

# Quantum Gates, Sensors, and Systems with Trapped Ions

by

Shannon Xuanyue Wang

B.A.Sc., Engineering Physics  
University of British Columbia (2007)

Submitted to the Department of Physics  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

© Massachusetts Institute of Technology 2012. All rights reserved.



Author .....  
Department of Physics  
May 18, 2012

Handwritten initials, possibly "S.W.", in a simple, sketchy style.

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

A handwritten signature in black ink, appearing to be "K. Rajagopal".

Accepted by .....  
Krishna Rajagopal  
Chair, Associate Department Head for Education  
Professor of Physics



# Quantum Gates, Sensors, and Systems with Trapped Ions

by

Shannon Xuanyue Wang

Submitted to the Department of Physics  
on May 18, 2012, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Physics

## Abstract

Quantum information science promises a host of new and useful applications in communication, simulation, and computational algorithms. Trapped atomic ions are one of the leading physical systems with potential to implement a large-scale quantum information system, but many challenges still remain. This thesis describes some experimental approaches to address several such challenges broadly organized under three themes: gates, sensors, and systems.

Quantum logic gates are the fundamental building blocks for quantum algorithms. Although they have been demonstrated with trapped ions previously, scalability requires miniaturizing ion traps by using a surface-electrode geometry. Using a single ion in a surface-electrode trap, we perform a two-qubit entangling gate and fully characterize it via quantum process tomography, as an initial validation of surface-electrode ion traps for quantum information processing.

Good logic gates are often good sensors for fast fluctuations and energy changes in their environment. Trapped ions are sensitive to fluctuating and static charges, leading to motional state decoherence (heating) and instabilities, problems exacerbated by the surface-electrode geometry. We investigate the material dependence of heating, specifically with aluminum and superconducting traps, to elucidate the physical origin of these fluctuating charges. Static charging is hypothesized to be caused by the trapping and cooling lasers due to the photoelectric effect. We perform systematic experiments with aluminum, gold, and copper traps with lasers at various wavelengths to validate this hypothesis.

Realizing quantum processors at the system level requires models and tools for predicting system performance, demonstration of good classical and quantum control, and techniques for integrating different quantum systems. We develop a modeling system for trapped ion quantum computing experiments and simulate the effect of physical and technical noise sources on practical realizations of quantum algorithms in a trapped ion system. We experimentally demonstrate several such algorithms, including the quantum Fourier transform, order-finding, and Shor's algorithm on up to 5 ions. These experiments highlight several unique advantages of ion trap systems and help identify needs for further development. Finally, we explore the integration of ion traps with optical elements including mirrors and photon detectors as key elements in creating future hybrid quantum systems.

Thesis Supervisor: Isaac L. Chuang

Title: Professor of Physics

Professor of Electrical Engineering and Computer Science





## Acknowledgments

This thesis work would not have been possible without the contributions and influences of many people. First of all I would like to thank my research advisor, Prof. Isaac Chuang, for his guidance and support during the 5 years of my PhD. His knowledge and enthusiasm about every topic in physics, engineering, and computer science was a great inspiration in light of the multi-disciplinary nature of my thesis work. Ike gave me the freedom to pursue new ideas, the encouragement to take risks and go beyond my comfort zone, and had faith in my abilities even when I had doubts. I especially appreciate his emphasis on developing my communication and management skills in addition to scientific skills.

I'd also like to thank the other members of my thesis committee: Vladan Vuletic, Ray Ashoori, and Karl Berggren. Karl was especially helpful in providing an engineering perspective in this field.

I have benefited from many scientific discussions and social interactions with all my colleagues at the Quanta Lab. In the beginning, Jaroslaw Labaziewicz was instrumental in bringing me up to speed in the lab and teaching me the essentials of atomic physics and ion trapping. I owe many of my experimental results to the groundwork that he laid out in designing and building the robust and flexible experimental apparatus and hardware/software infrastructure in the lab. He has also been a wonderful life coach and remains a good friend and mentor. During my first years in the group, I also had the pleasure of meeting David Leibbrandt, Rob Clark, Ruth Shewmon, Liz George, and Carter Lin, with whom we shared many fruitful scientific discussions and fun times outside of the lab.

The next generations of Quanta Lab members shaped the bulk of my remaining PhD years. Peter Herskind's arrival and guidance on the finer details of ion trapping during my second year and beyond was crucial after the previous generations of students had graduated. Peter's mastery of all things optics and meticulous editing of manuscripts were much appreciated. Yufei Ge fabricated most of the ion traps that were to become the core of my experimental work. Nathan Lachenmyer, Eric Dauler, Adam McCaughan, and Gleb Akselrod also assisted with trap fabrication at various times. Guang Hao Low performed key calculations for the laser-induced charging project. Jeffrey Russom constructed the new 461 nm laser and moved the experiment to its more permanent home on the optical table. Towards the end of my PhD, Amira Eltony took the lead in implementing the indium-tin oxide trap project. Hans Andersen contributed valuable proofreading on this thesis document. Over the years I've also enjoyed the company of other group members including: Paul Antohi, Xie Chen, Arghavan Safavi-Naini, Tony Kim, Molu Shi, David Meyer, Arolyn Conwill, Michael Gutierrez, Anders Mortensen, Tony Zhu, and Rich Rines.

