3:** (a) A single, trapped  $^{88}\text{Sr}^+$  ion above MLG-c as seen on the experiment CCD camera (false colors). (b) Shelving rate (averaged over 20 measurements, each with 100 trials per point) on the red sideband (shown in red) and the blue sideband (shown in blue). (c) Plot showing the inferred average phonon number  $\bar{n}$  as a function of time. The linear fit indicates a heating rate of  $\dot{\bar{n}} = 1020 \pm 30$  quanta/s.

## 5.7 Discussion

The heating rate measured in the graphene-coated trap is  $\approx 100$  times higher than expected for a bare metal trap operated under comparable conditions [LGA<sup>+</sup>08, LGL<sup>+</sup>08]. This is a significant change, indicating that the graphene coating had a major impact on motional heating. But contrary to our hypothesis, the graphene layer increased rather than lessened the heating rate. It is possible that the presence of imperfections such as holes, multi-layer flakes and grain boundaries in the graphene coating contributed crucially to the electric field noise at the ion position, but this was not anticipated because monolayer-graphene-coverage is  $\geq 90\%$ . Surface roughness could also potentially compound the effect of defects in the graphene layer.

### 5.7.1 Surface roughening during CVD synthesis

We observed some roughening of the underlying copper surface following graphene synthesis. Scans of the trap surface using a profilometer reveal that before graphene synthesis, the evaporated traps have a root mean square (RMS) roughness of less than

10 nm. After graphene synthesis, the RMS roughness is closer to 200 nm. Some representative profilometer scans can be seen in Figure 5-4. The change in surface quality is clearly caused by the elevated temperatures required for CVD synthesis (close to the melting point of copper). It is not clear what effect, if any, surface roughness might have on the properties of the trap, other than to increase the background light scattering from laser beams grazing the trap surface.


### 5.7.2 Related studies: modifying surface composition in situ

As we began our study of motional heating in traps with unique surface composition, other researchers pursued studies of motional heating in standard (metal) traps with surface composition altered in situ. In these experiments, the heating rate was measured before and after argon-ion bombardment [HCW<sup>+</sup>12, DGB<sup>+</sup>14] or laser-desorption cleaning [AGH<sup>+</sup>11] of the surface. These studies sought to determine whether removing surface contamination—such as oxide and adsorbates—to reveal pristine metal electrodes would decrease the heating rate. Remarkably, both investigations into argon-ion-beam cleaning of the trap surface in-vacuum reported a 100-fold reduction in the heating rate, confirming that surface effects dominate anomalous heating.

The most recent of these experiments employed an apparatus with capability to monitor surface composition, perform argon-ion-bombardment, and measure the heating rate in situ [DGB<sup>+</sup>14]. The investigators found that changes in oxygen or carbon adsorbate coverage in vacuum have little effect on heating. Instead, it is posited that hydrocarbon molecules adsorbed during the preparatory vacuum bake could be significant contributors to heating prior to their removal during argon-ion-beam cleaning [H<sup>+</sup>4, HCW<sup>+</sup>12, HCW<sup>+</sup>13]. This is particularly relevant to our investigation because graphene is highly lipophilic; regardless of the production method used, hydrocarbon molecules are found on graphene [MKE<sup>+</sup>08]. Thus, hydrocarbon contamination may

be partially responsible for the elevated heating rate we measured.

Hydrocarbon contamination on graphene films can be monitored via atomic-resolution, high-angle dark field imaging and electron-energy-loss spectroscopy [BGB<sup>+</sup>09]. Assuming that hydrocarbon contamination plays a role in anomalous heating, it would be instructive to compare the heating rates induced for different types and arrangements of hydrocarbon deposits on the graphene surface. Graphene could be a good material choice for further studies because adsorbates other than hydrocarbons are weakly bound to the surface, and graphene coating prevents reaction of underlying metal electrodes with the environment [SACS10, CBL<sup>+</sup>11]. Such a study could potentially illuminate the microscopic mechanism behind anomalous heating. Additionally, trapped ions could potentially be employed as extremely sensitive probes of hydrocarbon or other deposits on graphene [CGD<sup>+</sup>10], which would be valuable for understanding its surface chemistry.



**Figure 5-4:** (a) Profilometer scan across 20  $\mu\text{m}$  in a central area of an evaporated copper trap. (b) Profilometer scan across 20  $\mu\text{m}$  of a trap from the same batch following graphene synthesis (notice the axis scale has changed). (c) Profilometer scan across 400  $\mu\text{m}$  of the same graphene-coated trap.

## 5.8 Conclusion

In summary, we have synthesized monolayer graphene on the surface of an ion trap with evaporated copper electrodes, and demonstrated stable confinement of single ions in the trap. We measured an average heating rate of  $1020 \pm 30$  quanta/s, which is  $\approx 100$  times higher than expected for a bare metal trap operated under comparable conditions, at 4 K [LGA<sup>+</sup>08, LGL<sup>+</sup>08]. Instead of reducing electric field noise, as we originally expected, the added graphene layer appears to have increased it. This may be consistent with recent results suggesting that hydrocarbon contamination, which is ubiquitous on graphene, plays a prominent role in motional heating. Further studies incorporating atomic-level surface monitoring of hydrocarbon deposits on graphene could improve our understanding of anomalous heating on the microscopic level.





# Chapter 6

## Ion traps in CMOS

Can we leverage the performance and scale of the CMOS manufacturing platform for quantum information processing? In this chapter, we describe how we overcome materials challenges to create a surface-electrode ion trap fabricated in a 90-nm CMOS (complementary metal-oxide-semiconductor) foundry process. With one of the interconnect layers defining a ground plane between the trap electrode layer and the p-type doped silicon substrate, we find that ion loading is robust and trapping is stable. We measure a motional heating rate comparable to those seen in other surface-electrode traps of similar size. Finally, we note that the CMOS process employed in this work includes doped active regions and metal interconnect layers, allowing for co-fabrication of standard CMOS circuitry as well as devices for optical control and measurement.

### 6.1 CMOS technology for trapped-ion systems

All of the traps described thus far were constructed using specialized, non-standard fabrication processes. The benefit of specialized fabrication processes—which are the norm in the ion-trapping community—is that they allow each trap to be tailored to the materials or devices to be integrated. The downside of using custom pro-

cesses is that repeatability at different facilities is very difficult, and features like the lithographic resolution are typically limited. All of this is in contrast to how classical processors are manufactured, using standardized CMOS (complementary metal-oxide-semiconductor) technology, which allows for devices with critical dimension of  $\approx 10$  nm, in quantities of about  $10^{11}$ , to be fabricated reproducibly on a common silicon substrate.

A typical ion trap fabrication process involves patterning a single metal layer atop a quartz or sapphire substrate. By contrast, a standard, high-resolution CMOS fabrication process typically involves patterning of 8 or more copper interconnect layers (separated by oxide) plus an aluminum pad layer, on a doped polysilicon substrate; the resulting active and passive layers can be used to form dense, high-performance digital circuits. Silicon substrates have been used in trap construction before, but none of these traps have had doped, active device fabrication available, and only a few metal layers (four maximum) have been implemented to date [LLC<sup>+</sup>09, SFB<sup>+</sup>10, Bri08, AHJ<sup>+</sup>12, WSGS12, WAF<sup>+</sup>13, SRW<sup>+</sup>14, NLK<sup>+</sup>14]. To investigate whether the CMOS manufacturing platform can be harnessed for trapped-ion systems, we designed an ion trap to be constructed in a standard, high-resolution CMOS fabrication process.


## 6.2 Trap design

The Tower of London trap geometry was chosen (as described in Section 3.1.1) because this design offers flexibility to create various trapping potentials and allows scalability to multi-zone traps with complex geometries. Other advantages include comparative ease in selecting control voltages, and relatively narrow RF electrodes, allowing for lower capacitive coupling and RF power dissipation in the trap. The electrodes are scaled to create a trapping site  $50 \mu\text{m}$  above the surface of the trap chip. Trap electrodes are patterned into the top aluminum pad layer (for more de-

Trap feature	Size
Center ground width (not including gaps)	50 $\mu\text{m}$
RF width (not including gaps)	60 $\mu\text{m}$
DC width (not including gaps)	215 $\mu\text{m}$
Gap size	6 $\mu\text{m}$
Bending radius of the rounded edges	5 $\mu\text{m}$
Length of RF electrodes	2 mm
Width of DC electrode breakout wires	18 $\mu\text{m}$

**Table 6.1:** Trap electrode dimensions chosen to create a 50  $\mu\text{m}$  trap height.

tails, see Section 6.3). The mask for the traps is shown in Figure 6-1. Segmented dc control electrodes are routed to wire bonding pads in the corners of the trap chip to prevent wirebonds from obstructing laser access. Because devices are fabricated on a shared, multi-project wafer, the mask area for the traps is limited to  $3 \times 3 \text{ mm}^2$ . Each finished trap is 2.5 mm long and 1.2 mm wide. Detailed trap dimensions are shown in Table 6.1.



**Figure 6-1:** Trap mask. The mask is 3 mm long, and 1.5 mm wide.

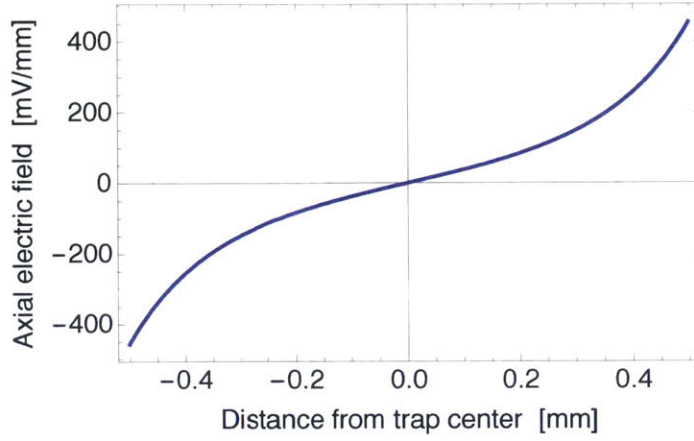
### 6.2.1 Edge-effects due to shortened RF electrode length

Each trap can be no longer than about 3 mm (including space for dicing the wafer), restricting the length of the RF electrodes considerably compared to a typical trap. The aspect ratio with respect to the ion height is still very large ( $\sim 30$ ), but it is worth checking if edge effects will be limiting. In the limit of very long RF electrodes, the RF electric field will be uniform along the trap axis (i.e. the field cross section in the plane orthogonal to the trap axis will be constant across the length of the axis), and the RF field null will form a line. In the case of shortened RF electrodes, the RF electric field  $\mathbf{E}_{RF}(z)$  will be non-zero at points along the trap axis further from the trap center, meaning that if the DC potential puts the trap site away from the geometric center of the trap, there will be some residual RF electric field, resulting in excess micromotion.

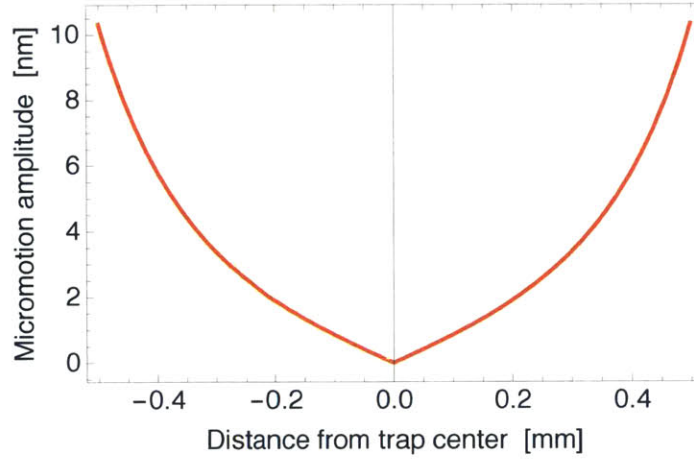
For a trap depth  $> 50$  meV, an RF amplitude of about 60 V is required on trap (at a frequency of 35 MHz). The resulting excess RF field due to edge-effects along the trap axis  $\mathbf{E}_{RF}(z)$  can be estimated using the model from Ref. [Hou08]). The z-component of this field plotted along the trap axis (z axis) is shown in Figure 6-2. The axial micromotion amplitude induced by this field,  $a_z(z)$ , as a function of distance from the geometric center of the trap, can then be calculated (plotted in Figure 6-3):

$$a_z(z) \approx \frac{e\mathbf{E}_{RF}(z) \cdot \hat{\mathbf{z}}}{m\Omega_{RF}^2} \quad (6.1)$$

The largest micromotion amplitude plotted in Figure 6-3 (half-way between the center of the trap and the end of the RF electrodes) is 10 nm, which suggests that excess micromotion due to shortened RF electrodes should be negligible. To get a feeling for what this means consider the following: one consequence of excess micromotion is broadening of the transition used for Doppler cooling, which reduces the rate at which the laser can cool the ion. To assure that the Doppler-cooling transition



**Figure 6-2:** Estimated RF field along the trap axis due to edge-effects ( $\mathbf{E}_{RF}(z) \cdot \hat{\mathbf{z}}$ ).

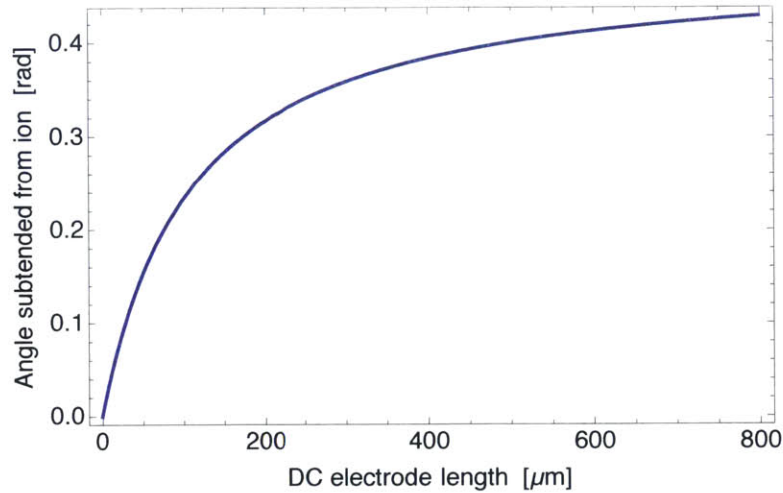


**Figure 6-3:** Estimated excess micromotion due to edge-effects along the trap axis.

(with  $\Gamma \approx 20$  MHz at 422 nm for  $^{88}\text{Sr}^+$ ), is broadened by less than 1% by micromotion, the micromotion amplitude needs to be  $< 40$  nm (assuming the RF frequency  $\Omega_{RF} = 35$  MHz). So, within the center half of the RF electrodes, broadening of the Doppler-cooling transition due to excess micromotion induced by edge effects is negligible.

### 6.2.2 Width of DC electrodes

The DC electrodes are also significantly shorter (DC width,  $w = 215 \mu\text{m}$ ) than in a typical trap. However, the angle subtended (in the x-direction) by the electrode from the ion saturates after  $\approx 2w$  (as can be seen in Figure 6-4) so this is not a major concern.



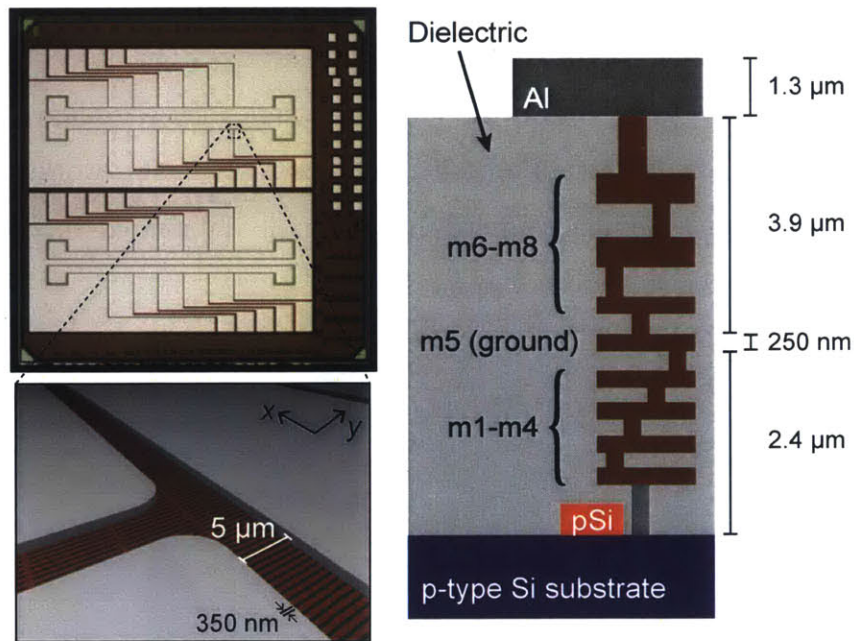
**Figure 6-4:** Angle subtended (in the x-direction) by the electrode from the ion as a function of electrode width.

## 6.3 Trap fabrication

Devices were fabricated on a  $3 \times 3 \text{ mm}^2$  die (Fig. 6-5) on a shared, 200-mm, multi-project wafer produced in a 90-nm CMOS process operated by IBM (9LP process designation). This process is primarily utilized for dense, high-performance digital circuits, and the trap die was one of many designs fabricated in parallel on the same wafer. The process allows for patterning of 8 copper interconnect layers, along with the top aluminum pad layer (right panel of Fig. 6-5). This  $1.3 \mu\text{m}$  thick pad layer was used for the trap electrodes. For several of the traps, an additional copper layer (m5) approximately  $4 \mu\text{m}$  below the aluminum layer's bottom surface and  $2 \mu\text{m}$  above




the silicon substrate was used to form a ground plane under the extent of the trap. Other traps were fabricated without a ground plane. Due to metal density constraints arising from chemical-mechanical polishing steps applied to these layers, this ground plane was patterned as a mesh of 600 nm strips separated by 350 nm along the  $x$  direction and 10  $\mu\text{m}$  along the  $y$  direction (see Fig. 6-5). Metal vias connect this copper ground plane to the center electrode of the trap through the upper metal layers m6–m8.



**Figure 6-5:** Die and process cross-section. A micrograph of the fabricated  $3 \times 3 \text{ mm}^2$  die is shown in the upper left panel. The lower left panel shows a perspective rendering of the top aluminum trap layer and the meshed ground plane in copper below, as designed; the gaps in the trap electrodes here are  $5 \mu\text{m}$ , and the ground mesh is formed of 600 nm wires with 350 nm gaps along  $x$  and  $10 \mu\text{m}$  gaps along  $y$ . A chip cross section is diagrammed at right, with approximate relevant dimensions labeled (“pSi” is polysilicon and metal interconnect layers are labelled m1 through m8). Vias shown between metal layers are only representative.

## 6.4 Trap characterization

After commercial foundry fabrication, trap chips are diced from the full die, to a size of approximately  $2.5 \times 1.2 \text{ mm}^2$ , and are bonded to a larger interposer. Trap testing was carried out in a cryogenic vacuum system at MIT Lincoln Laboratory that allows for variation of the trap-chip temperature. This system has been described in Ref. [SKC12]. Using a quarter-wave helical resonator, a 43 MHz RF signal of approximately 100 V amplitude was applied to the trap electrodes to produce radial trap frequencies of approximately 4–5 MHz. DC potentials of up to approximately 30 V were applied to the segmented control electrodes to produce an axial potential with frequencies near 1 MHz.  $^{88}\text{Sr}^+$  ions are loaded by accelerating precooled Sr atoms from a magneto-optical trap toward the ion trap where they are photo-ionized and Doppler cooled [SKC12]. Although not measured precisely in this work, loading efficiency into these traps is similar to more conventionally fabricated surface-electrode traps using the same loading method.



**Figure 6-6:** Micrograph of trap chip diced from die and mounted on the sapphire interposer of a cryogenic vacuum system. Aluminum wirebonds are used to make contact from the aluminum trap electrodes to the gold interposer leads. The chip is 2.5 mm long and 1.2 mm wide. The inset shows two ions trapped  $50 \mu\text{m}$  from the surface of the trap chip. The ions are approximately  $5 \mu\text{m}$  apart.

We noticed slightly more power dissipation in the foundry traps than in traps



fabricated from gold or niobium on sapphire for similar RF voltage amplitude[CS14], most likely due to higher dissipation in the metals or dielectrics. Ion lifetimes of more than an hour were observed in the presence of Doppler cooling light, equivalent to the best lifetimes seen in other traps measured in this vacuum system.

### 6.4.1 Dielectric breakdown

A possible limitation to the use of high-resolution CMOS processing is breakdown at large applied potentials, especially since typical ion trap voltage amplitudes are significantly higher than those used in CMOS electronics. However, for up to 200 V static bias applied, the leakage current was below 10 pA, and no sudden increase corresponding to a dielectric breakdown was observed. We performed these tests at room temperature (in a high-dielectric-strength fluid to prevent air breakdown), applying the potential between one of the RF electrodes and either the ground plane or one of the adjacent DC electrodes in both types of trap.

### 6.4.2 Laser-induced photo-effects

Traps without a ground plane displayed significant laser-induced photo-effects due to the excitation of carriers in the silicon by scattered light used for atom photoionization (PI, 405 nm) and ion Doppler cooling (422 nm). During trap loading, this manifested itself as variation of the RF voltage amplitude on the trap electrodes due to varying impedance of the trap when the 405-nm PI light was on. The effects on the trapping potential were visible as ion motion synchronized to the PI light switch state. We observed no photo-effects in traps with a ground plane.

### 6.4.3 Temperature-dependent nonlinearities

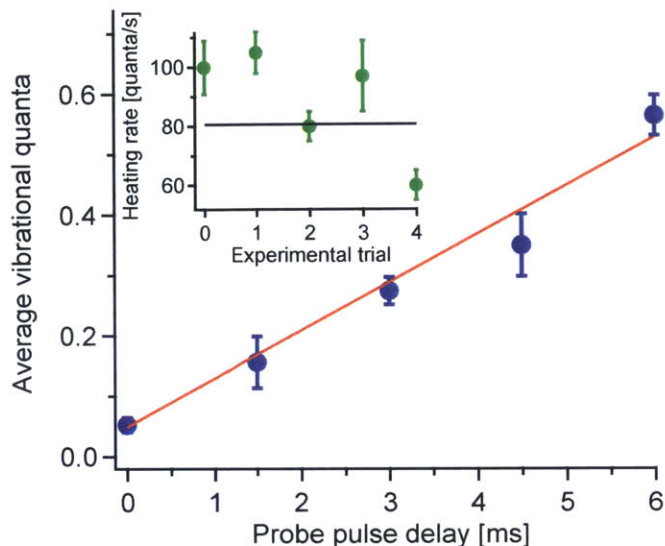
Traps without a ground plane also exhibit strong trap-temperature-dependent nonlinearities in the resonance response of the voltage-step-up resonator. A ground plane reduces RF leakage into the silicon substrate sufficiently to eliminate this effect, such that we observed stable trapping for chip temperatures from 300 K down to 8 K.

### 6.4.4 Motional heating

As discussed in Chapter 5, it is important to characterize the heating rate in potentially scalable trap technologies. We measured the heating rate of the 1.3 MHz axial vibrational mode using sideband amplitude spectroscopy on the  $S_{1/2} \rightarrow D_{5/2}$  transition (using the method described in Sections 2.4.3 and 5.5). Results of one such measurement are presented in Fig. 6-7 for a chip temperature of 8.4 K. Five measurements were recorded over a few days for nominally the same conditions; the average heating rate is 81(9) quanta/s. When scaled by  $1/d^4$  to compare traps of different sizes, this heating rate is lower than that reported in any other trap fabricated on a silicon substrate[LLC<sup>+</sup>09, SFB<sup>+</sup>10, Bri08, AHJ<sup>+</sup>12, WSGS12, WAF<sup>+</sup>13, SRW<sup>+</sup>14, NLK<sup>+</sup>14].

### 6.4.5 Excess micromotion

Excess micromotion (ion motion at the RF drive frequency) is caused by static electric fields that displace the ion from the RF null and can lead to ion heating. As described in Section 3.5.1, we compensate for this micromotion by applying an additional opposing static field. We observed typical stray field values on the order of 500 V/m, which appear stable over days. Although silicon oxide dielectric is exposed at the locations of gaps in the electrodes and may charge due to laser-induced photo-electron production, the stray field's stability suggests another cause. Wirebonds, which are asymmetric with respect to the ion location and also closer to the ion here than in the



**Figure 6-7:** Representative measurement of heating rate in a CMOS-foundry-fabricated ion trap. Average occupation of the axial mode of vibration in the linear trap is plotted as a function of delay time after preparation in the ground state at a trap temperature of 8.4 K. A linear fit (line shown) gives a heating rate for these data of 80(5) quanta/s where the uncertainty is due to statistical errors propagated through the fit. An average of five such measurements gives a heating rate of 81(9) quanta/s where the uncertainty is due to run-to-run variability. The inset shows all five measurements with a line indicating the weighted average value.

case of larger trap chips, may be responsible for the steady stray electric field. The use of through-silicon-via technology can eliminate wirebonds from the chip surface, as has recently been demonstrated for surface-electrode ion traps[GFH<sup>+</sup>14].

### 6.4.6 Light scatter

When compared to more typical single-metal-layer traps on sapphire substrates, these traps exhibit increased scatter of laser light, possibly due to higher as-deposited roughness of the aluminum layer. We examined the trap-electrode surface using atomic force microscopy and measured an RMS roughness of 35 nm, significantly larger than the 2 nm we have measured on single-metal-layer traps. Scatter from the surface can be reduced by focusing laser beams to a smaller diameter at the trap.

## 6.5 Conclusion

The experimental demonstration of an ion trap fabricated using standard CMOS processes opens up a tantalizing opportunity to make use of the wide ranging capabilities of modern CMOS systems to push trapped-ion quantum information systems to next levels of scalability and compactness. One could imagine making use of integrated silicon nanophotonic waveguides and grating couplers to produce quantum networks with ions and photonic interconnects on a chip [FLC<sup>+</sup>10, OKH<sup>+</sup>11]. A large-scale, integrated optics architecture might have single-mode waveguides distributing light to various locations in a dielectric layer in the same chip as the trap electrodes, and focusing grating couplers directing the light, through gaps in the trap electrodes to  $\mu\text{m}$ -scale focuses at the ion locations. This approach, which is explored further in Chapter 7, could allow practical integration of many couplers, avalanche photodiodes, electro-optic modulators, and control electronics into the same chip as the trap electrodes. It could potentially offer superior individual-ion addressing, efficient fluorescence read-out, and improved phase and pointing stability of control beams. Standardization of the foundry process also makes it possible for any group to repeatedly produce identical devices.

# Chapter 7

## A SIMD quantum Fourier transform

In Chapter 6, we described the development of a functional ion trap manufactured using a commercial CMOS process. This work opens the door to incorporating existing CMOS electronics and photonics technologies into ion traps, towards creating a fully-integrated trap module. Ideally, an integrated trap module could form the basis of a scalable system, versatile enough to implement a number of quantum algorithms. Because a single ion chain is limited to computations with about 10 ions [WMI<sup>+</sup>97], scalability requires an integrated trap module incorporate multiple trap zones, and hardware to transport ions between the zones. In this chapter we take a systems-oriented view as we investigate how quantum algorithms can be efficiently executed in such a network of distributed ion chains.

We begin by selecting an algorithmic primitive to serve as a building block for other algorithms. We choose the quantum Fourier transform (QFT) because it forms an inherently parallel basis for computational steps that tend to dominate the overhead of quantum algorithms (such as arithmetic or multi-qubit-controlled gates) [Sho95]. We show that the QFT primitive allows for efficient implementation of other important

algorithms, such as Shor’s factoring algorithm.

Next, we consider how to translate a quantum circuit for the QFT into a physically-realizable protocol of gate pulses and shuttling steps. We find that Flynn’s taxonomy [Fly72]—a classification scheme for computer architectures—provides a useful framework with which to map the QFT structure onto an efficient pattern for shuttling ions between different trap zones. We develop what Flynn would term a SIMD (Single Instruction Multiple Data) algorithm for the QFT, which greatly simplifies requirements for ion movement in a large array. In Chapter 8, this algorithm will guide the design of a CMOS-compatible, scalable trap module.

## 7.1 The discrete Fourier transform and classical parallelism

Before discussing the quantum Fourier transform (QFT), it is useful to consider its classical counterpart: the discrete Fourier transform (DFT), and how it is implemented efficiently in classical computer hardware. Later in this chapter, we develop an algorithm for the QFT analogous to the FFT algorithm (used to implement the DFT efficiently), and consider how the classical concept of “Single Instruction Multiple Data” (SIMD) processing can be interpreted in quantum information processing with ions. Hence, the logic for the rest of the chapter follows as a “quantum analog” to the material in this section.

### 7.1.1 Flynn’s taxonomy of parallelism

Flynn’s taxonomy is a classification scheme for computer architectures based on the interactions of their instruction and data streams [Fly72]. Flynn differentiates four types of architecture based on how many parallel data streams or instruction sets are possible in the architecture: Single Instruction Single Data (SISD); Single Instruc-

tion Multiple Data (SIMD); Multiple Instruction Single Data (MISD); and Multiple Instruction Multiple Data (MIMD). Computation proceeds sequentially in an SISD architecture: there are neither parallel instruction nor parallel data streams. In a SIMD architecture, on the other hand, operations are performed on multiple data streams in parallel, which can yield a significant reduction in computation time for problems in which the same operation is applied to multiple inputs.

Processors are mostly SISD, but nearly all include a SIMD extension to the instruction set. For example, all major Intel processors since the Pentium III have implemented streaming SIMD extensions (SSE).

### 7.1.2 The discrete Fourier transform

The discrete Fourier transform (DFT) acts on a finite sequence of  $N$  complex numbers:  $\{x_0, x_1, \dots, x_{N-1}\}$ , and is defined by its action on elements of the input sequence:

$$\text{DFT}_N : x_n \longrightarrow \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{2\pi i \left(\frac{n}{N}\right)k} x_k. \quad (7.1)$$

Evaluating the DFT naïvely from its definition on a digital computer requires  $N^2$  multiplications (where  $N$  is the transform length).

Like its continuous time counterpart, the discrete Fourier transform (DFT) converts temporal or spatial data into spectral data. The DFT is an essential component of digital signal processing, including image, audio, and video processing. Some problems, such as finding solutions to partial differential equations, can be more readily solved after transformation. Additionally, computations such as convolution or multiplication of large integers can be executed more efficiently using the DFT.

### 7.1.3 The fast Fourier transform algorithm

The fast Fourier transform (FFT) is a computer algorithm capable of evaluating the DFT in fewer time steps than those required to evaluate the definition directly (Equation 7.1). The result of the FFT is the exact DFT of the input data, but the time complexity required for the FFT is  $O(N \log N)$ , as opposed to  $O(N^2)$  working directly from the DFT definition. The algorithm recursively divides a DFT of size  $N = N_1 N_2$  into smaller DFTs of size  $N_1$  and  $N_2$  and takes advantage of common factors to speed the calculation. When  $N$  is a power of 2, the algorithm for the FFT is simplified: the  $N$ -point transform is divided into two  $N/2$ -point transforms, which in turn are divided into four  $N/4$ -point transforms, and so on, until division into 2-point transforms occurs. Microprocessor programs typically evaluate the FFT this way, known as a depth first traversal, because it streamlines the locations of successive memory reads (memory locality) [Sin67].

For algorithms in which computational steps can be completed more efficiently using the DFT (as described in Section 7.1.2), the FFT is typically used to evaluate the DFT. For example, the fastest known algorithms for multiplication of very large integers use the FFT [Fö9, DKSS13]

### 7.1.4 Hardware implementation

Processors in nearly all desktop and handheld machines today devote significant resources to computing the FFT for applications such as image, audio, and video processing. Because most of these devices are equipped with SIMD-capable processors, many generate SIMD code to exploit data-level parallelism inherent to the FFT [JMS86, Rod02]. By optimizing the performance of FFT algorithms on SIMD microprocessors, the FFT can be computed at the speed of, or faster than, state-of-the-art FFT libraries [Bla12, BWC13].



## 7.2 The quantum Fourier transform

The quantum Fourier transform (QFT) is effectively the same transformation as the DFT, expressed as an operation on quantum input states. Like the DFT, the QFT can be leveraged to perform other computations more efficiently. In particular, computational steps required for a number of important quantum algorithms (such as the modular multiplication steps in Shor's factoring algorithm [Sho95]) can be executed more efficiently using quantum Fourier transforms. In this sense, the QFT is a good choice for an algorithmic primitive.

### 7.2.1 Definition of the QFT

The quantum Fourier transform on  $N$  qubits is defined by the unitary transformation

$$\text{QFT}_N : |x\rangle \longrightarrow \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} e^{2\pi i \left(\frac{x}{N}\right)k} |k\rangle, \quad (7.2)$$

which is an interesting analogy to the discrete Fourier transform (Equation 7.1). If we choose  $N = 2^n$  for some integer  $n$ , the computational basis  $\{|0\rangle, |1\rangle, \dots, |2^n - 1\rangle\}$  can be expressed using the binary representation  $|x\rangle = |x_1 x_2 \dots x_n\rangle$ . So for example,  $|5\rangle \leftrightarrow |2^2 + 2^0\rangle \leftrightarrow |101\rangle$ . Binary fractions can be represented similarly:  $|\frac{x_1}{2^0} + \frac{x_2}{2^1} + \dots + \frac{x_n}{2^n}\rangle \leftrightarrow |0.x_1 x_2 \dots x_n\rangle$ . With this in mind, the action of the QFT can be expressed in a convenient form (as a tensor product):

$$\text{QFT}_{2^n} : |x_1 x_2 \dots x_n\rangle \longrightarrow \frac{\left(|0\rangle + e^{2\pi i(0.x_n)}|1\rangle\right) \otimes \left(|0\rangle + e^{2\pi i(0.x_{n-1}x_n)}|1\rangle\right) \otimes \dots \otimes \left(|0\rangle + e^{2\pi i(0.x_1 x_2 \dots x_n)}|1\rangle\right)}{\sqrt{2^n}} \quad (7.3)$$

In this form, we can see that the resulting state is not entangled, and that each qubit contains the lower  $k$  binary digits of  $x$ . This form is particularly helpful when deriving

a circuit for the QFT.

## 7.2.2 Circuit for the QFT

To develop a circuit for the QFT, let's begin with the simple case of 3 qubits. Our input state is  $|\psi_{\text{in}}\rangle = |x_1x_2x_3\rangle$ , and our desired output state is

$$|\psi_{\text{out}}\rangle = \frac{\left(|0\rangle + e^{2\pi i(0.x_3)}|1\rangle\right) \otimes \left(|0\rangle + e^{2\pi i(0.x_2x_3)}|1\rangle\right) \otimes \left(|0\rangle + e^{2\pi i(0.x_1x_2x_3)}|1\rangle\right)}{\sqrt{2^3}}. \quad (7.4)$$

If we apply a Hadamard gate

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (7.5)$$

to the third qubit, we obtain

$$H|x_3\rangle = \frac{|0\rangle + e^{2\pi i(0.x_3)}|1\rangle}{\sqrt{2}}, \quad (7.6)$$

which is the desired first qubit output. If  $x_3 = 0$ , we can also obtain the desired second qubit output by applying a Hadamard gate to the second qubit:

$$H|x_2\rangle = \frac{|0\rangle + e^{2\pi i(0.x_20)}|1\rangle}{\sqrt{2}} = \frac{|0\rangle + e^{2\pi i(0.x_2x_3)}|1\rangle}{\sqrt{2}}. \quad (7.7)$$


But if  $x_3 = 1$  we will need more than a Hadamard gate; we will also need to apply a phase shift gate  $R_\theta$ ,

$$R_\theta = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{bmatrix}, \quad (7.8)$$

with  $\theta = \pi/2$  to the second qubit, to produce:

$$R_{\pi/2}H|x_2\rangle = \frac{|0\rangle + e^{2\pi i(0.x_21)}|1\rangle}{\sqrt{2}}. \quad (7.9)$$

Hence, we need to apply controlled phase gate  $R_{\pi/2}$  targeted to the second qubit, controlled by the state of the third qubit. Similarly, for the first qubit, if we apply a Hadamard gate, followed by a phase rotation  $R_{\pi/2}$  controlled by the state of the second qubit, and another phase rotation  $R_{\pi/4}$  controlled by the state of the third qubit, the output will be  $\frac{|0\rangle + e^{2\pi i(0.x_1x_2x_3)}|1\rangle}{\sqrt{2}}$ , i.e. the desired third qubit output. Thus, up to reversing the order of the qubits, we have created a circuit to perform a 3-qubit QFT. This circuit is shown in Figure 7-1.




**Figure 7-1:** Circuit for a 3-qubit QFT, up to a permutation of the output qubits.

This approach can be generalized, resulting in the circuit for the QFT shown in Figure 7-2. This circuit requires  $n(n+1)/2$  gates (ignoring permutation of the output qubits) for a QFT on  $n$  qubits.

### 7.2.3 Approximate QFT

As the number of qubits  $n$  in the QFT grows, the smallest phase rotation angle in the circuit  $\theta = \frac{2\pi}{2^n}$  decreases, so that the corresponding rotation matrix approaches the identity. Eventually, we will reach a tolerance below which we cannot perform this rotation gate accurately, in a realistic implementation. In general, we define the approximate quantum Fourier transform (AQFT) as the QFT in the special case in



**Figure 7-2:** A circuit for the QFT, up to reversing the order of the output qubits. Note that in this figure, the notation  $R_k$  is used to refer to  $R_\theta$  with  $\theta = \frac{2\pi}{2^k}$  as defined in Equation 7.8.

which we eliminate the phase rotation operations  $R_\theta$  for  $\theta$  below a certain threshold,  $\theta \leq \frac{\pi}{2^m}$  for some  $m$ .

Repeatedly performing the AQFT (or the QFT) and then measuring the output register in the computational basis is a protocol to estimate the periodicity in probability amplitudes describing the input state of the register. Hence, the AQFT is useful for algorithms which involve periodicity estimations, such as Shor’s factoring algorithm. Because fewer gates are required in the AQFT, the AQFT can actually yield better results than the full QFT for periodicity estimation in the presence of decoherence [BEST96].

While we focus on the exact QFT in this chapter, much of our discussion will also apply to the AQFT.

## 7.2.4 Computing with the QFT

The interaction sequence of the QFT can be generalized by allowing variable rotation angles for each gate. For example, if the angle for every controlled rotation gate in the QFT of  $|x\rangle$  is multiplied by an odd number  $r$ , then the output state changes

(from that shown in Equation 7.3) to

$$\frac{\left(|0\rangle + e^{2r\pi i(0.x_n)}|1\rangle\right) \otimes \left(|0\rangle + e^{2r\pi i(0.x_{n-1}x_n)}|1\rangle\right) \otimes \dots \otimes \left(|0\rangle + e^{2r\pi i(0.x_1x_2\dots x_n)}|1\rangle\right)}{\sqrt{2^n}}, \quad (7.10)$$

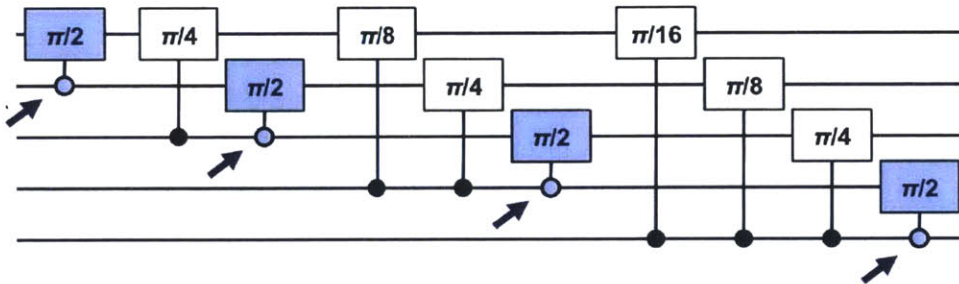
which is the output state that would be obtained if the initial state was  $|rx\rangle$ . Hence, this QFT circuit with each rotation angle multiplied by  $r$  performs a combined multiplication-by- $r$  and QFT. In a similar fashion, the addition of a classical number to a quantum superposition can also be performed using the QFT [KJL<sup>+</sup>08], and hence can be parallelized. Saeedi et al. [SP13] used this idea to design a linear-depth, quadratic-size quantum circuit to implement an  $n$ -qubit Toffoli gate. Their circuit consists of elementary two-qubit controlled-rotation gates around the  $x$  axis, and so can be constructed using QFT circuits. The  $n$ -qubit Toffoli gate is a key part of many quantum algorithms, such as quantum error correction codes [NC02].