I had rewarding experiences working with external collaborators. Prof. Kazuhiro Hayasaka and Prof. Shinji Urabe graciously welcomed me to their labs at NICT Kobe and Osaka University, and I owe a big part of my hands-on knowledge about laser systems to the times I spent there. With visitors to our lab at MIT, I also benefited from rewarding discussions as parts of experimental collaborations on trap testing: Richard Slusher, Jason Amini, Fayaz Shaikh, and Michael Ramm.

The extended collaboration with Prof. Rainer Blatt's group at the University of Innsbruck on the SQIP project deserves a special mention. Thomas Monz was instrumental

in providing the communication channel that made this collaboration as fruitful as it was. The entire team at Innsbruck, including Thomas, Philipp Schindler, Daniel Nigg, Matthias Brandl, Julio Barreiro, and Esteban Martinez, welcomed me during my stays there and gave me the chance to fully participate and learn the many fascinating details of their experiment. I feel privileged to have had this opportunity to work with their state-of-the-art ion trap quantum computer.

The success of my PhD experience owe a large part to the groundwork laid out by the numerous wonderful teachers and mentors I've had during high school and undergraduate years. At Magee High School, Stefan Merchant was the very first physics teacher I had and deserves full credit for providing the initial inspiration and motivation that launched me into this chosen area of study. Tom Katronis was equally instrumental in encouraging me to pursue extra academic opportunities that led to me becoming an International Physics Olympiad participant. At UBC, Prof. Andrea Damascelli provided me the first experience of working in an experimental physics lab, and the early mentorship and deep technical knowledge of Jeff Mottershead were invaluable. At University of Toronto, my summer in Prof. Aephraim Steinberg's lab helped me identify atomic physics as the system of choice to study, and I appreciated the advice on grad school and life from Chris Ellenor.

I would like to thank the MIT music department, including the Chamber Music Society, Chamber Chorus, and Emerson Scholarship program for providing a creative outlet during my time at MIT. Special thanks to Bill Cutter, David Deveau, Jean Rife, and Channing Yu for inspirations and encouragement to take my musical abilities to the next level. Outside of the department, the MIT G&S Players, the Cambridge Symphony Orchestra, Mercury Orchestra, and Lowell House Opera provided me with similar rewarding opportunities to meet and collaborate with musicians from all walks of life. There is very much in common in what it takes to become a better musician and a better scientist, and my PhD training was enhanced by the lessons I learned from making music.

I would also like to thank the many friends I made at MIT for the companionship and support: Gleb Akselrod, Kathy Miu, Ankur Mani, Ila Sheren, Cody Gilleland, Yvonne Moret, Lin Lee Cheong. And finally, I owe everything to the unfailing support and love from my family. I am forever indebted to my mother for the great sacrifices she made in exchange for bringing me to Canada and allowing me to pursue my dreams.

# Contents

<b>1</b>	<b>Introduction</b>	<b>17</b>
1.1	Classical computing . . . . .	18
1.1.1	A brief history . . . . .	18
1.1.2	Limitations . . . . .	20
1.2	Quantum information processing . . . . .	20
1.2.1	DiVincenzo criteria . . . . .	21
1.2.2	Physical implementations . . . . .	21
1.3	Trapped ion quantum computing . . . . .	23
1.3.1	Current state-of-the-art . . . . .	23
1.3.2	Surface-electrode ion trap . . . . .	24
1.4	Challenges to scaling . . . . .	25
1.4.1	Gates in surface-electrode traps . . . . .	25
1.4.2	Fluctuating and static charges . . . . .	26
1.4.3	System design and integration . . . . .	26
1.5	Contributions of this work . . . . .	27
1.6	Outline . . . . .	29
1.7	Contributions of coworkers . . . . .	30
1.8	Publications . . . . .	31
<b>2</b>	<b>Ion trapping fundamentals</b>	<b>33</b>
2.1	Ion traps . . . . .	33
2.1.1	Paul trap . . . . .	33
2.1.2	Surface-electrode ion trap . . . . .	35
2.2	$\text{Sr}^+$ qubit . . . . .	36
2.2.1	Laser-ion interactions . . . . .	36
2.2.2	Atomic structure . . . . .	38
2.2.3	Motional state . . . . .	39
2.3	Quantum state manipulation . . . . .	40
2.3.1	State initialization . . . . .	40
2.3.2	State detection . . . . .	41
2.3.3	Quantum operations . . . . .	42

<b>3</b>	<b>Experimental Apparatus</b>	<b>47</b>
3.1	Cryostat . . . . .	47
3.2	Laser system . . . . .	48
3.3	Surface-electrode trap . . . . .	53
3.4	Experiment control . . . . .	55
<b>I</b>	<b>Gates</b>	<b>57</b>
<b>4</b>	<b>Controlled-NOT gate in a surface trap</b>	<b>59</b>
4.1	CNOT gate . . . . .	60
4.2	Coherence . . . . .	60
4.2.1	Motional state coherence . . . . .	60
4.2.2	Magnetic field stabilization . . . . .	61
4.3	Process tomography on a single ion . . . . .	63
4.3.1	Two-qubit state tomography for one ion . . . . .	63
4.3.2	Process tomography: operator definitions . . . . .	64
4.3.3	State preparation . . . . .	65
4.3.4	Complete basis of measurements . . . . .	65
4.3.5	Construction of the process matrix . . . . .	67
4.4	Results . . . . .	67
4.4.1	Gate performance . . . . .	67
4.4.2	Simulation and error sources . . . . .	68
4.5	Outlook . . . . .	69
4.5.1	Disadvantages and alternatives . . . . .	70
<b>II</b>	<b>Sensors</b>	<b>71</b>
<b>5</b>	<b>Heating rates</b>	<b>73</b>
5.1	Electric field noise . . . . .	74
5.1.1	Models of noise sources . . . . .	74
5.1.2	Measurement method . . . . .	74
5.2	Aluminum traps . . . . .	76
5.2.1	Fabrication . . . . .	76
5.2.2	Heating rate . . . . .	77
5.2.3	SMIT-II trap . . . . .	79
5.2.4	Summary . . . . .	81
5.3	Superconductor . . . . .	81
5.3.1	Fabrication . . . . .	82
5.3.2	Superconductivity . . . . .	83
5.3.3	Heating rate . . . . .	85
5.3.4	Summary . . . . .	86
5.4	Discussion . . . . .	87

<b>6</b>	<b>Laser-induced charging</b>	<b>89</b>
6.1	Model . . . . .	90
6.1.1	Micromotion and ion displacement . . . . .	91
6.1.2	Ion displacement and charge distribution . . . . .	92
6.1.3	Charge accumulation & dissipation . . . . .	93
6.2	Experiment . . . . .	94
6.2.1	Trap fabrication . . . . .	94
6.2.2	Experimental setup & methods . . . . .	95
6.2.3	Calibration . . . . .	96
6.3	Results . . . . .	96
6.3.1	Material & wavelength dependence . . . . .	96
6.3.2	Quantifying charge . . . . .	97
6.4	Conclusion . . . . .	99
 <b>III Systems</b>		 <b>101</b>
<b>7</b>	<b>TIQC-SPICE</b>	<b>103</b>
7.1	System overview . . . . .	103
7.1.1	System components . . . . .	104
7.1.2	Prior art . . . . .	105
7.1.3	Operator set . . . . .	106
7.2	TIQC simulator . . . . .	108
7.2.1	State evolution . . . . .	108
7.2.2	Output and analysis . . . . .	110
7.2.3	Pulse and pulse sequence . . . . .	111
7.2.4	Decoherence modeling . . . . .	112
7.3	Examples . . . . .	114
7.3.1	Performance . . . . .	119
7.4	In-circuit fidelity . . . . .	120
7.4.1	Definition . . . . .	120
7.4.2	Simulation . . . . .	122
7.4.3	Results . . . . .	124
7.5	Summary and future work . . . . .	125
<b>8</b>	<b>Quantum Algorithms</b>	<b>127</b>
8.1	Motivation . . . . .	128
8.2	Implementations . . . . .	129
8.2.1	Prior art . . . . .	129
8.2.2	TIQC implementation . . . . .	130
8.2.3	Experimental setup . . . . .	130
8.3	Quantum Fourier transform . . . . .	132
8.3.1	Algorithm . . . . .	132
8.3.2	Pulse sequence . . . . .	133