More generally, the structure of the QFT (or the AQFT) can be used to parallelize quantum computations, including arithmetic and multi-qubit-controlled gates (as we have seen). Because for many quantum algorithms, the main overhead arises in these computational steps [Sho95], being able to build up algorithms from QFT circuits is a powerful tool.

### 7.2.5 Parallelizing the QFT

In general, finite coherence times mean that parallelization is an important way of minimizing the total computation time. The ability to perform parallel operations is also crucial for quantum error correction; without parallel operations, errors will accumulate in the circuit more quickly than error correction can handle [NC02]. The QFT is inherently parallel in the sense that every qubit receives a portion of the same instruction set, or in other words: a subset of the single, fixed sequence of controlled phase gates that are applied to each qubit.


The QFT requires a single entangling gate (controlled phase gate) between every pair of qubits, amounting to  $\frac{1}{2}(n^2 - n)$  gates for an  $n$ -qubit QFT. The bottleneck for parallelizing the QFT is a restriction on the ordering of these entangling gates (because of operations that do not commute). Specifically, each qubit can be the target of a controlled rotation only after all gates in which it acts as a control are complete. Or, explained another way, each qubit must interact with each higher-order (more significant) qubit (in any order) before interacting with every lower-order qubit. This creates a “staircase” structure (see Figure 7-3) with a depth of  $n - 1$  operations which is difficult to reorder [MN98]. This staircase restriction also applies to the AQFT.



**Figure 7-3:** Circuit for the 5-qubit QFT (Hadamard gates not shown) with the “staircases” through which gates are not commutable (see description in text) indicated in blue. In this diagram, the label  $\pi/2$  indicates a phase rotation gate  $R_{\pi/2}$  as defined in Section 7.2.2.

### 7.2.6 Implementing the QFT with a single chain of ions

The general layout considered in this chapter is a network of distributed ion chains, as shown in Figure 7-4. It is a 2D array of ion chains of height  $H$  (i.e.  $H$  ions per chain) and overall width  $W$  (meaning that there are  $W$  ion chains in total). However, the simplest implementation of the QFT involves only ions confined in one shared trap, forming a single linear chain ( $W = 1$ ). In this configuration the staircase limit (corresponding to a circuit depth of  $n - 1$ ) can be saturated.



**Figure 7-4:** Schematic showing the general distributed-ion-chain layout we consider in this chapter. It is a 2D array of ion chains of height  $H$  and overall width  $W$ . Each of the grey rectangles labelled 1 through  $W$  correspond to a single trap site (electrodes not shown) with  $H$  ions shown in blue; connection between chains is also not shown.

Within a single linear chain of  $H$  ions ( $H$  qubits), a full set of controlled rotation operations controlled by or targeting a single qubit can be performed with a constant-depth pulse sequence. As we show in detail in Appendix B.1, four entangling gates are required regardless of the chain size. In effect, each of these four entangling gates is equivalent to  $H - 1$  two-qubit controlled-rotation gates. This means that an  $n \leq H$  qubit QFT can be performed with  $n - 1$  global operations, saturating the staircase limitation on circuit depth.

### 7.3 Implementing the QFT in a scalable architecture

The number of ions  $H$  in a single ion chain is limited to something on the order of 10 ions [WMI<sup>+</sup>97]. Hence, a scalable system will require quantum algorithms to be distributed across multiple chains acting as independent “quantum co-processors” (as discussed in Section 1.2.2). In order to perform a QFT of size  $n > H$ , we must move to an architecture with  $W = \lceil n/H \rceil$  distributed ion chains, and incorporate communication steps between these chains into our QFT algorithm.

In Section 1.2.2, we introduced the two major categories of scalable architecture proposed to connect distributed ion chains. The two architectures can be differentiated by the means of communication employed to transmit information between ion chains; the different modes of communication are:

- **Communication mediated by physically transporting ions:** information is carried from one ion chain to another ion chain by ions that are physically transported between the chains. This mode of communication corresponds to an architecture known as a “quantum charge-coupled device” (QCCD) [KMW02].
- **Communication mediated by photons:** the quantum state of one ion qubit is converted to the quantum state of a photon that then carries this information through space to another trapped ion qubit. This mode of communication corresponds to an architecture known as a “quantum network” [CZKM96].

Currently, ion-photon mapping is relatively inefficient, with the mean connection times ranging from 5 to 250 ms [MRR<sup>+</sup>14], considerably slower than the  $\sim 10 \mu\text{s}$ -scale connection times possible via shuttling [WZR<sup>+</sup>12]. The photon-mediated communication approach has the advantage of being able to transmit quantum information over long distances, but for building a computer architecture, communication based on shuttling ions is clearly preferable.



Shuttling imposes two primary limitations on the overall coupling network. First, as qubits are passed physically, communication is limited by physical proximity (i.e. the network is limited to nearest-neighbor with dimension  $d \leq 3$ ). Generating entanglement between qubits across the system results in a shuttling depth of  $\sim n^{1/d}$  simply because it takes that many steps for them to meet. Second, only the equivalent of the quantum *SWAP* gate can be performed between independent traps (i.e. only permutations of qubits can occur between traps).

Because of the staircase limitation (discussed in Section 7.2.5), the circuit depth of the QFT is already at least  $\sim n$ . So as long as shuttling does not asymptotically increase the depth of the QFT, or the number of gates, shuttling will not change the overall (gate + shuttling) depth. Therefore, in this section we will focus on the QCCD architecture, in which ions are confined to a number of modular “trap zones”, and gates are performed between all possible qubits by shuttling ions between the different zones. We will determine the precise overhead added by shuttling and show that we can achieve a circuit depth of  $\sim 2n$  with a simple algorithm.

### 7.3.1 Notation

Before proceeding to discuss the overhead incurred through shuttling, we should fix our notation: our  $n$ -qubit system consists of  $W$  individual chains designated  $C_0$  through  $C_{W-1}$ , each representing an array of  $H = n/W$  qubits. Qubits are designated  $q_0$  through  $q_{n-1}$  with  $q_0$  being the most significant bit (MSB) of the input state. Further, each qubit is described as being in a “binary” (input) or “Fourier” (output) basis (indicating whether they are in the stage of controlling or being targeted by the QFT’s controlled rotation gates, respectively). Qubits switch from binary to Fourier after all of the rotation gates in which they are the control qubit have occurred, indicating a “step” down the staircase.

We further describe each chain as having  $H_s$  shuttling zones, from which ions

can be simultaneously shuttled into and out of that chain. Each zone is associated with a unique shuttling set ( $D_0$  through  $D_{H_s-1}$ ) consisting of the  $W$  qubits from the corresponding shuttling zones in each chain. Cyclic permutation operators  $A_{\alpha\beta\gamma\dots}$  act on an individual shuttling set  $D_k$ , representing qubits shuttled between chains  $C_\alpha$ ,  $C_\beta$ ,  $C_\gamma$ , etc. in that zone. Finally,  $V$  denotes the maximum number of qubits that can be permuted in a single cycle.

We will refer to the following figures of merit to quantify the efficiency of algorithms we describe for implementing the QFT:

- Shuttle depth  $L$ : the number of global shuttling steps
- Number of shuttles  $N_{shuttles}$ : the total number of times a single qubit is shuttled
- Number of gates  $N_{gates}$ : the total number of controlled-rotation sequences
- Number of swap gates  $N_{swaps}$ : the total number of swap gates

### 7.3.2 Shuttling overhead – performance bounds

We assume that each of  $W$  chains holds  $H$  ion qubits (so that  $n = HW$ ), of which  $H_s$  can be shuttled in a single time step. We can immediately construct lower bounds for the overhead from shuttling by counting the unique pairings of qubits within each trap (remember that the QFT consists of  $(n^2 - n)/2$  two-qubit gates, each between a unique pair of qubits), resulting in the following constraints:

- Within a single  $H$ -qubit chain, there exist  $(H^2 - H)/2$  unique qubit pairings. Therefore across a set of  $W$  distributed chains, there exist  $W \cdot (H^2 - H)/2$  pairings prior to any communication between chains.
- When a single ion is shuttled to a new chain, at most  $H - 1$  new pairings are created.

- At most  $H_s \cdot W$  qubits can be shuttled in a given time step.

Combining these constraints, in order to perform the  $(n^2 - n)/2$  gates of the QFT we obtain the bounds:

$$M = (n^2 - n)/2(H - 1) - n/2 \approx n^2/2H, \quad (7.11)$$

$$L = M/H_s W \approx n/2H_s, \quad (7.12)$$

where  $M$  is the total number of qubit-shuttles and  $L$  is the depth of shuttling steps. Notice that  $H$  appears in the denominator of Equation 7.11. This is due to the result described in Section 7.2.6: for a fixed  $n$ , increasing  $H$  allows each controlled-rotation gate to affect more qubits, decreasing the number of required gates (for  $H \leq n$ ).

### 7.3.3 Simple QFT algorithm for $H_s = 1$ , $V = 2$

As we will show, it is possible to create a shuttling protocol for the QFT which asymptotically achieves the lower bound on shuttling depth (Equation 7.12). For simplicity, we will assume that only one ion per chain can be shuttled in a given time step ( $H_s = 1$ ), and that permutations are limited to pairs of qubits ( $V = 2$ ).

In essence, the QFT shuttling procedure consists of alternating “even” and “odd” permutations  $P_0(m)$  and  $P_1(m)$  acting on  $D_0$  (the single shuttling set) and affecting a subset of  $m$  adjacent chains (always beginning with chain  $C_0$ ):

$$P_0(m) \triangleq \prod_{\alpha=0}^{\lfloor m/2-1 \rfloor} A_{2\alpha, 2\alpha+1}, \quad (7.13)$$

$$P_1(m) \triangleq \prod_{\alpha=1}^{\lfloor m/2-1 \rfloor} A_{2\alpha-1, 2\alpha}. \quad (7.14)$$

The permutations  $P_0(m)$  and  $P_1(m)$  depend on  $m$ , the number of “active chains”, i.e. the chain number corresponding to the next chain a Fourier ion will be swapped

into. Initially,  $m = 1$  (only the first chain is being swapped). If we have  $k$  chains filled with Fourier ions, then the next Fourier ion will be swapped into chain  $k + 1$  and so  $m = K + 1$ .

Complete pseudocode for this procedure can be found in Appendix B.2, Algorithm 1. Here, we provide a brief description of how the algorithm works.

Initially, qubits are arranged sequentially;  $C_0$  contains qubits  $q_{H-1}, \dots, q_1, q_0$ ,  $C_1$  contains qubits  $q_{(2H-1)}, \dots, q_{H+1}, q_H$ , etc. Or to generalize,

$$C_k = \{q_{(k+1)H-1} \dots q_{kH+1} q_{kH}\}. \quad (7.15)$$

The algorithm is then broken into three stages. The outline of our strategy is to recursively shuttle each binary-space qubit until it has passed through chains containing every more significant qubit (which it must target with a controlled rotation) and can be relabeled as in the “Fourier” basis. It is then swapped into the next available non-shuttling zone (in the process “dislodging”, or swapping into  $D_0$ , a new binary qubit). The next most significant binary qubit follows the same path, and so is able perform a controlled rotation gate on the new Fourier qubit and can then be swapped into the next non-shuttling location, and so on.

We begin (stage 1, below) by filling the chains  $C_0, C_1 \dots$  with new Fourier-space qubits sequentially, using the permutation  $P_0(m)P_1(m)$  using  $m$  to signify the chain number of the next available spot: no shuttling needs to occur in chains after  $m$  until all of the chains are filled. However, once the final chain is full ( $m \geq W$ ), we have only transformed  $W(H - 1)$  of the  $HW$  qubits (with the remaining binary qubits being left in the shuttling set  $D_0$ ). In order to transform the  $W$  remaining binary qubits, we must begin again at chain  $C_0$  and continue to use chain  $C_{(m \bmod W)}$  in place of chain  $C_m$ . In a 1D array of traps, this requires a unique shuttling stage to refresh our qubit arrangement (stage 2, below), before continuing with the algorithm (stage 3). Note that if we could instead arrange our traps in a “loop” on a 2D plane, we

could immediately continue to trap  $C_0$ , eliminating the refreshing stage.

With each permutation  $P_0(m)$  or  $P_1(m)$ , half of the qubits involved ( $m/2$ ) are shuttled in each direction. As each qubit will pass through every necessary chain sequentially after  $C_0$ , it is sufficient to only perform gates controlled by that qubit when it is moving “forward” (or toward higher trap indices). Therefore, in each iteration, there will be  $m/2$  parallel rotation gates: those in chains with even indices after  $P_0$ , and odd indices after  $P_1$ . Further, as no chain will require either a controlled rotation or a *SWAP* gate in two consecutive iterations, we can always postpone the controlled rotation in a chain in which a *SWAP* occurs, so that the only a single gate occurs per shuttling iteration and the *SWAP* gates do not add latency to our system.

To calculate the shuttling and circuit depths of this algorithm, we tally the required gates and shuttles step by step, as described below.

- **Stage 0:** Prior to any shuttling, we perform the available controlled rotations in chain  $C_0$ , so that the first  $H$  qubits are moved to Fourier space.
- **Stage 1:** Pseudocode for stage 1 begins on line 22 of 1. Initially (after stage 0), chain  $C_0$  is filled with Fourier-space qubits, and so the next available spot is in chain  $C_1$  ( $m = 1$ ). The repeated permutation operation  $P_0(m)P_1(m)$  shuttles each sequential qubit, once it is “dislodged”, in a path toward trap  $C_0$ , and subsequently through each consecutive trap. Starting with trap  $C_0$ , the qubit performs a controlled rotation gate on the non-shuttled Fourier-basis qubits in each trap it passes through. When the qubit has been permuted through traps with all higher-order qubits (so that it can be considered to be in Fourier basis), it is swapped out of  $D_0$  and into a non-shuttling trap zone, and a new binary qubit is swapped down (dislodged). The range variable  $m$  is increased whenever each of the  $H - 1$  spots in  $C_m$  have been filled with Fourier qubits, until  $m = W$

This stage performs a QFT on qubits  $q_0$  through  $q_{n-W}$ , as well as a subset of the controlled rotations of these qubits by the remaining binary qubits. It takes  $2H - 1$  shuttling iterations to fill the  $H - 1$  available spots in each chain, so that the total depth of this stage is  $W(2H - 1)$ . Each iteration involves an application of either  $P_0(m)$  or  $P_1(m)$ , parallel controlled rotation gates in every other chain from  $C_0$  to  $C_m$ , and a *SWAP* gate in chain  $m$  roughly every other iteration. In total, this stage amounts to  $N_{shuttles} = W^2(2H - 1)/2$ ,  $N_{gates} = W^2(2H - 1)/4$ , and  $N_{swaps} = W(H - 1)$ .

- **Stage 2:** Pseudocode for stage 2 begins on line 38 of 1. After stage 1, there are no new trap zones into which to swap Fourier qubits. In order for subsequent qubits to continue to pass through each chain of more-significant qubits, we must continue to iterate  $P_0P_1$  until the next Fourier qubit can be swapped into chain  $C_0$ . As no qubit is transformed to Fourier space in this stage,  $m$  is constant. This stage therefore consists of  $W$  global permutations ( $P_0(W)$  or  $P_1(W)$ ), again with gates in every other chain. In total, this amounts to  $N_{shuttles} = W^2$ ,  $N_{gates} = W^2/2$ , no *SWAP* gates, and a shuttle depth of  $W$ .
- **Stage 3:** Pseudocode for stage 3 begins on line 48 of 1. Here we finally restart our binary qubit permutation scheme at  $C_{(m \bmod W)} = C_0$ , performing a QFT on the remaining qubits  $q_{n-W}$  through  $q_{n-1}$ . After stages 1 and 2, we are left with  $W$  qubits in binary space (the number of qubits in shuttling zones, or,  $||D_0||$ ), requiring the reuse of  $W/(H - 1)$  chains, so that stage 3 has a depth of  $(2H - 1)W/(H - 1) \approx 2W$ . The first  $W$  iterations again require global permutations. It turns out, however, that after  $W$  we can decrease the range with each iteration until it reaches chain  $C_{W/(H-1)}$ . In total, the stage involves  $N_{shuttles} \approx 3W^2/2$ ,  $N_{gates} = 3W^2/4$ , and  $N_{swaps} = W$ .

Combining these steps, we see that the total QFT procedure entails  $\approx W^2(H+5/2)$

shuttling steps,  $W^2(2H + 5)/4$  controlled rotation gates,  $HW = n$  *SWAP* gates, and has an overall shuttle depth of  $L = 2W(H + 1) \approx 2n$ . So this algorithm achieves the lower bound on shuttling overhead (up to a constant factor).

## 7.4 A SIMD QFT algorithm

The shuttling overhead for the simple QFT algorithm described in Section 7.3.3 is close to the lower bound derived in Section 7.3.2, but it does not take into account another aspect of shuttling: the classical control and hardware required for ion movement. The ability to perform arbitrary shuttling steps requires a complex network of control electronics and computed voltage waveforms. Hence, even with a time-efficient algorithm, the physical hardware requirements may create a bottleneck to scaling-up.

In this section, we adapt the simple algorithm described in Section 7.3.3 to streamline the classical hardware required to realize the QFT algorithm in a distributed network of ion chains. We find that Flynn’s taxonomy [Fly72]—a classification scheme for computer architectures, introduced in Section 7.1.1—provides a useful framework for analyzing different shuttling patterns. If we can limit our QFT protocol to a minimum of shuttling sequences (“instructions”) that can be applied to all shuttling zones (“data”), it greatly simplifies the hardware involved. We develop what Flynn might term a SIMD (Single Instruction Multiple Data) protocol for the QFT, which reduces classical-control overhead significantly for ion movement in a large array.

### 7.4.1 SIMD algorithm for $H_s = 1$ , $V = 2$

In the algorithm described in Section 7.3.3, the entire QFT is performed via the repeated application of the permutation  $P_0(m)P_1(m)$ , used to recursively shuttle qubits until they are shifted into Fourier space and swapped into chain  $m$ . The algorithm

does not effect the qubits in chains  $> m$ , so we can imagine instead performing the permutation  $P_0(W)P_1(W)$ , and reorganizing our input arrangement to compensate for the additional permutations qubits in chains  $> m$  will experience with each step. This would vastly simplify a hardware implementation because it would require just two forms of shuttling, which would involve identical ion movements across all the trap zones. In other words: the movements of all ions would be controlled by repeatedly applying just one instruction  $(P_0(W)P_1(W))$ —our QFT protocol would become SIMD!

The permutation operators  $P_0$  and  $P_1$  can be made constant throughout the algorithm, so the shuttling algorithm for the QFT can be considered SIMD. The new, constant permutation operators are defined by fixing  $n = W$  in Equations 7.13 and 7.14:

$$P_0 \triangleq \prod_{\alpha=0}^{\lfloor W/2 \rfloor} A_{2\alpha, 2\alpha+1}, \quad (7.16)$$

$$P_1 \triangleq \prod_{\alpha=1}^{\lfloor W/2-1 \rfloor} A_{2\alpha-1, 2\alpha}. \quad (7.17)$$

Then, the set of operators  $\{(P_0P_1)^k; k = 0, 1, 2, \dots\}$  forms a cyclic permutation group of order  $W$ . As such, by arranging the first  $W$  most significant qubits ( $q_0$  through  $q_{W-1}$ ) correctly in the single shuttling set  $D_0$ , each qubit will consecutively shuttle through traps  $C_0, C_1$ , etc. as in the previous procedure. Further, as each qubit is swapped out of  $D_0$  sequentially, the qubits swapped into their place will follow their path and will also pass through the consecutive chains in the order they are swapped into  $D_0$ . These qubits are therefore arranged consecutively in the order that they will be swapped.



Thus, an algorithm for what we call the SIMD QFT results from simply replacing  $P_0$  and  $P_1$  in the steps previously defined in Algorithm 1 and rearranging the initial ion placement. The first few shuttling steps of this algorithm are shown in Figure 7-





---

5. Remarkably, the depth of and number of gates in the circuit are unchanged. The SIMD QFT greatly simplifies the hardware required to implement the QFT, and in turn, other quantum algorithms which employ the QFT, without adding any depth or gates to the algorithm (although the total number of shuttling steps is increased).