8.3.3	Experiment and simulation . . . . .	134
8.3.4	Semiclassical version . . . . .	138
8.4	Order-finding . . . . .	139
8.4.1	Algorithm . . . . .	140
8.4.2	Pulse sequence . . . . .	141
8.4.3	Experiment and simulation . . . . .	142
8.5	Shor's algorithm . . . . .	146
8.5.1	Algorithm . . . . .	146
8.5.2	Pulse sequence . . . . .	147
8.5.3	Experiment and simulation . . . . .	148
8.6	Summary and future work . . . . .	152
<b>9</b>	<b>Integration with optics</b>	<b>155</b>
9.1	Mirror trap . . . . .	156
9.1.1	Trap design and fabrication . . . . .	156
9.1.2	Results . . . . .	158
9.1.3	Summary . . . . .	159
9.2	Photon detectors . . . . .	159
9.2.1	Indium-tin oxide traps . . . . .	160
9.2.2	Superconducting nanowire single photon detectors . . . . .	164
<b>10</b>	<b>Conclusion</b>	<b>169</b>
<b>A</b>	<b>Ion trap history</b>	<b>173</b>
<b>B</b>	<b>Pulse sequencer</b>	<b>179</b>
B.1	Architecture . . . . .	179
B.2	Instruction set . . . . .	179
B.3	Higher-level constructs . . . . .	181
B.4	Standard pulse sequences . . . . .	182
<b>C</b>	<b>TIQC-SPICE code</b>	<b>191</b>
C.1	TIQC-SPICE module list . . . . .	191
C.2	Example simulations . . . . .	191
<b>D</b>	<b>Pulse sequences for algorithms</b>	<b>197</b>

# List of Figures

1-1	Examples of classical and quantum gates, sensors, and systems. . . . .	18
1-2	Schematic of Paul trap and camera images of trapped ions. . . . .	23
2-1	The four-rod and surface-electrode Paul traps. . . . .	35
2-2	Atomic structure of the $^{88}\text{Sr}^+$ ion. . . . .	39
2-3	Tensor product of atomic and motional states describing the complete Hilbert space for quantum computation with a single ion. . . . .	40
2-4	State detection via the shelving technique. . . . .	41
2-5	Schematic of the Mølmer-Sørensen gate. . . . .	45
3-1	Cross section and photo of the cryostat. . . . .	48
3-2	Block diagram of RF chain for the trap RF electrode. . . . .	49
3-3	Cryogenic ion trap experimental setup. . . . .	49
3-4	Beam distribution optics for the 1091 and 1033 nm lasers. . . . .	50
3-5	Beam distribution optics for the 422 nm laser. . . . .	51
3-6	Beam shaping, cavity lock, and distribution optics for the 674 nm qubit laser. . . . .	51
3-7	Drift of the 674 nm cavity frequency over several years. . . . .	52
3-8	Beam shaping and distribution optics for the 405 nm and 460 nm photoionization lasers. Arrows lead to beam paths for other experiments. . . . .	53
3-9	Laser delivery optics to the cryostat. . . . .	54
3-10	Schematic of the surface-electrode ion trap. . . . .	54
3-11	Block diagram for the experiment and control systems. . . . .	56
4-1	Rabi and Ramsey oscillations on the blue sideband. . . . .	61
4-2	Superconducting disks for magnetic field stabilization. . . . .	63
4-3	Process tomography on the CNOT gate. (a), (b), and (c) show the absolute, real, and imaginary parts of the $\chi$ matrix, respectively. . . . .	68
5-1	Example of heating rate data. (a) Plot of red sideband and blue sideband population as a function of delay time. (b) Average quanta vs delay time. . . . .	75
5-2	Summary of all heating rate data obtained in 7 aluminum traps. . . . .	78
5-3	Heating rate as a function of radial DC compensation voltage for three aluminum traps. Voltages are relative to their mean value. . . . .	79
5-4	Heating rate as a function of oxide thickness for aluminum oxide traps. . . . .	79
5-5	Photographs of the SMIT trap. . . . .	80