Initial ion configuration:





Permutation 1:

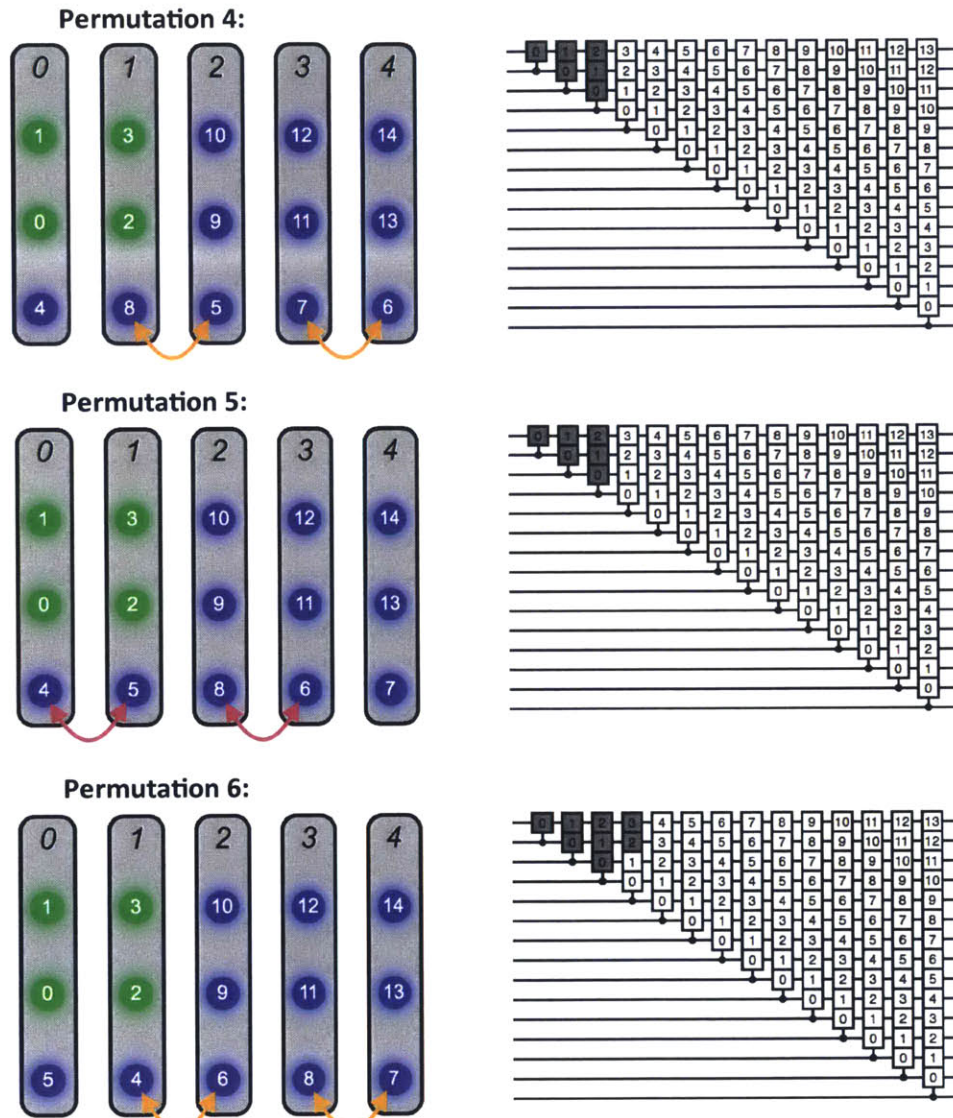


Permutation 2:



Permutation 3:





**Figure 7-5:** The first six shuttling steps for a SIMD QFT with  $H = 3$ ,  $W = 5$ ,  $H_s = 1$  and  $V = 2$  (see text for details). Rotation gates and in-chain *SWAP* steps are not shown, but progress through the QFT circuit is indicated schematically on the right as completed gates turn from white to grey. Input (binary) ions are shown in blue, and transformed (Fourier) ions are shown in green. Permutation  $P_0$  is shown in pink (applied to chains 0, 1, 2, and 3 simultaneously), and permutation  $P_1$  is shown in orange (applied to chains 1, 2, 3, and 4 simultaneously). In this simple case (of small  $W$ ), parallelism is apparent in the four ion movements occurring in each shuttling step.

## 7.5 Conclusion

In this chapter we took a distinctly systems-oriented approach, exploring how quantum algorithms can be implemented efficiently in a scalable trap array. We ascertained that the quantum Fourier transform is a good choice of algorithmic primitive, and devised protocols to implement the QFT efficiently in both a single ion chain, and a scalable trap array. We discovered that we can vastly simplify the hardware required for shuttling without adding any depth or gates to the algorithm by modifying the permutation operators so that the ion movements are universal; or, in other words, by making the overall QFT shuttling protocol SIMD.

In this work, we have only considered the simple case where each ion chain contains a single shuttling zone, and connections to two other chains ( $H_s = 1$  and  $V = 2$ ), but there is potential to further reduce shuttling overhead for ion chains with more shuttling zones and more connections to other chains. The shuttling depth  $L \approx n/2H_s$ , indicating that adding shuttling zones ( $H_s > 1$ , i.e. allowing more than one ion per chain to be shuttled) would reduce shuttling overhead. Switching to a 2D arrangement of traps (such as a square grid) would allow us organize our trap in a loop, eliminating the need for stage 2 of the QFT algorithm (the refreshing stage). More exotic permutations made possible by increasing the connections at each chain ( $V > 2$ , for example by creating an array in which three chains intersect at a point) may also reduce the shuttling depth.

# Chapter 8

## An integrated trap module

As discussed in Chapter 6, we realized stable ion confinement in a trap manufactured using a commercial complementary metal-oxide-semiconductor (CMOS) foundry process. The trap functions favorably, with electric-field noise comparable to that seen in other surface-electrode traps of similar size. A primary motivation for developing CMOS traps is the ability to leverage extensive existing libraries of integrated circuits for digital logic, memory, and photonics technology; the epitome of which would be the development of a scalable trap module incorporating components for local optical and electronic control.

In Chapter 7, we showed that the quantum Fourier transform (QFT) can form a useful building block from which parallel circuits for other quantum algorithms (such as Shor's factoring algorithm) can be constructed. We investigated how best to map the QFT onto a planar trap array, ultimately developing a SIMD (Single Instruction Multiple Data) algorithm for the QFT involving successive shuttling and gate steps.

In this chapter, we connect the results of Chapters 6 and 7, exploring their implications for the design of a scalable trap module. We discuss potential hardware for such a module, focusing on CMOS-compatible technologies. Our goal is to provide a perspective on what is possible with currently-realizable surface-electrode traps and

integrated optics. We summarize the hardware constraints arising from the QFT, finding that requirements for ion movement, qubit manipulation (gates), quantum state detection, and motional cooling are simplified for a SIMD algorithm. Overall, we find that the prospects for a modular trap architecture based on the QFT are good.

## 8.1 Prototype system

It will be fruitful to review one of the most advanced prototype systems available today: the linear ion trap experiment (“LinTrap”) at the University of Innsbruck, where many quantum protocols have been realized (see, for example, Refs. [SBM<sup>+</sup>11, MSB<sup>+</sup>11, BR12]). Details of prior work implementing the QFT with trapped ions in this prototype system [SNM<sup>+</sup>13] will guide us as we estimate the physical requirements for a prospective, scalable architecture.

At a high level, the QFT demonstration in the Innsbruck system relied on accomplishing the following elements:

- Qubit initialization
- Global Mølmer-Sørensen (MS) gate operations [SM98, MS98]
- Global x- and y-rotations
- Controlled z-rotations individually targeting each ion
- Ion movements: splitting, shuttling, recombination
- Readout/measurement

Some of these requirements are specific to the Innsbruck experiment, but many of them will also apply to our proposed module.

<b>Qubits</b>	$^{40}\text{Ca}^+$ ions	up to 14 ions in a single chain
<b>State initialization</b>	Sideband cooling Raman cooling	$\bar{n} = 0$ with 99.5% probability in $\sim 2$ ms $\bar{n} \approx 1$ in $\sim 400$ $\mu\text{s}$
<b>Quantum gates</b>	$R(\pi, \phi)$ $Z(\pi)$ $MS(\pi/2, \phi)$	gate fidelity >99%, gate time 20 $\mu\text{s}$ gate fidelity >99%, gate time 20 $\mu\text{s}$ gate fidelity 99% (2 ions), 94% (5 ions), gate time 45 $\mu\text{s}$
<b>Readout</b>	PMT CCD	measurement time 400 $\mu\text{s}$ measurement time 8 ms
<b>Coherence time</b>	$T_1$ $T_2$	1.1 s 5 to 50 ms

**Table 8.1:** Capabilities of the ion trap experiment at the University of Innsbruck, including durations and fidelities of various operations. See Ref. [Wan12]

### 8.1.1 Capabilities of the Innsbruck experiment

The capabilities of the ion trap experiment at the University of Innsbruck (“LinTrap”; see Ref. [SNM<sup>+</sup>13]) are summarized in Table 8.1.

As indicated in Table 8.1, the fidelity of the MS gate depends strongly on the number of ions. When generating N-particle Greenberger-Horne-Zeilinger (GHZ) states, the fidelity of the  $MS(\pi/2)$  gate ranges from 99% for two qubits to about 51% for 14 qubits, as shown in Table 8.2. The errors are primarily attributed to state initialization, inhomogeneous illumination of the quantum register, and correlated dephasing [MSB<sup>+</sup>11]. As these factors will no doubt also affect our proposed module, this may create an upper bound on the number of ions we can store in each trap zone ( $H$ ).

### 8.1.2 Laser requirements of the Innsbruck system


The University of Innsbruck experiment is based on  $^{40}\text{Ca}^+$  qubits. The  $^{40}\text{Ca}^+$  ion has a  $\Lambda$  atomic structure, similar to the  $^{88}\text{Sr}^+$  ion, and its qubit transition is at 729 nm

Number of ions	Fidelity [%]
2	99.50(7)
3	97.6(2)
4	97.5(2)
5	96.0(4)
6	91.6(4)
8	84.7(4)
10	67.0(8)
12	53.3(9)
14	56.2(11)

**Table 8.2:** Table showing how the fidelity of generating N-particle GHZ states with the maximally-entangling  $MS(\pi/2)$  interaction depends on the ion number (data from the ion trap experiment at the University of Innsbruck). See Ref. [MSB<sup>+</sup>11].

with a lifetime of about 1 second. For this ion, the relevant atomic transitions are shown in Figure 8-1 (compare to Figure 2-3 for  $^{88}\text{Sr}^+$ ). The laser intensities required to implement quantum protocols like the QFT in this system are shown in Table 8.3. We will refer to this data later in the chapter to make estimates about our prospective system.





**Figure 8-1:** Level scheme for  $^{40}\text{Ca}^+$  (only relevant transitions are shown).

Laser purpose	Wavelength	Transition	Max. Intensity
Qubit manipulation Sideband cooling	729 nm	$4S_{1/2} \longleftrightarrow 3D_{5/2}$	$\approx 3.5 \times 10^7 \text{ W/m}^2$
State initialization Readout Doppler cooling	397 nm	$4S_{1/2} \longleftrightarrow 4P_{1/2}$	$\approx 5.0 \times 10^2 \text{ W/m}^2$
Repumping	866 nm 854 nm	$3D_{3/2} \longleftrightarrow 4P_{1/2}$ $3D_{5/2} \longleftrightarrow 4P_{3/2}$	$\approx 1.6 \times 10^3 \text{ W/m}^2$

**Table 8.3:** Required lasers, their purpose and the maximum intensity required at each ion. See Ref. [Rie05].

## 8.2 Relevant requirements for the SIMD QFT

In this section, we construct a set of hardware requirements based on the SIMD algorithm we developed for the QFT in Section 7.4. These requirements are discussed below, organized by function.

### 8.2.1 Requirements for ion movement

Shuttling is an essential facet of the SIMD QFT algorithm, and hence is central to any hardware implementation. An ion is transported between two different trap zones by applying a sequence of voltages (typically called “waveforms”) to the segmented DC electrodes (“control electrodes”) spanning the two zones. The constant permutation operators  $P_0$  and  $P_1$  defined in Equation 7.16 and 7.17 correspond to voltage waveforms that are the same in every trap zone, significantly simplifying the requirements for ion movement.

For the SIMD QFT algorithm, the shuttling depth  $L$  scales with the circuit depth, so for this to be executed efficiently, it is essential that the shuttling times be comparable to the gate times. Additionally, heating of the motional state of the ions during transport (and potentially during splitting of the ion from the initial chain) should be minimized to prevent gate errors and reduce re-cooling time.

In summary, the SIMD QFT algorithm requires that control voltages be designed to implement the permutation operators  $P_0$  and  $P_1$  with duration on the order of a gate time, and with minimal energy gain.

### 8.2.2 Requirements for qubit manipulation (gates)

The SIMD QFT algorithm entails application of the Hadamard gate (single qubit gate), application of controlled-rotations on up to  $H$  qubits, and application of the quantum *SWAP* gate on qubits within a chain. In terms of gates readily implemented

in a practical ion system, these requirements can be framed equivalently as: the ability to perform an MS gate (an entangling gate acting on 2 or more qubits), and the ability to perform arbitrary, parallel single-qubit rotations.

Physically, quantum gates are applied to the ions in the form of laser pulses, requiring optical hardware: laser beams capable of individually addressing each ion, modulators capable of generating pulse sequences, and waveguides or optical fibers to transmit laser light to the trap zones. In the SIMD QFT algorithm, gates are only performed on ions located within the trap zones (not ions being shuttled between zones), so laser beams need only address these ion locations.

### 8.2.3 Requirements for quantum state detection

The result of the algorithm is readout via fluorescence emitted by the ions (see Section 4.1). Hence, it is essential that the fluorescence emitted by each individual ion be collected efficiently, and in such a way that it can be differentiated from fluorescence emitted by adjacent ions. The signal-to-noise ratio of light captured from each ion must be large enough to permit quantum state detection of the individual ion within a reasonable time frame. For typical background light levels, measurement with greater than 99% fidelity requires about 10 photons (per ion) to be detected. Physically, this amounts to a requirement for optics able to collect at least 1% of the fluorescence emitted by each ion (with a background photon count rate of  $\leq \sim 10^3$ ), and to transmit this light to some form of photodetector, with sensitivity near or at the single-photon level.

### 8.2.4 Requirements for motional cooling

Provided that shuttling between trap zones can be executed with minimal motional heating, the SIMD QFT algorithm only requires that ions be cooled when they are within a trap zone. It is also necessary to be able to initialize trapped-ion qubits in

the motional ground state. This amounts to a requirement for Doppler cooling and sideband (or Raman) cooling capabilities in each trap zone. Physically, this may not correspond to individually addressed cooling beams if sympathetic cooling is used. However, the atomic transition chosen for Doppler cooling may also be used for state readout, in which case the Doppler cooling beam should have a projection onto each ion.

Depending on the method of cooling, the time required to prepare ions in the motional ground state can range from 100s of  $\mu\text{s}$  to a few ms, which could add significant overhead to our algorithm.

### 8.3 Capabilities of existing hardware

With a good understanding of the requirements, we will now investigate how currently available hardware could be used to implement the SIMD QFT algorithm in a planar trap array. We focus on CMOS-compatible technologies (assuming no particular manufacturing scale), although we will include some other integrated photonics references. Parallel to the overview provided in Section 8.2, this discussion is focused on: ion movement, qubit manipulation (gates), quantum state detection, and motional cooling. The following information is collected to show what is possible with surface-electrode traps and integrated optics.

#### 8.3.1 Ion movement

The ability to separate an ion from a chain (“splitting”), and to translate the ion to another location (“shuttling”) are essential elements of our QFT protocol. For the algorithm to work, these movements need to be efficient, with short duration (on the order of the gate time or less) and minimal energy gain (less than one motional quanta). In this section, we will describe what has been realized so far in terms of

ion shuttling and splitting.

Blakestad et al. [BOV<sup>+</sup>09, BOV<sup>+</sup>11] demonstrated reliable ion transport through a two-dimensional trap array incorporating an X-junction. An ion was transported 3.52 mm—traversing the X-junction twice—in 910  $\mu\text{s}$  (i.e. at a speed of about 3 m/s) with an energy increase of just  $0.18 \pm 0.02$  vibrational quanta per trip (for a motional frequency of 3.6 MHz). Following millions of successive traverses, they inferred a transport success probability of greater than 0.999995. Walther et al. [WZR<sup>+</sup>12] realized even faster ion transport in a segmented microchip trap: an ion was shuttled over a distance of 280  $\mu\text{m}$  in 3.6  $\mu\text{s}$  (corresponding to a speed of about 78 m/s) with energy gain as low as  $0.10 \pm 0.01$  quanta (for a motional frequency of 1.4 MHz). So, the requirement of shuttling times on the order of gate times ( $\sim 10 \mu\text{s}$ ) with minimal energy gain ( $\ll 1$  motional quanta) has already been realized. Further, theoretical investigations indicate that shuttle times on the order of a trap oscillation period ( $\sim 1 \mu\text{s}$ ) or less should be possible [LMC<sup>+</sup>14, FGP<sup>+</sup>14].


Splitting has been implemented as part of a quantum teleportation protocol [BCS<sup>+</sup>04] and the experimental entanglement purification of two-ion entanglement [RLK<sup>+</sup>06]. Ruster et al. [RWK<sup>+</sup>14] have experimentally demonstrated a protocol for fast ion separation which retains the ions in the Lamb-Dicke regime. Separation of a two-ion crystal in a segmented trap was achieved within 80  $\mu\text{s}$  (comparable to the time required to implement an MS gate) with a mean excitation of about 4 quanta per ion. This realization is consistent with the adiabatic limit corresponding to their trap parameters. Bowler et al. [BGL<sup>+</sup>12] have realized separation of two ions in just 55  $\mu\text{s}$ , with resulting excitations of about 2 quanta per ion. It is expected that technical improvements, together with newly developed techniques, will enable separation durations in the 10  $\mu\text{s}$  range, with energy transfers below the single-phonon level [KRS<sup>+</sup>14].

### 8.3.2 Qubit manipulation – waveguides and couplers

Integrated grating couplers could be used to emit beams tightly focused on individual ions, replacing conventional free-space laser beams grazing the trap. Space permitting, grating couplers could provide separate cooling and control beams for each individual ion. To bring light to the grating couplers, waveguides can be incorporated into the trap.

Orcutt et al. [OKH<sup>+</sup>11] demonstrated photonic integration within a state-of-the-art 28 nm CMOS foundry process (see Figure 8-2). They used the 80 nm thick polysilicon (typically deposited for transistor gates and local electrical interconnects) as the high-index waveguide core, yielding suitably strong confinement for transverse-electric-polarized light from 1.2  $\mu\text{m}$  to 1.6  $\mu\text{m}$  (this work focused on telecommunications wavelengths). The waveguides are polarization maintaining and typically less than 100 nm wide. Because these waveguides rely on deposited polysilicon that has not been optimized for photonics, the losses are higher, scaling with decreasing confinement factor down to approximately 30% [OMS<sup>+</sup>12]. Orcutt et al. [OMS<sup>+</sup>12] measured waveguide propagation losses below 15 dB/cm across the telecommunications spectrum, with a minimum of about 6 dB/cm at a wavelength of 1550 nm. This is higher than can be achieved outside of CMOS-integrated photonics; for example, propagation losses as low as 0.6 dB/cm have been achieved at telecom wavelengths using sputter-deposited aluminum nitride (AlN) film on (crystalline) silicon substrate [XPT12].

In the same CMOS process, Orcutt et al. [OKH<sup>+</sup>11] created broadband grating couplers for surface-normal optical input and output with bandwidth as great as 150 nm and insertion loss below 10 dB (tested over a range of 1280 nm to 1630 nm). The minimum insertion loss of 4.8 dB was obtained at a wavelength 1560 nm, with a 1 dB bandwidth of 93 nm. For comparison, the typical coupling efficiency is 5 dB at 1550 nm, and 10 dB at 770 nm for more specialized, silicon-integrated AlN grating



**Figure 8-2:** (a) SEM image showing a vertical grating coupler designed by Orcutt et al. [OKH<sup>+</sup>11]. (b) TEM image showing the cross-section of a CMOS-integrated waveguide. The waveguide has a  $670 \times 80$  nm polysilicon core, clad with a conformal 50 nm silicon nitride liner, surrounded by oxide. (Images reproduced from Ref. [OKH<sup>+</sup>11]).

couplers [XPT12]. The bandwidths possible in the CMOS couplers [OMS<sup>+</sup>12] are such that multiple beams could potentially be transmitted together, such as 729 nm and 866 nm beams, but they will exit the coupler at different angles, which may pose a problem. In terms of footprint, for a  $5 \mu\text{m}$  spot size (as may be desired for single-ion addressing)  $50 \mu\text{m}$  away from the surface, the grating coupler extent would be about  $15 \mu\text{m}$ . The maximum power that can be transmitted is likely to be limited by phase error rather than by thermal effects.

Integrated waveguides can be made large enough so that thermal effects due to high powers (in waveguides transmitting light to many ions, for example) will not be an issue. However, because of the Kerr effect (refractive index change with applied E-field) power changes in the waveguide will result in phase errors. Hence, optical pulses traveling through a waveguide will be distorted. The intensity-dependent phase shift (which is in addition to that from the linear refractive index) accumulated over some length  $L$  of waveguide transmission, is  $\phi = n_2 I k_0 L$ , where  $k_0 = 2\pi/\lambda$  is the vacuum wavenumber,  $I$  is the intensity in the waveguide, and  $n_2$  is the intensity-

dependent refractive index of the core material. The Kerr nonlinear properties of plasma-deposited silicon nitride were characterized for the first time in 2008. For IR wavelengths (1558 nm), the intensity-dependent refractive index  $n_2$  was measured to be  $2.4 \times 10^{-15} \text{ cm}^2/\text{W}$ , which is 10 times larger than that of silicon dioxide [ISAF08] (we have yet to find  $n_2$  measurements for visible wavelengths).

We can estimate the phase error expected due to the Kerr effect using laser parameters from the prototype system, summarized in Table 8.3. From this table, we see that the maximum laser intensity required at the ion position is  $\approx 3.5 \times 10^7 \text{ W/m}^2$  at a wavelength of 729 nm (the qubit laser). Assuming a  $5 \mu\text{m}$  spot size is achieved using a focusing grating coupler (as may be desired for single-ion addressing), this laser intensity corresponds to a laser power of about 1 mW being transmitted to the coupler. In the waveguide feeding the coupler, the intensity is given by  $I = P/A_{eff}$ , where  $A_{eff}$  is an effective area of the waveguide for this nonlinearity, which will be on the order of  $0.1 \mu\text{m}^2$  in our devices. So the laser intensity in this waveguide will be approximately  $10^{10} \text{ W/m}^2$ . After propagating through 1 cm of waveguide, the intensity-dependent phase shift incurred in this waveguide will be approximately  $\phi = 0.01$  degrees. So this phase shift will likely be negligible so long as modulators are placed fairly close to the ions they will address.

### 8.3.3 Qubit manipulation – modulators

To apply gate pulses, integrated phase/intensity modulators would be convenient (otherwise a zoo of stand-alone AOMs would be required on an adjacent laser table). In this section we discuss silicon-compatible, integrated modulators suitable for applying laser pulses to ions.


Because silicon has a narrow indirect bandgap (1.1 eV) and centrosymmetric crystal structure, it does not exhibit a second order ( $\chi^{(2)}$ ) nonlinearity (the Pockels effect). Commercial (non-integrated) high speed electro-optic modulators, such as devices



made from lithium niobate, rely on the the  $\chi^{(2)}$  field-dependent refractive index to achieve wide band modulation. In silicon, modulation can be achieved using the free carrier plasma dispersion effect, in which free carriers are injected or removed to change the free carrier density, which in turn alters the refractive index. Phase modulation in silicon using the free carrier plasma dispersion effect can be achieved in several different device configurations: forward biased p-i-n diode [TR95], MOS capacitor [LJL<sup>+</sup>04, LSRM<sup>+</sup>05], and reverse-biased pn junction [GREP05]. The forward biased p-i-n diode configuration offers superior modulation depth (defined as the resulting phase change when 1 V is applied to the modulator), while the MOS capacitor and reverse-biased pn junction configurations offer superior modulation speed. The modulation depth for a reverse-biased pn junction could potentially be improved through appropriate device design.

Intensity modulation can be realized using an interferometric Mach-Zehnder configuration which can be integrated into the trap. For example, Liu et al. [LLR<sup>+</sup>07] developed a high-speed silicon modulator based on a Mach-Zehnder interferometer (MZI) with a reverse-biased pn diode embedded in each of the two arms (see Figure 8-3). Each pn diode phase shifter is 3 mm long, a length chosen to keep the bias voltage low ( $< 5$  V). Liu et al. [LLR<sup>+</sup>07] demonstrate high-frequency modulator optical response with a 3 dB bandwidth of  $\approx 20$  GHz, for a continuous-wave laser beam at 1550 nm (this work focused on telecommunications wavelengths).

Unfortunately, the extinction of the integrated modulators is about 20 dB (at best), which means that the Rabi frequency when the modulator is quiescent will still be 10% of the maximum Rabi frequency. One solution might be to cascade several modulators to increase the extinction. For example, cascading 3 modulators would result in about 60 dB extinction, and a quiescent Rabi frequency of 0.1%. The extinction of the integrated Mach-Zehnder modulators could also likely be improved through device design changes, such as optimizing the dopant concentration and



**Figure 8-3:** Diagram showing the silicon modulator developed by Liu et al. [LLR<sup>+</sup>07]. It is a Mach-Zehnder interferometer containing two pn junction based phase shifters. The waveguide splitter is a  $1 \times 2$  multi-mode interference (MMI) coupler. (Image reproduced from Ref. [LLR<sup>+</sup>07]).


profile [LLR<sup>+</sup>07]. An alternative “software” solution could be to implement pulse sequences that correct for the quiescent drive.

### 8.3.4 Quantum state detection

For efficient state-detection, fluorescence capture from the ions should be maximized, and the subsequent photon-to-electron conversion optimized. The capability to discern the state of individual ions (by detecting their fluorescence individually) is also necessary. In this section we outline some different approaches to light collection suitable for a large-scale system.

Photodetectors can be integrated in CMOS; Field et al. [FLC<sup>+</sup>10] fabricated a single-photon avalanche diode (SPAD) within a standard  $0.13 \mu\text{m}$  CMOS process (see Figure 8-4). The fabricated device has a photosensitive area of about  $5 \mu\text{m} \times 5 \mu\text{m}$ , and the complete device occupies an area of about  $30 \mu\text{m} \times 30 \mu\text{m}$ . The photon detection probability peaks at just below 30% for a wavelength of 425 nm (measured at room temperature), while exhibiting a dark count rate of only 231 Hz and an impulse response of 198 ps [FLC<sup>+</sup>10]. The dead time is prohibitively long at  $15 \mu\text{s}$ , but could likely be improved a great deal by changing the quenching circuit used [FLC<sup>+</sup>10]. A previous SPAD fabricated using a high-voltage  $0.8 \mu\text{m}$  CMOS process has a dead

time of  $< 40$  ns [NRBC05].




**Figure 8-4:** (a) Micrograph of SPAD structure fabricated by Field et al. [FLC<sup>+</sup>10]. (b) Illustration of the expected pn diode structure after fabrication of the SPAD. (Images reproduced from Ref. [FLC<sup>+</sup>10]).

For integrated photodetectors like the CMOS SPAD, the collection area would likely need to be very small in order to fit within the trap electrodes (probably beneath an opening in the ground electrode) and to keep the detector capacitance low (to allow for sufficient bandwidth). For an ion-surface distance of  $50 \mu\text{m}$  and an inter-ion spacing of  $5 \mu\text{m}$ , the detector area would likely be limited to about  $200 \mu\text{m}^2$  per ion (to fit within the central ground electrode), which corresponds to fluorescence collection from just 0.6% of the solid angle. Solid-angle capture could be increased by reducing the ion-surface distance to  $20 \mu\text{m}$  (and scaling the opening to about  $80 \mu\text{m}^2$  to fit within a scaled-down ground electrode), resulting in fluorescence collection from about 1.6% of the solid angle, comparable to that achieved in conventional systems (see Section 4.1). To further increase solid-angle capture, the ground electrode could be made of a transparent conductor (such as indium tin oxide, as we demonstrated in Chapter 4), increasing the extent of the collection area (to span the outer-edges of the gaps separating the RF electrodes from the center ground electrode), resulting in

fluorescence collection from about 3% of the solid angle.


As an alternative to light collection through the electrodes, it might be more practical to use an array of Fresnel lenses above the trap (as in [SNJ<sup>+</sup>11], see Figure 8-5) or similar, and a CCD array. A downside of this approach is that light emitted by grating couplers on the trap surface will contribute some additional background signal, but this may not be significant because the beams from the couplers will be diffuse by this point.



**Figure 8-5:** Illustration showing the scheme proposed by Streed et al. [SNJ<sup>+</sup>11] for highly parallel readout of trapped-ion qubits. Fluorescence from ions trapped at multiple sites on a surface trap is efficiently coupled into single optical modes by an array of microfabricated phase Fresnel lenses on a single substrate. (Image reproduced from Ref. [SNJ<sup>+</sup>11]).

Another alternative for light detection would be to collect fluorescence from beneath each ion and couple it into a waveguide. The light could then be transmitted to an optimal detector mounted elsewhere. This brings to mind, in particular, the recent demonstration of superconducting-nanowire-single-photon detectors (SNSPDs) integrated with silicon photonic circuits (see Figure 8-6. In this demonstration, SNSPDs were transferred onto SOI waveguides using a micron-scale flip-chip process, resulting in on-chip detection efficiencies up to 52% for single photons with a wavelength of 1550 nm, sub-50-ps jitter, and nanosecond-scale reset time [NMH<sup>+</sup>15].





**Figure 8-6:** Optical micrograph of 10 waveguide-integrated detectors (D1 – D10) assembled on the same photonic chip by Najafi et al. [NMH<sup>+</sup>15]. The waveguides are marked by red arrows. The length of the scale bar (shown in blue) is 100  $\mu\text{m}$ . (Image reproduced from Ref. [NMH<sup>+</sup>15]).

### 8.3.5 Motional cooling

An important consideration is how much overhead cooling and re-cooling ions will add to the protocol. Depending on the method of cooling, the time required to prepare ions in the motional ground state can range from 100s of  $\mu\text{s}$  to a few ms (see Table 8.1). So, we would like to determine how often we can expect to re-cool ions during implementation of the algorithm. Let’s begin by assuming that we want to re-cool the ions whenever the mean number of motional quanta  $\bar{n}$  exceeds 2 quanta. The average heating rate measured in the CMOS trap was about 80 quanta/s (for a chip temperature of 8.4 K; see Section 6.4.4). Referring to the prototype system, it is reasonable to assume that most quantum gates can be completed in under 45  $\mu\text{s}$  (see Table 8.1). So, we will need to re-cool the ions approximately every 500 gates. From recent demonstrations of shuttling (see Section 8.3.1), we see that ideally each shuttling step will add energy corresponding to 0.1 quanta. Hence, we will need to re-cool the ions about every 20 shuttling steps.

Because the trap is typically much larger than the confined ion chain, we expect that electric field noise will be nearly uniform across the chain. Uniform electric field noise will only heat modes involving the center-of-mass (COM) motion, in which the ion chain moves as a rigid body. Non-COM modes, such as the collective breathing mode, can only be excited by field gradients, hence the heating of these modes is suppressed [KWM<sup>+</sup>98]. Thus, we should choose a non-COM mode for quantum logic.

However, heating of modes other than the logic mode will result in uncontrollable changes in the total wavepacket spread of the ion, causing the Rabi frequency of the transition between logic-mode motional states to also undergo uncontrollable changes. Hence, it is still necessary to keep the COM modes cool.


By combining an even number of qubit ions with a single cooling ion placed in the center of the chain, it is possible to use sympathetic cooling to reduce the thermal wavepacket spread of the ions, without disturbing the coherences of the internal qubits or the motional mode used for quantum logic. If the mass of the central, cooling ion is different from that of the qubit ions, all of the COM modes must be cooled (by addressing the central ion) to keep the ions in the Lamb-Dicke regime [KKM<sup>+</sup>00]. However, because we have access to individual addressing beams for each ion, we can choose for our central, cooling ion to be of the same species as our qubit ions, making their masses the same. In this limit, only the lowest axial mode will heat significantly, and we can maintain all ions in the Lamb-Dicke regime by cooling just this mode (by addressing only the central ion) [KKM<sup>+</sup>00]. Using the same ion species for sympathetic cooling also significantly simplifies ion loading into the trap array, and reduces the number of different lasers needed.

## 8.4 Visualizing an integrated trap module

With an understanding of the requirements, and the hardware available to construct an integrated trap module for implementing the QFT, we are now in a position to imagine what such a module might look like. This brief “design” section is intended more as an outlook than as a blueprint, but we hope that it will convey the potential of a modular trap architecture.

In Figure 8-7, we give a top-view of a hypothetical trap zone housing 5 ions, with fully-integrated optical components. Just beyond the segmented DC electrodes, openings allow laser beams emitted by focusing grating couplers to address ions from


one or more waveguides. Modulators apply pulse sequences and direct light. In the center of the trap zone, beneath the location of each ion, a grating coupler collects fluorescence into a waveguide. Each in-coupler is followed by a modulator, so that the fluorescence from the ions can be distinguished. Integrated electronics to control the optical devices and supply DC voltages to the segmented trap electrodes are not shown. For a more detailed look, the integrated electronics and optical devices to manipulate a single ion can be seen in Figure 8-8, which is a cross-sectional view.



**Figure 8-7:** Top view of a hypothetical trap zone for 5 ions. Ions would be trapped  $\sim 50 \mu\text{m}$  above the in-coupler sites, with an inter-ion spacing of  $\sim 5 \mu\text{m}$ . Integrated optical components are shown in place below trap electrodes (see text for details). Picture is not drawn to scale, and square bends (not practical for waveguides) are used only for clarity. In a real system, the width of the central ground electrode might be  $\sim 50 \mu\text{m}$  and the grating couplers and EO modulators can be expected to have a scale of  $\sim 10 \mu\text{m}$ .

	trap electrode
	waveguide
	focusing out-coupler
	in-coupler
	EO modulator





**Figure 8-8:** Conceptual drawing showing cross-sectional view of a hypothetical trap zone (not drawn to scale). A  $^{88}\text{Sr}^+$  ion is confined about  $50\ \mu\text{m}$  above the trap surface. Light emitted by two CMOS-compatible grating couplers comes to a focus at the ion position, delivering light for Doppler cooling and qubit control. Integrated waveguides carry light to these focusing grating couplers, and away from an in-coupler placed below the ion to collect fluorescence. Modulators are also integrated to direct light and apply pulses. Beneath a ground plane formed in a buried metal layer, additional metal layers are used to bias optical devices, and supply control voltages to the trap DC electrodes.

## 8.5 Conclusion

In this chapter, we took the SIMD QFT algorithm developed in Chapter 7 and asked how it might be realized in a modular trap architecture using CMOS-compatible integrated electronics and optical devices. Drawing on existing ion systems, and prior demonstrations, we assessed how close current hardware is to meeting the requirements posed by the SIMD QFT. We found that methods for ion transport—shuttling between trap zones and splitting within a zone—are already well-developed, and can meet the constraints for the algorithm. Appropriate waveguides and grating couplers have also been realized in a CMOS process, although not yet at visible wavelengths (but this is expected to be relatively straight-forward). CMOS-compatible EO modulators have been demonstrated, but further device development will be required to achieve the extinction ratio needed for pulse sequences. There are some promising options for efficient, scalable fluorescence collection from ions, but further study will also be necessary.

On the whole, the tools needed to implement the SIMD QFT seem to be within reach. Almost all of the requisite hardware has already been realized in some form. The prospects for a scalable trap module incorporating components for local optical and electronic control are good. With some further device development, a prototype module will be well worth constructing.

# Chapter 9

## Conclusion

Over the past decade, ion traps have morphed from macroscopic arrays of machined rods, blades, or needles—not much different from the original four-rod trap invented in the 1950s—to compact devices made up of planar electrodes on a single chip. Despite having lower trap depths and greater trap potential anharmonicity than their three-dimensional counterparts, microfabricated surface traps have been used to realize ground state cooling, high fidelity quantum gates, and entangling operations with ions confined 20 to 200 microns from the electrodes. A compelling proposal for a large-scale quantum computer is based on a so-called “quantum charge-coupled device” (QCCD), in which two-dimensional arrays of ions are positioned above a microfabricated surface. Gates are performed across different zones by shuttling ions between them. Movement of ions in a multiplexed trap array has been demonstrated experimentally, including splitting of ion crystals, shuttling of ions, transport of ions through junctions, and recombination of ion crystals.

With the QCCD blueprint, moving beyond proof-of-principle experiments to a machine encompassing thousands of manipulable ions still requires a great deal of technology development. Forming a scalable ion-light interface is crucial, as laser light is needed for cooling, state preparation, and gate operations. Additionally,

fluorescence emitted by the ions must be detected for quantum state measurement. To deliver and collect light efficiently from each individual trapped ion, it will likely be necessary to integrate optical components such as photodetectors and waveguides into the trap itself. In a large-scale system, classical electronics for automated ion control and movement will also need to be integrated into the trap. Thus the current challenge is to design an integrated ion system combining all the requisite components. The contributions of this thesis are towards developing the technology required to build trapped-ion systems up to a scale where useful computation is possible, through addressing a number of challenges, and developing modular building blocks.

A first challenge is gathering light efficiently from many trapped ions. We established a novel approach: capturing fluorescence transmitted through the transparent electrodes of a surface trap. If a transparent, multi-zone trap is combined with a photodetector array, massively parallel photon collection becomes possible. With the right photodetector mounted below the trap, this would also speed up quantum state detection significantly. Our demonstration of an ion trap with transparent electrodes opens the door to integration with other optical components. In the long run, however, light collection optics may instead become truly integrated with the trap, through fabrication of ion traps that incorporate optical waveguides or photodetectors in the substrate.

Next, we sought to better understand the sources of motional heating in surface-electrode traps, where the ion's proximity to the electrodes exacerbates the problem. Although we found that a graphene coating did not reduce motional heating as predicted, this investigation has shown how motional heating is strongly modulated by surface properties. We believe that this approach of experimenting with the material properties of trap electrodes holds promise for discovering scalable solutions to reduce heating. One potential idea might be to passivate the trap surface with a thin layer of a lipophobic material that prevents hydrocarbon molecules from sticking to the trap

---

during the vacuum bake. Another approach would be to clean or chemically break down (crack) hydrocarbon contamination in-situ.

Shifting to a more system-oriented view, another interesting challenge is to exploit existing CMOS fabrication processes—that have enabled scaling to billions of transistors—for quantum information processing. We have demonstrated basic functionality, including stable trapping and low electric-field noise, of a trap fabricated in an unmodified CMOS foundry process, opening the door to co-fabrication of advanced CMOS and photonics technology on an ion-trap chip, including the extensive existing libraries of integrated circuits for digital logic and memory. We believe that further development building on CMOS fabrication technologies offers a promising path towards scalable, local optical and electronic control and readout of trapped-ion arrays.

Finally, we turned our attention to mapping quantum algorithms onto ion trap hardware. We found that the quantum Fourier transform is a powerful algorithmic primitive with which to construct efficient quantum circuits for other quantum algorithms, such as Shor’s factoring algorithm. We also developed a framework for optimizing trap layout and shuttling patterns for the QFT based on an interpretation of Flynn’s taxonomy for a quantum computer architecture. Lastly, we discussed how CMOS and silicon-integrated photonics technology could be employed to build fully-integrated QFT modules. This work is an initial step toward quantum processing at scale with trapped ions. We think that the modular approach developed here will be valuable for the continued development of scalable trap architectures.

We hope that some of the techniques demonstrated in this thesis will prove useful to others working towards a large-scale quantum computer based on trapped ions. Perhaps the lasting effect of this work will be to show how versatile surface-electrode ion traps can be, and, ideally, to encourage continued explorations into exotic trap materials and integration with optical elements. It is an exciting moment for re-

searchers of quantum information processing, with many fundamental quantum operations already demonstrated with trapped ions, and a multitude of possibilities for further trap technology development. A majority of the current challenges to scaling up are purely technological, leading this author to believe that if sufficient resources and system-level efforts are devoted to further ion trap technology development, a quantum computer on the scale of 100s of qubits may not be far off.

# Appendix A

## Additional details about SuperDAC boards

The 3 SuperDAC PCBs are designed to snap together (using connectors on the front and back sides of each board) to make a 3-layer stack with the DAC board on the bottom, the ADC board in the middle, and the amplifier board on the top (since it may get warm). This design should save time in assembling the boards because no inter-board connection wiring needs to be prepared. (Each of the boards can also function completely stand-alone).