5-6	Heating rates measured in the SMIT-II trap. . . . .	81
5-7	Photographs of NbN and Nb traps. . . . .	82
5-8	Microscope image of Nb trap with wire structure for measuring electrode resistivity. . . . .	83
5-9	Resistance and critical current demonstrating superconducting transition .	84
5-10	Summary of all heating rate data obtained in 4 niobium traps. . . . .	85
5-11	Heating rate in normal and superconducting states of a niobium trap. . . .	86
5-12	Heating rates measured in traps Nb-I and Nbg-Ia over a period of 1.5 years.	87
6-1	Diagram of charging experiment setup. . . . .	90
6-2	Block diagram converting micromotion amplitude to ion displacement and total charge. . . . .	91
6-3	Measurement of the micromotion signal. . . . .	92
6-4	Illustration of constants and parameters in the charging model. . . . .	94
6-5	Photographs of (a) aluminum, (b) gold, and (c) copper traps. . . . .	95
6-6	Micromotion amplitude vs voltage on a compensation electrode. . . . .	96
6-7	Typical plot of ion displacement vs time in aluminum, copper, and gold traps.	98
6-8	Charging rate $K$ at $t = 0$ measured in an aluminum trap. . . . .	100
7-1	Trapped ion quantum computing simulator system overview. . . . .	104
7-2	Module dependency diagram for the TIQC simulator. . . . .	108
7-3	Example simulations of standard experiments. The code to generate these plots are in Appendix C. . . . .	116
7-4	Illustration of a generic quantum algorithm for in-circuit fidelity. . . . .	121
7-5	In-circuit fidelity: fidelity vs gate and matching experiment to simulation errors to obtain gate bounds. . . . .	123
8-1	General quantum circuit for Shor's algorithm. . . . .	128
8-2	Atomic structure for the $^{40}\text{Ca}$ ion. . . . .	131
8-3	Quantum circuit for the 3-qubit inverse quantum Fourier transform algorithm.	133
8-4	Optimized pulse sequence for the 3-qubit QFT. . . . .	134
8-5	Initial and final state populations in the QFT pulse sequence for input states of various periods. . . . .	135
8-6	Comparison of population after each pulse for ideal, experiment, simulation in the QFT pulse sequence. . . . .	136
8-7	State fidelity vs pulse for the QFT sequence with input state $\psi =  000\rangle$ for 4 distance measures. . . . .	137
8-8	Semiclassical QFT circuits. . . . .	139
8-9	Measured state populations in the semiclassical QFT pulse sequence for input states of various periods. . . . .	139
8-10	Example of a permutation on 8 elements. . . . .	140
8-11	Quantum circuits for the order-finding algorithm. . . . .	141
8-12	Ideal, experimental, and simulated state populations for the order finding algorithm on 3 qubits. . . . .	144



8-13	(a) Simplified quantum circuit for Shor's algorithm for factoring $N = 15$ . (b) Semiclassical version. . . . .	147
8-14	Circuits for the modular exponentiation in Shor's algorithm. . . . .	147
8-15	State populations measured as the output of Shor's algorithm on $N = 15$ for (a) $a = 11$ and (b) $a = 7$ . . . . .	149
9-1	Photographs of the MirrorTrap. . . . .	157
9-2	Images of various iterations of developing the ITO trap. . . . .	162
9-3	Photograph of trap ITO Ia mounted in a CPGA for testing. . . . .	162
9-4	Ion fluorescence detection setup. . . . .	163
9-5	Detected signals from the ITO trap. . . . .	164



# List of Tables

3.1	Typical beam waists, cryostat viewport transmission, and power at the ion.	53
3.2	Typical trap parameters for traps of various geometries. . . . .	55
4.1	State preparation operations. . . . .	65
4.2	State measurement functions. . . . .	66
4.3	Measured gate fidelities for the identity gate, the single CNOT gate, and two concatenated CNOT gates. . . . .	68
4.4	Error budget listing the major sources of errors on the process fidelity of the single CNOT gate. . . . .	69
5.1	Summary of ion height and fabrication parameters for all aluminum traps tested. . . . .	77
5.2	Lowest heating rate of traps made of superconducting and normal metals measured at cryogenic temperatures. . . . .	85
6.1	Summary of fit parameters from charging data in aluminum, copper, and gold traps. . . . .	99
7.1	Operator set for implementing the quantum algorithms in this chapter. . . .	107
7.2	List of members for the <code>hspace</code> class. . . . .	109
7.3	List of operators in the <code>hspace</code> operator dictionary. . . . .	109
7.4	List of member variables and functions of the <code>database</code> object. . . . .	110
7.5	Pulses defined in the simulation and their mapping to operations. . . . .	111
7.6	List of error sources, variable names within the class <code>parameters</code> , and typical values used for the Innsbruck experiment. . . . .	113
7.7	Run times for simulations with various dimensions and pulse sequence length.	120
7.8	Possible definitions for the difference between simulated and experimental results, $\Delta(F_{\text{sim}}, F_{\text{exp}})$ , to be minimized during ICF evaluation. . . . .	123
7.9	Gate fidelity bounds obtained with ICF for 3 types of gates in two pulse sequences. . . . .	124
7.10	Fidelity bounds on individual gates obtained via ICF from density matrix ( $\rho$ ) and population (p) measurements. . . . .	125
8.1	Capabilities of the ion trap experiment at the University of Innsbruck. . . .	131
8.2	Examples of input strings with different periodicities and the output strings after applying the DFT operation. . . . .	133

8.3	Definitions of various distance measures. . . . .	136
8.4	Error budget for the QFT pulse sequence. . . . .	138
8.5	Representative permutations for order-finding on 4 elements. . . . .	142
8.6	Experimental and simulated fidelities (sso and trace distance) for the population outputs of the order-finding algorithm. . . . .	144
8.7	Measured and simulated fidelities (sso and trace distance) of the QFT output of Shor's algorithm for $N = 15$ , $a = 11, 7$ . . . . .	149
8.8	Error budget on the final population measurement for the Shor $a = 11$ pulse sequence. . . . .	150
8.9	Error budget on the final state fidelity for the Shor $a = 11$ pulse sequence. . . . .	150
9.1	Summary of fabrication parameters and results for all tested ITO traps. . . . .	161
A.1	Gold . . . . .	173
A.2	Aluminum, Aluminum Oxide, Copper . . . . .	174
A.3	Niobium and Niobium Nitride . . . . .	175
A.4	MirrorTrap . . . . .	176
A.5	Indium-tin oxide . . . . .	177
B.1	Instruction Set – DDS controls . . . . .	179
B.2	Instruction Set – Timing controls . . . . .	180
B.3	Instruction Set – Logic and arithmetic . . . . .	180
B.4	Instruction Set – Flow control . . . . .	180
B.5	Higher-level instructions . . . . .	181
C.1	Core modules, classes, and functions in TIQC-SPICE. . . . .	196
D.1	List of symbols used to denote pulse types in pulse sequences. . . . .	197
D.2	Color legend for pulse sequences. . . . .	198

# Chapter 1

## Introduction

Computing has been with humanity since the very beginning. Ever since the age of the abacus, humans were interested in using physical and technical means to solve practical problems and enable new technologies. Today, it is difficult to imagine our society existing and functioning without the constant presence of computers and the technologies that depend on them. While the development of computer algorithms to solve both old and new problems continues to take place, suggestions of an alternate method of computing that promises exponential speedup on certain algorithms have captured much interest. This method relies on the principles of quantum mechanics to describe physical systems and is known as quantum computing. This thesis explores one potential physical device to implement quantum computing, the ion trap. It is motivated by three major themes: gates, sensors, and systems, and bridges the world of traditional computing to the field of quantum computing using these themes.

The technologies that enabled today's computers relied on the development and commercialization of a single physical device, the transistor. The transistor is used to construct logic gates, the fundamental building blocks for computations and algorithms. Trapped ions are used to implement quantum logic gates, which are likewise the fundamental building blocks in quantum algorithms.

Good logic gates employ the least energy and operate fast, and therefore are sensitive to small energy changes and fast fluctuations in their environment. Thus good gates often make good sensors for such effects. Building a logic gate out of a physical device requires detailed measurements to characterize the interactions between the device and its environment. Understanding this interaction, in the case of the transistor for classical computing and the ion trap for quantum computing, is crucial to being able to build better logic gates.

Such knowledge gained with individual devices and prototype systems are used to contribute to concrete models, which lead to tools for designing and predicting the behavior of large-scale systems. For classical computing, such tools were indispensable once the system complexity exceeded the capabilities of analytical models. For quantum computing, continued progress in the development of viable physical systems will require a similar approach in system design and realization.

These three challenges (Figure 1-1) are driven by a need to thoroughly understand the physics of quantum information as a prelude to engineering quantum systems. The goal

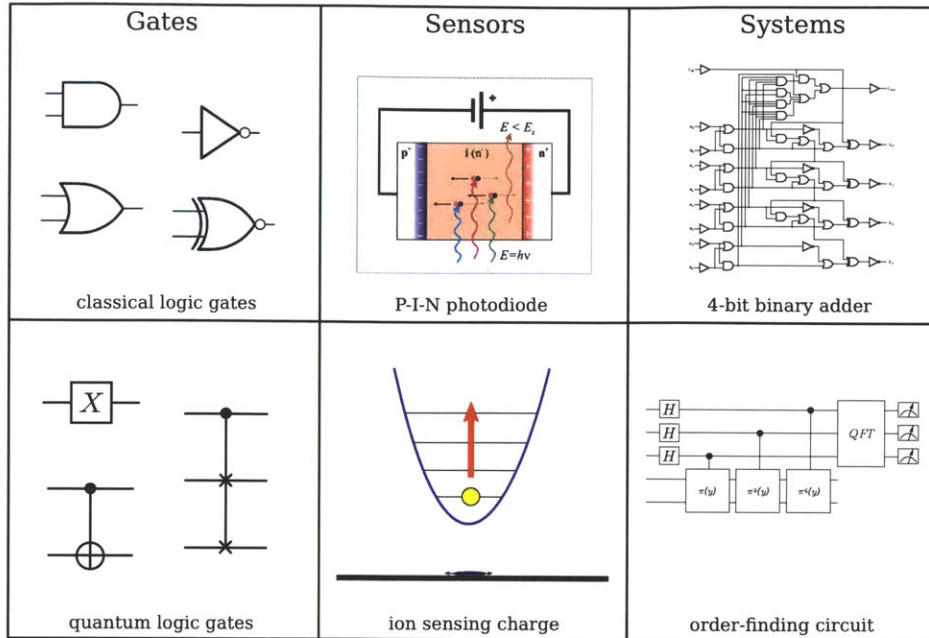


Figure 1-1: Examples of classical and quantum gates, sensors, and systems. Top left: logic gates in Boolean algebra: AND, NOT, OR, XOR. Top middle: a P-I-N photodiode which senses light, based on the same semiconductor physics responsible for transistors. (Image courtesy of Kirnehkrib [CC-BY-SA-3.0] via Wikimedia Commons.) Top right: circuit diagram for a 4-bit binary full adder. Bottom left: logic gates for quantum computation: X, CNOT, Fredkin. Bottom middle: trapped ion whose motional state acts as a sensor for charges on a metal surface. Bottom right: quantum circuit for the order-finding algorithm (discussed in Chapter 8).