The dimensions of the 3 boards are:

- PCB size: 4 x 4 inches,
- Mounting hole size: 2.8 mm diameter.

Details of each board can be found in the following sections.

### A.1 DAC board

The DAC board features 4 buffered outputs, each with 20-bit resolution over a voltage range of -10 V to +10 V. For each output, there are 3 digital input pins on the board

which provide an SPI interface to set the output voltage. The board requires +15 V and -15 V power rails.

- The DAC board is based on the Analog Devices AD5791 DAC, which has a resolution of 20 bits over a -10 V to +10 V output range (in the configuration I chose);
- The op amp used for the output buffers is the Analog Devices AD8597 (low noise, slew rate: 14 V/us);
- The voltage reference for the DACs is the Analog Devices AD688 +/-10 V reference (which has a temperature coefficient of 3 ppm/degree C);
- For isolation on the digital inputs, the Analog Devices ADUM 1300 digital isolator is used;
- The op amp used for the power supply buffers is the Analog Devices AD8676 (dual op amp);
- Voltage regulator: NJM7805 +5 V regulator;
- Passives: All resistors are 0805 package; 10 uF capacitors are 0805 package; 1 uF and 0.1 uF capacitors are 0603 package.

The DAC board schematic and layout can be seen in Figures A-1 and A-2.



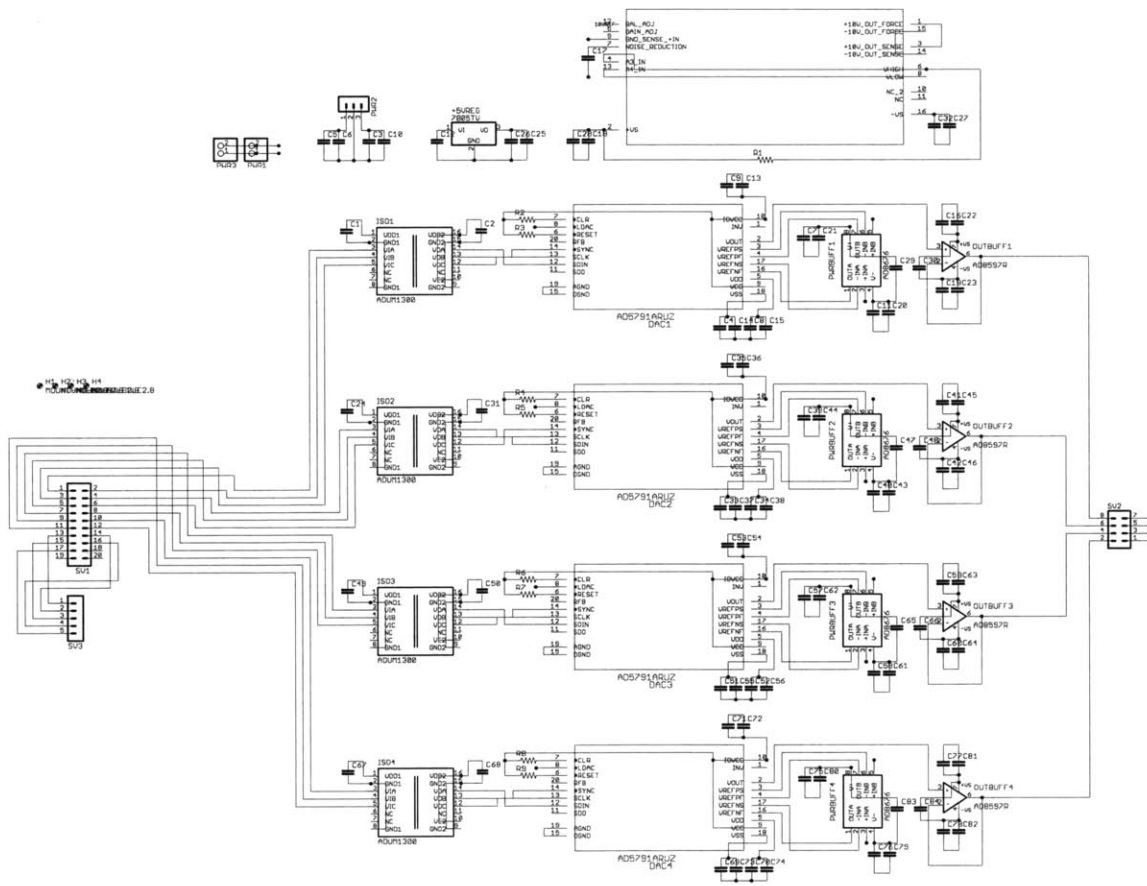


Figure A-1: Schematic for the DAC board (see text for details).




Figure A-2: Board layout for the DAC board (see text for details).

## A.2 High-voltage amplifier board

The high-voltage amplifier board provides a gain of 10 for 4 inputs up to  $\pm 75$  V rails. A bootstrapped op amp circuit is used to provide these large output swings. The board requires  $+80$  V and  $-80$  V power rails.

- The op amp used is the Analog Devices OP37, (low noise, slew rate of: 17 V/ $\mu$ s);
- The BJTs used are: Zetex ZTX657 (npn transistor) and ZTX757 (pnp transistor). They have a (collector-emitter) breakdown voltage of 300 V, and a DC current gain (hFE) of 50;
- All components on this board are through-hole.

The high-voltage amplifier board schematic and layout can be seen in Figures A-3 and A-4.

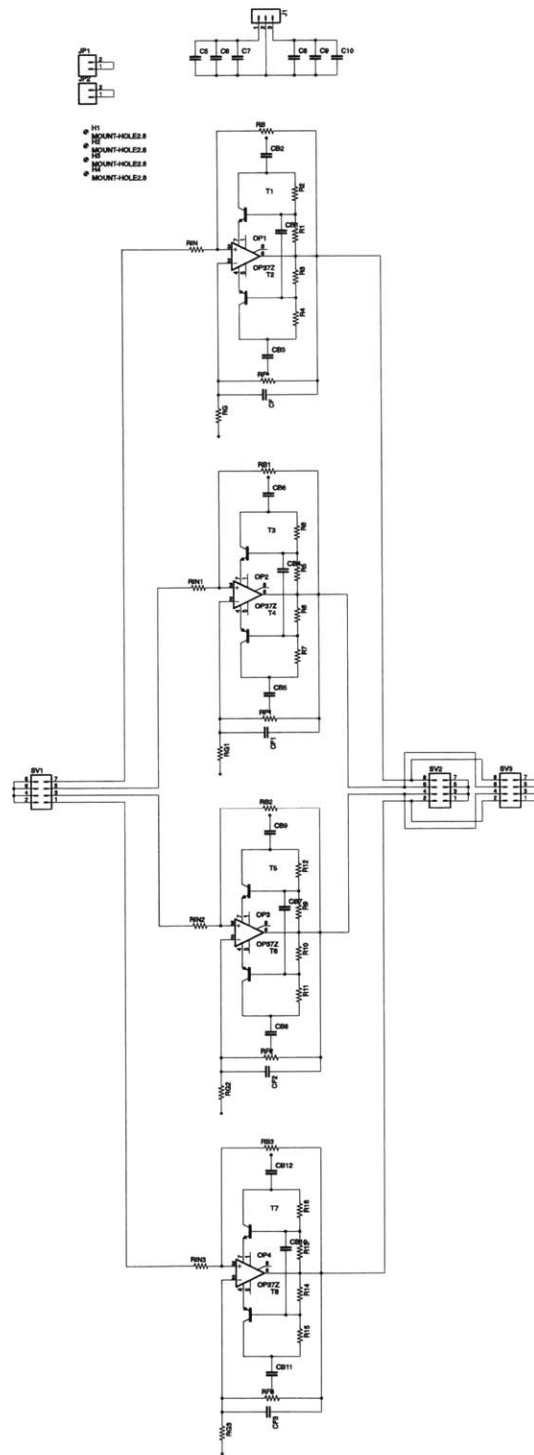


Figure A-3: Schematic for the high-voltage amplifier board (see text for details).




Figure A-4: Board layout for the high-voltage amplifier board (see text for details).

## A.3 ADC board

The ADC board has 4 inputs for measurement. A voltage ladder brings the inputs into the -2.5 V to +2.5 V measurement range. The board requires +15 V and -15 V power rails.

- The ADC board is based on the Texas Instruments ADS1248 ADC, which has a resolution of 24 bits over a -2.5 V to +2.5 V input range (in the configuration I chose);
- The voltage reference for the ADCs is the Analog Devices ADR4525 +2.5 V reference (which has a temperature coefficient of 2 ppm/degree C);
- For isolation on the digital inputs and outputs, the Analog Devices ADUM 1402 digital isolator is used;
- Voltage regulators:
  - NJM7805 +5 V regulator,
  - NJM7905 -5 V regulator,
  - Texas Instruments TPS79925 +2.5 V regulator,
  - Texas Instruments TPS72325 -2.5 V regulator;
- Passives: All resistors are 0805 package; 10 uF capacitors are 0805 package; 1 uF and 0.1 uF capacitors are 0603 package.

The ADC board schematic and layout can be seen in Figures A-5 and A-6.

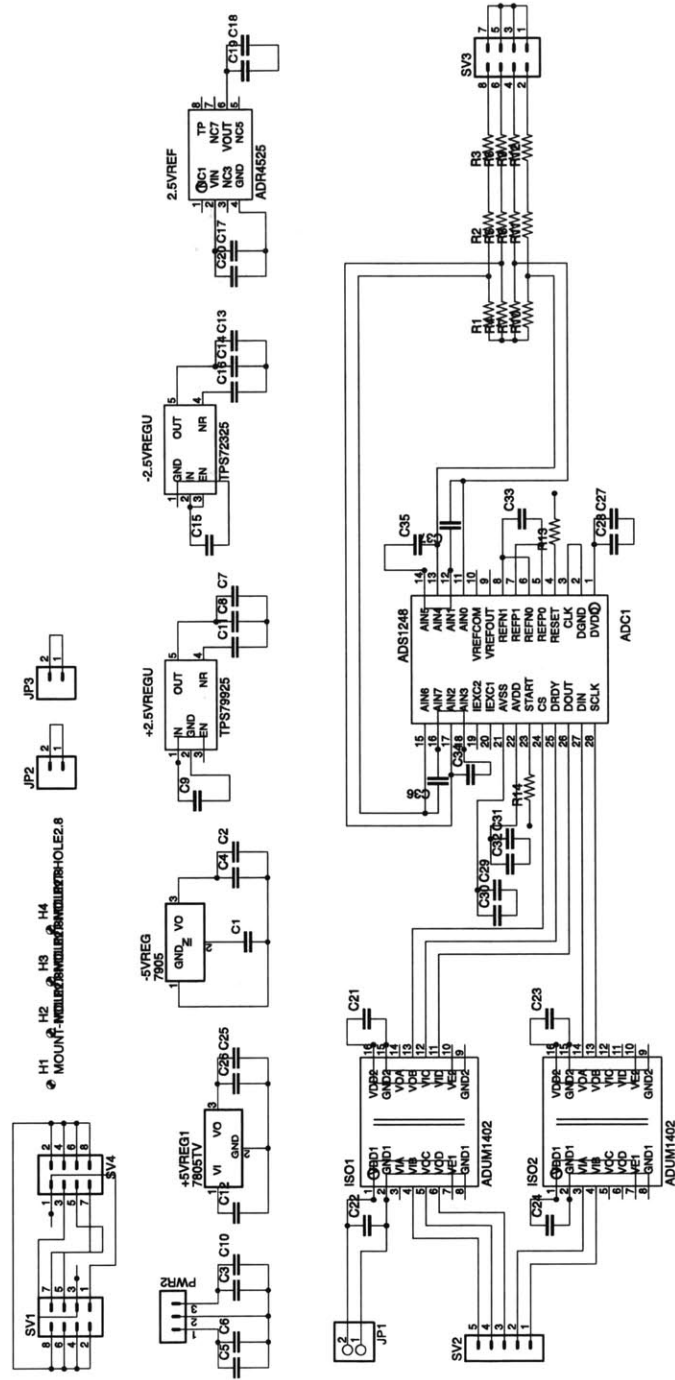


Figure A-5: Schematic for the ADC board (see text for details).






Figure A-6: Board layout for the ADC board (see text for details).



# Appendix B

## Additional details about QFT algorithm

### B.1 Pulse sequences for QFT with constant circuit depth

Within a single  $H$ -qubit linear chain, a full set of controlled rotation operations controlled by or targeting a single qubit can be performed with a constant-depth pulse sequence (with four MS gates regardless of the chain size). The following sequence performs the operation represented by the circuit in Figure B-1.

$$\left[ H^{(0)}, MS(\pi/4, 0), Z^{(0)}(\pi), MS(-\pi/4, 0); \left( \prod_j Z^{(j)}(\theta_j) \right), \right. \\ MS(\pi/4, 0), Z^{(0)}(\pi), MS(-\pi/4, 0), \left( \prod_j X^{(j)}(-\theta_j) \right), \quad (B.1) \\ \left. X^{(0)} \left( -\sum_j \theta_j \right), H^{(0)} \right]$$




Figure B-1: A set of simultaneous controlled rotations with  $H = 4$ .

## B.2 Details of simple QFT algorithm for $H_s = 1$ , $V = 2$

---

---

**Algorithm 1** Simple 2D QFT

---

---

**Input:** Qubits  $q_0 \dots q_{n-1}$  in chains  $C_0 \dots C_{W-1}$  of size  $H = n/W$ **Output:** QFT of qubits  $q_0 \dots q_{n-1}$ : controlled rotation  $R^{(c,t)}(\pi/2^{c-t})$  of each qubit  $\{q_t : t = 0 \dots (n-1)\}$  by each qubit  $\{q_c : c > t\}$ 

```

1: define  $P_0(m)$ : ▷ even permutation operator
2:    $P_0 \leftarrow \prod_{\alpha=0}^{\lfloor (m+1)/2 \rfloor} A_{2\alpha, 2\alpha+1}$ 
3: end define

4: define  $P_1(m)$ : ▷ odd permutation operator
5:    $P_1 \leftarrow \prod_{\alpha=0}^{\lfloor m/2 \rfloor} A_{2\alpha+1, 2\alpha+2}$ 
6: end define

7: define  $R_G^{(c)} C_\alpha$ : ▷ full-chain simultaneous controlled rotations operator
8:    $R_G^{(c)} C_\alpha \leftarrow \prod_{\{t: q_t \in C_\alpha, t < c\}} R^{(c,t)}(\pi/2^{c-t}) \cdot C_\alpha$ 
9: end define

10: function  $EvenGates(m)$  ▷ parallelized gates following  $P_0$ 
11:   parallel for  $\alpha \leftarrow 0$  to  $\lfloor (m-1)/2 \rfloor$  do
12:      $C_{2\alpha+1} \leftarrow R_G^{(c)} C_{2\alpha+1}$  ▷ simultaneous rotations in chain  $C_{2\alpha+1}$ 
13:   end parallel
14: end function

15: function  $OddGates(m)$  ▷ parallelized gates following  $P_1$ 
16:   parallel for  $\alpha \leftarrow 0$  to  $\lfloor m/2 \rfloor$  do
17:      $C_{2\alpha} \leftarrow R_G^{(c)} C_{2\alpha+1}$  ▷ simultaneous rotation op. in  $C_{2\alpha+1}$ ,
18:   end parallel ▷ controlled by shuttled qubit  $q_c$ 
19: end function

```

---

---



---

**Algorithm 1** Simple 2D QFT (Cont.)

---



---

▷▷ **Global variables:**  
 20:  $T \leftarrow 0$  ▷ global counter (or, step parity)  
 21:  $m \leftarrow 0$  ▷ next chain for swapping Fourier ion

▷▷ **Stage 1:**  
 22: **for**  $i \leftarrow 1$  to  $W - 1$  **do** ▷ loop over chains  
 23:     **for**  $j \leftarrow 0$  to  $2H - 2$  **do** ▷ shuttling iterations per chain

24:         **if**  $T \bmod 2$  is 0 **then** ▷ even permutation  
 25:              $D_0 \leftarrow P_0(m)D_0$   
 26:              $EvenGates(m)$   
 27:         **else** ▷ odd permutation  
 28:              $D_0 \leftarrow P_1(m)D_0$   
 29:              $OddGates(m)$   
 30:         **end if**

31:         **if**  $j \bmod 2$  is 0 and  $j < 2H - 2$  **then**  
 32:              $C_m \leftarrow SWAP(0, j/2)C_m$  ▷ swap Fourier ion to slot  $j/2$  in  $C_m$   
 33:         **end if**

34:          $T \leftarrow T + 1$  ▷ increment counter (or, switch parity)  
 35:     **end for**  
 36:      $m \leftarrow m + 1$  ▷ increment current chain  
 37: **end for**

---

**Algorithm 1** Simple 2D QFT (Cont.)

---



---

▷▷ **Stage 2:**  
 38: **for**  $i \leftarrow 0$  to  $W - 1$  **do** ▷ loop over chains  
 39:     **if**  $T \bmod 2$  is 0 **then** ▷ even permutation  
 40:          $D_0 \leftarrow P_0(W)D_0$   
 41:          $EvenGates(W)$   
 42:     **else** ▷ odd permutation  
 43:          $D_0 \leftarrow P_1(W)D_0$   
 44:          $OddGates(W)$   
 45:     **end if**  
  
 46:      $T \leftarrow T + 1$  ▷ increment counter (or, switch parity)  
 47: **end for**  
  
 ▷▷ **Stage 3:**  
 48: **for**  $i \leftarrow 0$  to  $\lfloor W/(H - 1) \rfloor - 1$  **do** ▷ loop over chains  
 49:     **for**  $j \leftarrow 0$  to  $2H - 2$  **do** ▷ shuttling iterations per chain  
  
 50:          $r \leftarrow \max(3W - m, \lceil W/(H - 1) \rceil)$   
 51:          $r \leftarrow \min(r, W)$  ▷ permutations and gates affect first  $r$  chains  
  
 52:         **if**  $T \bmod 2$  is 0 **then** ▷ even permutation  
 53:              $D_0 \leftarrow P_0(r)D_0$   
 54:              $EvenGates(r)$   
 55:         **else** ▷ odd permutation  
 56:              $D_0 \leftarrow P_1(r)D_0$   
 57:              $OddGates(r)$   
 58:         **end if**  
  
 59:         **if**  $j \bmod 2$  is 0 and  $j < 2H - 2$  **then** ▷ swap in chain  $C_{m \bmod W}$   
 60:              $C_{m \bmod W} \leftarrow SWAP(0, j/2) C_{m \bmod W}$   
 61:         **end if**  
  
 62:          $T \leftarrow T + 1$  ▷ increment counter (or, switch parity)  
 63:     **end for**  
 64:      $m \leftarrow m + 1$  ▷ increment current chain  
 65: **end for**

---



# Bibliography

- [AGH<sup>+</sup>11] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Reduction of heating rate in a micro-fabricated ion trap by pulsed-laser cleaning. *New Journal of Physics*, 13(12):123023, December 2011. 104
- [AHJ<sup>+</sup>12] D. T. C. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki, S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring. Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap. *Applied Physics B: Lasers and Optics*, 107(4):913–919, November 2012. 110, 118
- [ASA<sup>+</sup>09] P. B. Antohi, D. Schuster, G. M. Akselrod, J. Labaziewicz, Y. Ge, Z. Lin, W. S. Bakr, and I. L. Chuang. Cryogenic ion trapping systems with surface-electrode traps. *Review of Scientific Instruments*, 80(1):013103, January 2009. 59, 65
- [Bak06] W. S. Bakr. *Towards a Cryogenic Planar Ion Trap for Sr-88*. B.S. Thesis, Department of Physics, Massachusetts Institute of Technology, 2006. 59, 65
- [BCL<sup>+</sup>07] K. R. Brown, R. J. Clark, J. Labaziewicz, P. Richerme, D. R. Leibbrandt, and I. L. Chuang. Loading and characterization of a printed-circuit-