of this chapter is to provide a broad perspective on the questions underlying this thesis work, to describe the approach taken to address these questions in the thesis research, and to set the stage for detailed discussions of the main results of this thesis in the following chapters. Section 1.1 begins with a historical overview of the computers we know of today, which I will refer to as classical computers, along with the current status and their limitations. Then Section 1.2 introduces quantum computation, along with the list and status of all physical implementations currently known, providing some perspective for the roles of trapped ions. Sections 1.3 and 1.4 summarizes some of the major challenges facing the trapped ion quantum computing community including some suggested solutions. The remaining sections conclude the chapter with a thesis outline and the contributions I have made in this thesis.

## 1.1 Classical computing

### 1.1.1 A brief history

In the early 1800s, Charles Babbage observed the error-prone process of manually creating logarithm tables and invented the idea of a mechanical computer. For years, he worked

on the design and construction of a mechanical Difference Engine for calculating values of polynomial functions. Subsequently he also designed and tried to construct the Analytical Engine, one consisting of multiple physical machines, including the use of punch cards for input and output. Sadly both of those machines were never completed during Babbage's lifetime, but they were among the first known attempted implementations of concepts familiar to modern computer architectures.

The invention of the transistor in 1947 at Bell Labs is often considered the birth of modern day computing. It was followed by the first silicon version in 1954 and the advancement of semiconductor device fabrication techniques in the next decades. The transistor was the physical device that most successfully implemented logic gates, a concept that existed well before the transistor's invention – they were initially described by Charles Peirce in 1886. Logic gates are the building blocks of Boolean algebra, which is the foundation for the operation of modern computers.

The logic gate is an idealized concept, but a transistor is a physical implementation, whose initial development relied on key concepts in solid-state physics at the electron level. Much theoretical understanding and experimental work was required to produce practical devices that implement the ideal behavior [Sho52]. Still, practical devices suffer from various noise and imperfections such as flicker ( $1/f$ ) noise, stray capacitance and inductance, current leakage, and statistical process variation. Parallel to the development of the technology, extensive studies were taking place to understand the physical origins of various types of noise found in transistors and other solid-state devices [vdZ70, CAV94].

The measurements of the behavior and noise of transistors and similar devices provided data and models that eventually bridged the gap between building a single or several such devices and building a full system capable of general-purpose computing. As the complexity of designs grew, analytical solutions to circuits were becoming impractical, and prototyping every circuit was expensive and time-consuming. Tools to automate the process of designing, simulating, testing, and eventually mass-producing practical devices were essential. For example, SPICE (Simulation Program with Integrated Circuit Emphasis) was a simulation software developed for this purpose, one which explicitly included physical imperfections and noise in the device [NP73, MN73], and was used to predict the behavior of large systems. The sophisticated software toolchain that evolved in parallel with advancements in hardware was instrumental at every stage of the development process for modern processors.

Computers of today are the results of decades of technological development and market forces. A typical 2012 processor contains over 1 billion transistors, a smallest feature size of 22 nm, and a clock speed of 3 GHz capable of orders of magnitude faster computations than the computers of the 1940s, all contained in a package of a few  $\text{cm}^2$  in size and at a cost of  $\sim$ \$200 each. This success story is far from finished: alternative physical implementations of logic gates are being explored, studies of noise in solid-state devices continue to the present day, and the simulation and verification of the next-generation processors prior to manufacturing requires non-trivial amounts of computation hours on the best of the current generation of computers.

### 1.1.2 Limitations

Despite the wide-ranging success of classical computers, their powers are fundamentally limited by quantifiable physical resources. There exists certain classes of problems for which the resources required scales exponentially with system size. One example of such a problem is the simulation of quantum systems. A quantum system consisting of  $N$  quantum objects requires a number of variables exponential in  $N$  to completely specify its state. Such a system quickly exhausts the physical resources of classical computer to store and compute its quantum state evolution. Presently, various approximation techniques such as Hartree-Fock, density functional theory, and others exist to enable solutions for problems of practical interest.

More broadly speaking, there exist hard problems for which no polynomial-time classical algorithms are known. Finding the ground state of the Ising spin model is one such example [BBT09]. Another well-known example is factoring integers. Factoring is believed to be such a hard problem that the RSA public key system, a cryptographic system that guarantees the security of many commercial transactions, relies on the difficulty of factoring large numbers.