- board atomic ion trap. *Physical Review A - Atomic, Molecular, and Optical Physics*, 75(1):015401, January 2007. 34
- [BCS<sup>+</sup>04] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Deterministic quantum teleportation of atomic qubits. *Nature*, 803(1982):802–803, 2004. 32, 153
- [BEM<sup>+</sup>11] G. R. Brady, A. R. Ellis, D. L. Moehring, D. Stick, C. Highstrete, K. M. Fortier, M. G. Blain, R. A. Haltli, A. A. Cruz-Cabrera, R. D. Briggs, J. R. Wendt, T. R. Carter, S. Samora, and S. A. Kemme. Integration of fluorescence collection optics with a microfabricated surface electrode ion trap. *Applied Physics B: Lasers and Optics*, 103(4):801–808, June 2011. 45, 82
- [BEST96] A. Barenco, A. Ekert, K.-A. Suominen, and P. Törmä. Approximate Quantum Fourier Transform and Decoherence. *Physical Review A*, 54(1):14, July 1996. 128
- [BGB<sup>+</sup>09] U. Bangert, M. H. Gass, A. L. Bleloch, R. R. Nair, and J. Eccles. Nanotopography of graphene. *Physica Status Solidi (A) Applications and Materials Science*, 206(9):2115–2119, September 2009. 105
- [BGL<sup>+</sup>12] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. Coherent diabatic ion transport and separation in a multizone trap array. *Physical Review Letters*, 109(8):1–4, 2012. 153
- [BHL<sup>+</sup>05] K. A. Brickman, P. C. Haljan, P. J. Lee, M. Acton, L. Deslauriers, and C. Monroe. Implementation of Grover’s quantum search algorithm in



- a scalable system. *Physical Review A - Atomic, Molecular, and Optical Physics*, 72(5):050306, November 2005. 33
- [Bla12] A. M. Blake. *Computing the fast Fourier transform on SIMD microprocessors*. PhD Thesis, Department of Computer Science, University of Waikato, June 2012. 124
- [BMB<sup>+</sup>98] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10):5025–5033, 1998. 77
- [BMS<sup>+</sup>11] J. T. Barreiro, M. Müller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt. An open-system quantum simulator with trapped ions. *Nature*, 470(7335):486–491, February 2011. 32
- [BOV<sup>+</sup>09] 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):1–4, 2009. 153
- [BOV<sup>+</sup>11] 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. *Physical Review A - Atomic, Molecular, and Optical Physics*, 84(3):1–14, 2011. 153
- [BR12] R. Blatt and C. F. Roos. Quantum simulations with trapped ions. *Nature Physics*, 8(4):277–284, April 2012. 32, 146

- [Bri08] J. W. Britton. *Microfabrication techniques for trapped ion quantum information processing*. PhD Thesis, Department of Physics, University of Colorado at Boulder, January 2008. 110, 118
- [Bru96] Bruce Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, 2nd edition, 1996. 29
- [BSK<sup>+</sup>12] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger. Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins. *Nature*, 484(7395):489–492, April 2012. 32
- [BVA<sup>+</sup>08] J. S. Bunch, S. S. Verbridge, J. S. Alden, A. M. Van Der Zande, J. M. Parpia, H. G. Craighead, and P. L. McEuen. Impermeable atomic membranes from graphene sheets. *Nano Letters*, 8(8):2458–2462, August 2008. 96
- [BWC13] A. M. Blake, I. H. Witten, and M. J. Cree. The Fastest Fourier Transform in the South. *IEEE Transactions on Signal Processing*, 61(19):4707–4716, October 2013. 124
- [CBB<sup>+</sup>05] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Surface-Electrode Architecture for Ion-Trap Quantum Information Processing. *Quantum Information & Computation*, 5(6):21, September 2005. 34, 45
- [CBL<sup>+</sup>05] J. Chiaverini, J. Britton, D. Leibfried, E. Knill, M. D. Barrett, R. B. Blakestad, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, T. Schaetz, and D. J. Wineland. Implementation of the semiclassical quantum Fourier transform in a scalable system. *Science*, 308(5724):997–1000, May 2005.

- [CBL<sup>+</sup>11] S. Chen, L. Brown, M. Levendorf, W. Cai, S. Y. Ju, J. Edgeworth, X. Li, C. W. Magnuson, A. Velamakanni, R. D. Piner, J. Kang, J. Park, and R. S. Ruoff. Oxidation resistance of graphene-coated Cu and Cu/Ni alloy. *ACS Nano*, 5(2):1321–1327, February 2011. 95, 96, 105
- [CDBM08] J. Coraux, A. T. Diaye, C. Busse, and T. Michely. Structural coherency of graphene on Ir(111). *Nano Letters*, 8(2):565–570, February 2008. 97
- [CGD<sup>+</sup>10] C. R. Clark, J. E. Goeders, Y. K. Dodia, C. R. Viteri, and K. R. Brown. Detection of single-ion spectra by Coulomb-crystal heating. *Physical Review A - Atomic, Molecular, and Optical Physics*, 81(4):043428, April 2010. 105
- [CS14] J. Chiaverini and J. M. Sage. Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range. *Physical Review A - Atomic, Molecular, and Optical Physics*, 89(1):012318, January 2014. 95, 117
- [CZKM96] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Physical Review Letters*, 78(16):4, April 1996. 34, 132
- [CZL<sup>+</sup>09] W. Cai, Y. Zhu, X. Li, R. D. Piner, and R. S. Ruoff. Large area few-layer graphene/graphite films as transparent thin conducting electrodes. *Applied Physics Letters*, 95(12):123115, 2009. 97
- [DBIW89] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Laser Cooling to the Zero-Point Energy of Motion. *Physical Review Letters*, 62(4):403–406, January 1989. 55, 58, 101

- [DBMM04] L. M. Duan, B. B. Blinov, D. L. Moehring, and C. Monroe. Scalable Trapped Ion Quantum Computation with a Probabilistic Ion-Photon Mapping. *Quantum Information & Computation*, 4(3):6, 2004. 35
- [DCG09] R. Dubessy, T. Coudreau, and L. Guidoni. Electric field noise above surfaces: A model for heating-rate scaling law in ion traps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 80(3):031402(R), September 2009. 95
- [DCT99] J. Dalibard and C. Cohen-Tannoudji. Atomic motion in laser light: connection between semiclassical and quantum descriptions. *Journal of Physics B: Atomic and Molecular Physics*, 18(8):1661–1683, April 1999. 55
- [Deu85] D. Deutsch. Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 400(1818):97–117, July 1985. 29
- [DFL08] Y. S. Dedkov, M. Fonin, and C. Laubschat. A possible source of spin-polarized electrons: The inert graphene/Ni(111) system. *Applied Physics Letters*, 92(5):052506, 2008. 96
- [DFRL08] Y. S. Dedkov, M. Fonin, U. Rüdiger, and C. Laubschat. Graphene-protected iron layer on Ni(111). *Applied Physics Letters*, 93(2):022509, 2008. 96
- [DGB<sup>+</sup>14] N. Daniilidis, S. Gerber, G. Bolloten, M. Ramm, A. Ransford, E. Ulin-Avila, I. Talukdar, and H. Häffner. Surface noise analysis using a single-ion sensor. *Physical Review B - Condensed Matter and Materials Physics*, 89(24):245435, June 2014. 104

- [dHBW<sup>+</sup>07] W. A. de Heer, C. Berger, X. Wu, P. N. First, E. H. Conrad, X. Li, T. Li, M. Sprinkle, J. Hass, M. L. Sadowski, M. Potemski, and G. Martinez. Epitaxial graphene. *Solid State Communications*, 143(1-2):92–100, July 2007. 96
- [DHL<sup>+</sup>04] L. Deslauriers, P. C. Haljan, P. J. Lee, K. A. Brickman, B. B. Blinov, M. J. Madsen, and C. Monroe. Zero-point cooling and low heating of trapped  $^{111}\text{Cd}^+$  ions. *Physical Review A - Atomic, Molecular, and Optical Physics*, 70(4):043408, October 2004. 94
- [DiV00] D. P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschritte der Physik*, 48(9-11):771–783, September 2000. 29
- [DJ92] D. Deutsch and R. Jozsa. Rapid Solution of Problems by Quantum Computation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 439(1907):553–558, December 1992. 29
- [DK02] R. DeVoe and C. Kurtsiefer. Experimental study of anomalous heating and trap instabilities in a microscopic  $^{137}\text{Ba}$  ion trap. *Physical Review A*, 65(6):063407, June 2002. 94
- [DKSS13] A. De, P. P. Kurur, C. Saha, and R. Saptharishi. Fast Integer Multiplication Using Modular Arithmetic. *SIAM Journal on Computing*, 42(2):685–699, April 2013. 124
- [DNM<sup>+</sup>11] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, S. Schulz, F. Schmidt-Kaler, and H. Häffner. Fabrication and heating rate study of microscopic surface electrode ion traps. *New Journal of Physics*, 13(7):079504, July 2011. 94, 95

- [DOS<sup>+</sup>06] 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):103007, September 2006. 94, 101
- [DS13] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: an outlook. *Science*, 339(6124):1169–74, March 2013. 31
- [FÖ9] M. Fürer. Faster Integer Multiplication. *SIAM Journal on Computing*, 39(3):979–1005, January 2009. 124
- [Fey82] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982. 29
- [FGP<sup>+</sup>14] H. A. Fürst, M. H. Goerz, U. G. Poschinger, M. Murphy, S. Montangero, T. Calarco, F. Schmidt-Kaler, K. Singer, and C. P. Koch. Controlling the transport of an ion: classical and quantum mechanical solutions. *New Journal of Physics*, 16(075007), 2014. 153
- [FKO05] T. Fujita, W. Kobayashi, and C. Oshima. Novel structures of carbon layers on a Pt(111) surface. *Surface and Interface Analysis*, 37(2):120–123, February 2005. 97
- [FLC<sup>+</sup>10] R. M. Field, J. Lary, J. Cohn, L. Paninski, and K. L. Shepard. A low-noise, single-photon avalanche diode in standard 0.13  $\mu\text{m}$  complementary metal-oxide-semiconductor process. *Applied Physics Letters*, 97(21):211111, 2010. 25, 120, 158, 159
- [Fly72] M. J. Flynn. Some Computer Organizations and Their Effectiveness. *IEEE Transactions on Computers*, C-21(9):948–960, 1972. 122, 139

- [GFH<sup>+</sup>14] N. D. Guise, S. D. Fallek, H. Hayden, C.-S. Pai, C. Volin, K. R. Brown, J. T. Merrill, A. W. Harter, J. M. Amini, L. M. Lust, K. Muldoon, D. Carlson, and J. Budach. In-Vacuum Active Electronics for Microfabricated Ion Traps. *Review of Scientific Instruments*, 85:063101, 2014. 119
- [Gho95] P. K. Ghosh. *Ion Traps*. Oxford University Press, USA, 1995. 44
- [GLK<sup>+</sup>11] R. Gerritsma, B. P. Lanyon, G. Kirchmair, F. Zähringer, C. Hempel, J. Casanova, J. J. García-Ripoll, E. Solano, R. Blatt, and C. F. Roos. Quantum simulation of the Klein paradox with trapped ions. *Physical Review Letters*, 106(6):060503, February 2011. 32
- [GN07] A. K. Geim and K. S. Novoselov. The rise of graphene. *Nature materials*, 6(3):183–191, March 2007. 96
- [GREP05] F. Gardes, G. Reed, N. Emerson, and C. Png. A sub-micron depletion-type photonic modulator in Silicon On Insulator. *Optics express*, 13(22):8845–8854, 2005. 157
- [GRL<sup>+</sup>03] S. Gulde, M. Riebe, G. P. T. Lancaster, C. Becher, J. Eschner, H. Häffner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt. Implementation of the Deutsch-Jozsa algorithm on an ion-trap quantum computer. *Nature*, 421(6918):48–50, January 2003. 33
- [Gro96] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings, 28th Annual ACM Symposium on the Theory of Computing (STOC)*, page 8, 1996. 28, 29
- [Hĭ4] H. Häffner. Reducing Motional Decoherence in Ion Traps with Surface Science Methods. *March Meeting of the American Physical Society, Denver, CO, MAR.M34.1*, March 2014. 104

- [HBHB10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New Journal of Physics*, 12(9):93035, 2010. 84
- [HCW<sup>+</sup>12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-Fold Reduction of Electric-Field Noise in an Ion Trap Cleaned With In Situ Argon-Ion-Beam Bombardment. *Physical Review Letters*, 109(10):103001, September 2012. 104
- [HCW<sup>+</sup>13] D. A. Hite, Y. Colombe, A. Wilson, D. T. C. Allcock, D. Leibfried, D. Wineland, and D. Pappas. Surface science for improved ion traps. *MRS Bulletin*, 38(10):826–833, October 2013. 104
- [HHJ<sup>+</sup>09] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. Complete methods set for scalable ion trap quantum information processing. *Science*, 325(5945):1227–1230, September 2009. 32
- [HHL09] A. W. Harrow, A. Hassidim, and S. Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502, October 2009. 29
- [Hou08] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 78(3), September 2008. 46, 60, 112
- [HS75] T. Hänsch and A. L. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13(1):68–69, January 1975. 52
- [Hug01] R. Hughes. Quantum Computation. *Computing in Science & Engineering*, 3(2):255–261, 2001. 32



- [HWS<sup>+</sup>11] P. F. Herskind, S. X. Wang, M. Shi, Y. Ge, M. Cetina, and I. L. Chuang. Microfabricated surface ion trap on a high-finesse optical mirror. *Optics letters*, 36(16):3045–3047, August 2011. 82, 83
- [ISAF08] K. Ikeda, R. E. Saperstein, N. Alic, and Y. Fainman. Thermal and Kerr nonlinear properties of plasma-deposited silicon nitride/silicon dioxide waveguides. *Optics express*, 16(17):12987–12994, 2008. 156
- [ISC<sup>+</sup>13] R. Islam, C. Senko, W. C. Campbell, S. Korenblit, J. Smith, A. Lee, E. E. Edwards, C.-C. J. Wang, J. K. Freericks, and C. Monroe. Emergence and frustration of magnetism with variable-range interactions in a quantum simulator. *Science*, 340(6132):583–7, 2013. 32
- [JMS86] L. H. Jamieson, P. T. Mueller, and H. J. Siegel. FFT algorithms for SIMD parallel processing systems. *Journal of Parallel and Distributed Computing*, 3(1):48–71, March 1986. 124
- [JSN<sup>+</sup>11] A. Jechow, E. W. Streed, B. G. Norton, M. J. Petراسiunas, and D. Kielpinski. Wavelength-scale imaging of trapped ions using a phase Fresnel lens. *Optics letters*, 36(8):1371–1373, April 2011. 82
- [KCK<sup>+</sup>10] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe. Quantum simulation of frustrated Ising spins with trapped ions. *Nature*, 465(7298):590–593, June 2010. 32
- [KHC11] T. H. Kim, P. F. Herskind, and I. L. Chuang. Surface-electrode ion trap with integrated light source. *Applied Physics Letters*, 98(21):214103, 2011. 45, 82
- [Kim08] H. J. Kimble. The quantum internet. *Nature*, 453(7198):1023–1030, June 2008. 15, 35, 36

- [KJL<sup>+</sup>08] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences of the United States of America*, 105(48):18681–18686, August 2008. 129
- [KKM<sup>+</sup>00] D. Kielpinski, B. King, C. Myatt, C. Sackett, Q. Turchette, W. Itano, C. Monroe, D. Wineland, and W. Zurek. Sympathetic cooling of trapped ions for quantum logic. *Physical Review A*, 61(3):032310, February 2000. 162
- [KLH<sup>+</sup>04] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431(7012):1075–1078, 2004. 82
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, June 2002. 15, 34, 35, 45, 132
- [KPH<sup>+</sup>99] H. Kim, A. Pique, J. S. Horwitz, H. Mattoussi, H. Murata, Z. H. Kafafi, and D. B. Chrisey. Indium tin oxide thin films for organic light-emitting devices. *Applied Physics Letters*, 74(23):3444, 1999. 84, 88
- [KPM<sup>+</sup>05] J. Kim, S. Pau, Z. Ma, H. McLellan, J. Gates, A. Kornblit, R. Slusher, R. Jopson, I. Kang, and M. Dinu. System design for large-scale ion trap quantum information processor. *Quantum Information & Computation*, 5(7):515–537, November 2005. 45
- [KRS<sup>+</sup>14] H. Kaufmann, T. Ruster, C. T. Schmiegelow, F. Schmidt-Kaler, and U. G. Poschinger. Dynamics and control of fast ion crystal splitting in segmented Paul traps. *New Journal of Physics*, 16(073012), 2014. 153

- [KSDL10] W. J. Kim, A. O. Sushkov, D. A. R. Dalvit, and S. K. Lamoreaux. Surface contact potential patches and Casimir force measurements. *Physical Review A - Atomic, Molecular, and Optical Physics*, 81(2):022505, February 2010. 94
- [KWM<sup>+</sup>98] B. E. King, C. S. Wood, C. J. Myatt, Q. A. Turchette, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. Cooling the Collective Motion of Trapped Ions to Initialize a Quantum Register. *Physical Review Letters*, 81(7):5, August 1998. 32, 161
- [KZJ<sup>+</sup>09] K. S. Kim, Y. Zhao, H. Jang, S. Y. Lee, J. M. Kim, K. S. Kim, J.-H. Ahn, P. Kim, J.-Y. Choi, and B. H. Hong. Large-scale pattern growth of graphene films for stretchable transparent electrodes. *Nature*, 457(7230):706–710, February 2009. 98
- [Lab08] J. Labaziewicz. *High Fidelity Quantum Gates with Ions in Cryogenic Microfabricated Ion Traps*. PhD Thesis, Department of Physics, Massachusetts Institute of Technology, 2008. 61
- [LCA<sup>+</sup>09] X. Li, W. Cai, J. An, S. Kim, J. Nah, D. Yang, R. Piner, A. Velamakanni, I. Jung, E. Tutuc, S. K. Banerjee, L. Colombo, and R. S. Ruoff. Large-area synthesis of high-quality and uniform graphene films on copper foils. *Science*, 324(5932):1312–1314, June 2009. 97, 99
- [LDM<sup>+</sup>03] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(6930):412–415, March 2003. 33

- [LGA<sup>+</sup>08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. *Physical Review Letters*, 100(1):013001, January 2008. 39, 59, 65, 69, 74, 86, 94, 97, 101, 102, 103, 107
- [LGL<sup>+</sup>08] J. Labaziewicz, Y. Ge, D. R. Leibbrandt, S. X. Wang, R. Shewmon, and I. L. Chuang. Temperature dependence of electric field noise above gold surfaces. *Physical Review Letters*, 101(18):180602, October 2008. 94, 101, 103, 107
- [LHC11] G. H. Low, P. F. Herskind, and I. L. Chuang. Finite-geometry models of electric field noise from patch potentials in ion traps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 84(5):053425, November 2011. 95
- [LHM<sup>+</sup>09] L. Luo, D. Hayes, T. A. Manning, D. N. Matsukevich, P. Maunz, S. Olmschenk, J. D. Sterk, and C. Monroe. Protocols and techniques for a scalable atom-photon quantum network. *Fortschritte der Physik*, 57(11-12):1133–1152, September 2009. 36, 81, 91
- [LHN<sup>+</sup>11] B. P. Lanyon, C. Hempel, D. Nigg, M. Müller, R. Gerritsma, F. Zähringer, P. Schindler, J. T. Barreiro, M. Rambach, G. Kirchmair, M. Hennrich, P. Zoller, R. Blatt, and C. F. Roos. Universal digital quantum simulation with trapped ions. *Science*, 334:13, 2011. 32
- [LJL<sup>+</sup>04] A. Liu, R. Jones, L. Liao, D. Samara-rubio, D. Rubin, O. Cohen, R. Nicolaescu, and M. Paniccia. A high-speed silicon optical modulator based on a metaloxide semiconductor capacitor. *Nature*, 427(February):615–618, 2004. 157

- [LLC<sup>+</sup>09] D. R. Leibbrandt, J. Labaziewicz, R. J. Clark, I. L. Chuang, R. J. Epstein, C. Ospelkaus, J. H. Wesenberg, J. J. Bollinger, D. Leibfried, D. J. Wineland, D. Stick, J. Sterk, C. Monroe, C. S. Pai, Y. Low, R. Frahm, and R. E. Slusher. Demonstration of a scalable, multiplexed ion trap for quantum information processing. *Quantum Information & Computation*, 9(11):18, 2009. 110, 118
- [LLR<sup>+</sup>07] A. Liu, L. Liao, D. Rubin, H. Nguyen, B. Ciftcioglu, Y. Chetrit, N. Izhaky, and M. Paniccia. High-speed optical modulation based on carrier depletion in a silicon waveguide. *Optics express*, 15(2):660–8, 2007. 25, 157, 158
- [LMC<sup>+</sup>14] X. J. Lu, J. G. Muga, X. Chen, U. G. Poschinger, F. Schmidt-Kaler, and A. Ruschhaupt. Fast shuttling of a trapped ion in the presence of noise. *Physical Review A - Atomic, Molecular, and Optical Physics*, 89(6):1–10, 2014. 153
- [LMM<sup>+</sup>97] D. Leibfried, D. Meekhof, C. Monroe, B. King, W. Itano, and D. Wineland. Engineering And Tomography Of Quantum States Of Motion Of A Trapped Atom. *QELS '97., Summaries of Papers Presented at the Quantum Electronics and Laser Science Conference*, 44(11):2485–2505, 1997. 57
- [LRVGP09] M. P. Levendorf, C. S. Ruiz-Vargas, S. Garg, and J. Park. Transfer-free batch fabrication of single layer graphene transistors. *Nano Letters*, 9(12):4479–4483, December 2009. 97, 98
- [LSRM<sup>+</sup>05] L. Liao, D. Samara-Rubio, M. Morse, A. Liu, D. Hodge, D. Rubin, U. Keil, and T. Franck. High speed silicon Mach-Zehnder modulator. *Optics express*, 13(8):3129–3135, 2005. 157

- [LTCV04] Y.-J. Lin, I. Teper, C. Chin, and V. Vuletić. Impact of the Casimir-Polder potential and Johnson noise on Bose-Einstein condensate stability near surfaces. *Physical Review Letters*, 92(5):050404, February 2004. 94
- [LZC<sup>+</sup>09] X. Li, Y. Zhu, W. Cai, M. Borysiak, B. Han, D. Chen, R. D. Piner, L. Colomba, and R. S. Ruoff. Transfer of large-area graphene films for high-performance transparent conductive electrodes. *Nano Letters*, 9(12):4359–4363, December 2009. 97
- [MGF<sup>+</sup>12] R. Maiwald, A. Golla, M. Fischer, M. Bader, S. Heugel, B. Chalopin, M. Sondermann, and G. Leuchs. Collecting more than half the fluorescence photons from a single ion. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(4):043431, October 2012. 82
- [MKB<sup>+</sup>02] A. B. Mundt, A. Kreuter, C. Becher, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Coupling a single atomic quantum bit to a high finesse optical cavity. *Physical Review Letters*, 89(10):103001, August 2002. 82
- [MKE<sup>+</sup>08] J. C. Meyer, C. Kisielowski, R. Erni, M. D. Rossell, M. F. Crommie, and A. Zettl. Direct imaging of lattice atoms and topological defects in graphene membranes. *Nano Letters*, 8(11):3582–3586, November 2008. 104
- [MKH09a] K. S. McKay, J. Kim, and H. H. Hogue. Enhanced quantum efficiency of the visible light photon counter in the ultraviolet wavelengths. *Optics express*, 17(9):7458–7464, April 2009. 89
- [MKH<sup>+</sup>09b] T. Monz, K. Kim, W. Hänsel, M. Riebe, A. S. Villar, P. Schindler, M. Chwalla, M. Hennrich, and R. Blatt. Realization of the quantum

- Toffoli gate with trapped ions. *Physical Review Letters*, 102(4):040501, January 2009. 33
- [MMK<sup>+</sup>95a] C. Monroe, D. Meekhof, B. King, W. Itano, and D. Wineland. Demonstration of a Fundamental Quantum Logic Gate. *Physical Review Letters*, 75(25):4714–4717, 1995. 33
- [MMK<sup>+</sup>95b] C. Monroe, D. Meekhof, B. King, S. Jefferts, W. Itano, D. Wineland, and P. Gould. Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy. *Physical Review Letters*, 75(22):4011–4014, 1995. 32
- [MN98] C. Moore and M. Nilsson. Parallel Quantum Computation and Quantum Codes. *SIAM Journal on Computing*, 31(3):799–815, January 1998. 130
- [MOH<sup>+</sup>09] P. Maunz, S. Olmschenk, D. Hayes, D. N. Matsukevich, L. M. Duan, and C. Monroe. Heralded quantum gate between remote quantum memories. *Physical Review Letters*, 102(25):250502, June 2009. 91
- [MRR<sup>+</sup>14] 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, February 2014. 34, 132
- [MS98] K. Molmer and A. Sorensen. Multi-particle entanglement of hot trapped ions. *Physical Review Letters*, 82(9):4, March 1998. 146
- [MSB<sup>+</sup>11] 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):1–4, March 2011. 33, 146, 147, 148

- [MSW<sup>+</sup>08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100(20):2–5, May 2008. 33
- [MVL<sup>+</sup>11] 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(10):103005, October 2011. 82, 83
- [NC02] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2002. 32, 129
- [NKM<sup>+</sup>10] R. Noek, C. Knoernschild, J. Migacz, T. Kim, P. Maunz, T. Merrill, H. Hayden, C. S. Pai, and J. Kim. Multiscale optics for enhanced light collection from a point source. *Optics letters*, 35(14):2460–2462, July 2010. 82, 83
- [NLK<sup>+</sup>14] M. Niedermayr, K. Lakhmanskiy, M. Kumph, S. Partel, J. Edlinger, M. Brownnutt, and R. Blatt. Cryogenic surface ion trap based on intrinsic silicon. *New Journal of Physics*, 16(11):113068, November 2014. 110, 118
- [NLR<sup>+</sup>99] H. Nägerl, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Laser addressing of individual ions in a linear ion trap. *Physical Review A*, 60(1):145–148, July 1999. 33
- [NMH<sup>+</sup>15] F. Najafi, J. Mower, N. C. Harris, F. Bellei, A. Dane, C. Lee, X. Hu, P. Kharel, F. Marsili, S. Assefa, K. K. Berggren, and D. Englund.



- On-chip detection of on-classical light by scalable integration of single-photon detectors. *Nature Communications*, 6:5873, 2015. 25, 160, 161
- [NRBC05] C. L. Niclass, A. Rochas, P. A. Besse, and E. Charbon. Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes. *IEEE Journal of Solid-State Circuits*, 40(9):1847–1854, 2005. 159
- [OKH<sup>+</sup>11] J. S. Orcutt, A. Khilo, C. W. Holzwarth, M. A. Popović, H. Li, J. Sun, T. Bonifield, R. Hollingsworth, F. X. Kärtner, H. I. Smith, V. Stojanović, and R. J. Ram. Nanophotonic integration in state-of-the-art CMOS foundries. *Optics express*, 19(3):2335–2346, January 2011. 25, 120, 154, 155
- [OMM<sup>+</sup>09] S. Olmschenk, D. N. Matsukevich, P. Maunz, D. Hayes, L.-M. Duan, and C. Monroe. Quantum teleportation between distant matter qubits. *Science*, 323(5913):486–489, 2009. 32
- [OMS<sup>+</sup>12] J. S. Orcutt, B. Moss, C. Sun, J. Leu, M. Georgas, J. Shainline, E. Zraggen, H. Li, J. Sun, M. Weaver, S. Urošević, M. Popović, R. J. Ram, and V. Stojanović. Open foundry platform for high-performance electronic-photonic integration. *Optics express*, 20(11):12222, May 2012. 154, 155
- [Pau90] W. Paul. Electromagnetic traps for charged and neutral particles. *Reviews of Modern Physics*, 62(3):531–540, 1990. 42
- [PAY<sup>+</sup>09] Y. A. Pashkin, O. Astafiev, T. Yamamoto, Y. Nakamura, and J. S. Tsai. Josephson charge qubits: A brief review. *Quantum Information Processing*, 8(2-3):55–80, February 2009. 94

- [PH53] W. Paul and S. H. A. A new mass spectrometer without a magnetic field. *Naturforsch*, 8a(8):448–450, July 1953. 42
- [PJT<sup>+</sup>05] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309(5744):2180–2184, September 2005. 31
- [PMO09] A. Politi, J. C. F. Matthews, and J. L. O’Brien. Shor’s quantum factoring algorithm on a photonic chip. *Science*, 325(5945):1221, September 2009. 31
- [PSG08] S. E. Pollack, S. Schlamming, and J. H. Gundlach. Temporal extent of surface potentials between closely spaced metals. *Physical Review Letters*, 101(7):071101, August 2008. 94
- [RCB<sup>+</sup>07] M. Riebe, M. Chwalla, J. Benhelm, H. Häffner, W. Hänsel, C. F. Roos, and R. Blatt. Quantum teleportation with atoms: Quantum process tomography. *New Journal of Physics*, 9(7):211–211, July 2007. 32
- [RDN<sup>+</sup>12] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Realization of three-qubit quantum error correction with superconducting circuits. *Nature*, 482(7385):382–385, February 2012. 31
- [Rie05] M. Riebe. *Preparation of Entangled States and Quantum Teleportation with Atomic Qubits*. PhD Thesis, Department of Physics, University of Innsbruck, 2005. 149
- [RJH<sup>+</sup>09] A. Reina, X. Jia, J. Ho, D. Nezich, H. Son, V. Bulovic, M. S. Dresselhaus, and K. Jing. Large area, few-layer graphene films on arbitrary

- substrates by chemical vapor deposition. *Nano Letters*, 9(1):30–35, January 2009. 97, 99
- [RKM<sup>+</sup>01] M. A. Rowe, D. Kielpinski, V. Meyer, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland. Experimental violation of a Bell’s inequality with efficient detection. *Nature*, 409(6822):791–794, 2001. 33
- [RLK<sup>+</sup>06] R. Reichle, D. Leibfried, E. Knill, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Experimental purification of two-atom entanglement. *Nature*, 443(7113):838–841, 2006. 153
- [Rod02] P. Rodriguez. A radix-2 FFT algorithm for Modern Single Instruction Multiple Data (SIMD) architectures. In *2002 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 3, pages III–3220–III–3223. IEEE, May 2002. 124
- [Roo00] C. Roos. *Controlling the quantum state of trapped ions*. PhD Thesis, Department of Physics, University of Innsbruck, 2000. 51, 52, 77
- [RTJ<sup>+</sup>09] A. Reina, S. Thiele, X. Jia, S. Bhaviripudi, M. S. Dresselhaus, J. A. Schaefer, and J. Kong. Growth of large-area single- and Bi-layer graphene by controlled carbon precipitation on polycrystalline Ni surfaces. *Nano Research*, 2(6):509–516, March 2009. 97, 99
- [RWK<sup>+</sup>14] T. Ruster, C. Warschburger, H. Kaufmann, C. T. Schmiegelow, A. Walther, M. Hettrich, A. Pfister, V. Kaushal, F. Schmidt-Kaler, and U. G. Poschinger. Experimental realization of fast ion separation in segmented Paul traps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 033410(90):1–11, 2014. 153

- [RZR<sup>+</sup>99] C. Roos, T. Zeiger, H. Rohde, H. C. Naegerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt. Quantum state engineering on an optical transition and decoherence in a Paul trap. *Physical Review Letters*, 83(23):4, December 1999. 32, 33
- [SACS10] E. Sutter, P. Albrecht, F. E. Camino, and P. Sutter. Monolayer graphene as ultimate chemical passivation layer for arbitrarily shaped metal surfaces. *Carbon*, 48(15):4414–4420, December 2010. 95, 96, 97, 105
- [SBM<sup>+</sup>11] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt. Experimental repetitive quantum error correction. *Science*, 332(6033):1059–1061, May 2011. 33, 146
- [SCB<sup>+</sup>13] A. Stute, B. Casabone, B. Brandstätter, K. Friebe, T. E. Northup, and R. Blatt. Quantum-state transfer from an ion to a photon. *Nature Photonics*, 7(3):219–222, March 2013. 32
- [SCK<sup>+</sup>11] G. Shu, C. K. Chou, N. Kurz, M. R. Dietrich, and B. B. Blinov. Efficient fluorescence collection and ion imaging with the “tack” ion trap. *Journal of the Optical Society of America B*, 28(12):9, November 2011. 82
- [SCR<sup>+</sup>06] S. Seidelin, J. Chiaverini, R. Reichle, J. J. Bollinger, D. Leibfried, J. Britton, J. H. Wesenberg, R. B. Blakestad, R. J. Epstein, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Physical Review Letters*, 96(25):1–4, June 2006. 34
- [SDA<sup>+</sup>12] E. Simoen, M. G. C. De Andrade, M. Aoulaiche, N. Collaert, and C. Claeys. Low-Frequency-Noise Investigation of n-Channel Bulk Fin-

- FETs Developed for One-Transistor Memory Cells. *IEEE Transactions on Electron Devices*, 59(5):1272–1278, May 2012. 94
- [SFB<sup>+</sup>10] D. Stick, K. M. Fortier, M. G. Blain, C. Highstrete, D. L. Moehring, R. Haltli, and C. Tigges. Demonstration of a microfabricated surface electrode ion trap. *arXiv:1008.0990*, August 2010. 110, 118
- [SFS08] P. W. Sutter, J.-I. Flege, and E. A. Sutter. Epitaxial graphene on ruthenium. *Nature Materials*, 7(5):406–411, May 2008. 97
- [Sho94] P. W. Shor. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994. 28, 29
- [Sho95] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):28, October 1995. 29, 121, 125, 129
- [SHO<sup>+</sup>06] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab, and C. Monroe. Ion Trap in a Semiconductor Chip. *Nature Physics*, 2(1):4, December 2006. 34
- [Sin67] R. C. Singleton. On computing the fast Fourier transform. *Communications of the ACM*, 10(10):647–654, October 1967. 124
- [SKC12] J. M. Sage, A. J. Kerman, and J. Chiaverini. Loading of a surface-electrode ion trap from a remote, precooled source. *Physical Review A - Atomic, Molecular, and Optical Physics*, 86(1):013417, July 2012. 116
- [SKHR<sup>+</sup>03] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt. Re-

- alization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422(6930):408–411, March 2003. 33
- [SKK<sup>+</sup>00] C. Sackett, D. Kielpinski, B. King, C. Langer, V. Meyer, C. Myatt, M. Rowe, Q. Turchette, W. Itano, D. Wineland, and C. Monroe. Experimental entanglement of four particles. *Nature*, 404(6775):256–9, March 2000. 33
- [SLM<sup>+</sup>12] J. D. Sterk, L. Luo, T. A. Manning, P. Maunz, and C. Monroe. Photon collection from a trapped ion-cavity system. *Physical Review A - Atomic, Molecular, and Optical Physics*, 85(6):062308, June 2012. 84
- [SM98] A. Sorensen and K. Molmer. Quantum computation with ions in thermal motion. *Physical Review Letters*, 82(9):5, March 1998. 146
- [SMS<sup>+</sup>01] B. C. Stipe, H. J. Mamin, T. D. Stowe, T. W. Kenny, and D. Rugar. Noncontact Friction and Force Fluctuations Between Closely Spaced Bodies. *Physical Review Letters*, 87(9):096801, August 2001. 94
- [SNJ<sup>+</sup>11] E. W. Streed, B. G. Norton, A. Jechow, T. J. Weinhold, and D. Kielpinski. Imaging of Trapped Ions with a Microfabricated Optic for Quantum Information Processing. *Physical Review Letters*, 106(1):010502, 2011. 25, 82, 160
- [SNM<sup>+</sup>13] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt. A quantum information processor with trapped ions. *New Journal of Physics*, 15(12):123012, December 2013. 33, 146, 147

- [SNRWS11] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour. Microscopic model of electric-field-noise heating in ion traps. *Physical Review A - Atomic, Molecular, and Optical Physics*, 84(2):023412, August 2011. 95
- [SNTK03] A. I. Stognij, N. N. Novitskii, S. D. Tushina, and S. V. Kalinnikov. Preparation of ultrathin gold films by oxygen-ion sputtering and their optical properties. *Technical Physics*, 48(6):745–748, June 2003. 85
- [SP13] M. Saeedi and M. Pedram. Linear-depth quantum circuits for n-qubit Toffoli gates with no ancilla. *Physical Review A - Atomic, Molecular, and Optical Physics*, 87(6):062318, June 2013. 129
- [SPM<sup>+</sup>10] K. Singer, U. Poschinger, M. Murphy, P. Ivanov, F. Ziesel, T. Calarco, and F. Schmidt-Kaler. Colloquium: Trapped ions as quantum bits: Essential numerical tools. *Reviews of Modern Physics*, 82(3):2609–2632, September 2010. 45
- [SRW<sup>+</sup>14] R. C. Sterling, H. Rattanasonti, S. Weidt, K. Lake, P. Srinivasan, S. C. Webster, M. Kraft, and W. K. Hensinger. Fabrication and operation of a two-dimensional ion-trap lattice on a high-voltage microchip. *Nature Communications*, 5:3637, January 2014. 110, 118
- [Ste86] S. Stenholm. The semiclassical theory of laser cooling. *Reviews of Modern Physics*, 58(3):699–739, July 1986. 55
- [Ste02] A. M. Steane. Overhead and noise threshold of fault-tolerant quantum error correction. *Physical Review A*, 68(4):21, October 2002. 29
- [TKK<sup>+</sup>00] Q. A. Turchette, D. 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*, 61(06):24, May 2000. 36, 94, 101
- [TR95] C. Tang and G. Reed. Highly efficient optical phase modulator in SOI waveguides. *Electronics Letters*, 31(6):451, 1995. 157
- [Tur37] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1):230–265, 1937. 30
- [VCA<sup>+</sup>10] A. P. Vandevender, Y. Colombe, J. Amini, D. Leibfried, and D. J. Wineland. Efficient fiber optic detection of trapped ion fluorescence. *Physical Review Letters*, 105(2):023001, July 2010. 45, 83
- [VSB<sup>+</sup>00] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, R. Cleve, and I. L. Chuang. Experimental realization of an order-finding algorithm with an NMR quantum computer. *Physical Review Letters*, 85(25):5452–5455, December 2000. 31
- [VvL03] E. J. D. Vredenburg and K. A. H. van Leeuwen. Laser cooling and trapping visualized. *American Journal of Physics*, 71(8), 2003. 50, 53
- [WAF<sup>+</sup>13] 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, March 2013. 110, 118
- [Wan12] S. X. Wang. *Quantum Gates, Sensors, and Systems with Trapped Ions*. PhD Thesis, Department of Physics, Massachusetts Institute of Technology, 2012. 59, 61, 66, 67, 70, 73, 147



- [WGL<sup>+</sup>10] S. X. Wang, Y. Ge, J. Labaziewicz, E. Dauler, K. Berggren, and I. L. Chuang. Superconducting microfabricated ion traps. *Applied Physics Letters*, 97(24):244102, 2010. 95, 102
- [WHL<sup>+</sup>11] S. X. Wang, G. Hao Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, and I. L. Chuang. Laser-induced charging of microfabricated ion traps. *Journal of Applied Physics*, 110(10):104901, 2011. 84, 86
- [WI79] D. J. Wineland and W. M. Itano. Laser cooling of atoms. *Physical Review A*, 20(4):1521–1540, October 1979. 52, 101
- [WMI<sup>+</sup>97] 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, May 1997. 32, 33, 34, 43, 121, 132
- [WMM<sup>+</sup>96] D. Wineland, D. Meekhof, C. Monroe, B. King, W. Itano, and D. Leibfried. Nonclassical motional states of a trapped atom. *Summaries of Papers Presented at the Quantum Electronics and Laser Science Conference*, 76(11):1796–1799, March 1996. 57
- [WSGS12] G. Wilpers, P. See, P. Gill, and A. G. Sinclair. A monolithic array of three-dimensional ion traps fabricated with conventional semiconductor technology. *Nature Nanotechnology*, 7(9):572–576, September 2012. 110, 118
- [WTD08] X. Wang, S. M. Tabakman, and H. Dai. Atomic layer deposition of metal oxides on pristine and functionalized graphene. *Journal of the American Chemical Society*, 130(26):8152–8153, July 2008. 95

- [WZR<sup>+</sup>12] A. Walther, F. Ziesel, T. Ruster, S. T. Dawkins, K. Ott, M. Hettrich, K. Singer, F. Schmidt-Kaler, and U. Poschinger. Controlling fast transport of cold trapped ions. *Physical Review Letters*, 109(8):1–5, 2012. 132, 153
- [XPT12] C. Xiong, W. H. P. Pernice, and H. X. Tang. Low-loss, silicon integrated, aluminum nitride photonic circuits and their use for electro-optic signal processing. *Nano Letters*, 12(7):3562–3568, 2012. 154, 155