## 1.2 Quantum information processing

Richard Feynman was among the first to postulate that a computing device built with help from quantum mechanics might be capable of more than their classical counterparts [Fey82]. Since then, a wealth of proposals and experimental demonstrations have realized such predictions with varying degrees of success. Three prominent examples of quantum information processing and their current status are:

*Quantum Communication.* Cryptography and key distribution can be made provably secure using quantum protocols [BB84]. Commercial QKD systems have been developed by several companies (MagiQ, idQuantique among others), though the market reception has been slow.

*Quantum Simulation.* Many practical problems, such as protein folding, design of drugs and materials, and emergent behaviors in condensed matter physics can be reduced to simulating many-body quantum systems. Such simulations are intractable on classical computers due to the exponentially growing system size, but can be efficiently simulated on quantum systems which naturally store such exponential number of variables. True to Feynman's visions, a variety of physical systems such as magnetic spins have recently been simulated with neutral atom systems and trapped ions [BPT<sup>+</sup>10, LHN<sup>+</sup>11].

*Quantum Algorithms.* A number of quantum algorithms that provides significant speed-ups compared to the currently available classical algorithms have been discovered. A few of the well-known ones include Shor's algorithm for factoring [Sho94], Grover's algorithm for unstructured search [Gro97], and a recent quantum algorithm for solving linear systems of equations [HHL09]. New algorithms are still actively investigated. However, no one has yet constructed a physical quantum computer that is capable of solving problems currently intractable on a classical computer.

How big of a quantum system is needed to exceed the capacity of classical systems? Estimates vary depending on the application. For quantum communication, a few qubits is



sufficient due to its serial nature. For quantum simulations, a quantum system of even tens of qubits would surpass the simulation abilities of classical computers, due to exponential scaling of the number of variables needed to represent a quantum state in a classical computer. For quantum algorithms however, the number of qubits required to surpass classical computers is much greater. For example, clusters of classical computers are able to factor numbers up to about 1000 (classical) bits at the present [AFK<sup>+</sup>07]. The standard quantum circuit implementation of Shor’s algorithm for a 1000-bit number requires thousands of qubits [Sho97]. Above this, the requirements for quantum error correction can easily demand orders of magnitude more qubits [Ste03a].

Implementation of quantum algorithms requires more than a large collection of qubits; it also requires well-defined methods to manipulate and measure them with high enough fidelity. DiVincenzo succinctly formulated a short list of such requirements to guide the choice of a physical system for realizing quantum information processing.

### 1.2.1 DiVincenzo criteria

DiVincenzo stated the following five requirements for the physical realization of quantum computation [DiV00]:

1. A scalable physical system with well characterized qubits.
2. The ability to initialize the qubits to a well-defined state, such as  $|000\dots\rangle$ .
3. Long relevant coherence times, much longer than the gate operation time.
4. A “universal” set of quantum gates.
5. A qubit-specific measurement capability

Here, coherence time refers to the ability of a qubit to remain in a desired quantum state over a certain time scale. A universal gate set means a set of gate operations which can be combined to construct any arbitrary quantum operation.

Two additional criteria are named for building quantum networks:

6. the ability to interconvert stationary and flying qubits.
7. the ability to faithfully transmit flying qubits between specified locations.

Here, a flying qubit refers to quantum information that can be transmitted over a physical distance.

Various physical implementations satisfying these criteria have been proposed and tested. Section 1.2.2 describes the basic principles and current status of some of these physical systems.

### 1.2.2 Physical implementations

One of the earliest demonstrations of various quantum computing primitives was with nuclear magnetic resonance (NMR). Nuclear spins of molecules in a liquid solution are used for qubits, and are controlled via radiofrequency pulses. Typical two-qubit gate times are on the order of ms, while coherence times are on the order of seconds [Ste03b]. Complete quantum algorithms including Grover’s search [VSS<sup>+</sup>99] and Shor’s algorithm for factoring

15 using 7 qubits [VSB<sup>+</sup>01] have been demonstrated with NMR quantum computing. However, numerous limitations to scaling NMR systems are known. For example, to increase the number of qubits, one needs to synthesize new molecules; in addition, the coherence time, probability of ground state cooling, frequency selectivity and total signal all decrease rapidly as the system size grows [NC00]. It also does not allow fast qubit measurements nor fast classical feedback from quantum measurements, and thus cannot readily be made fault-tolerant to quantum noise.

Another early implementation uses optical photons as qubits, where the spatial mode or the photon's polarization is used to store the quantum information. Optical components such as phase shifters, beamsplitters, and phase modulators are used to perform operations. A typical coherence time is 0.1 ms [LJL<sup>+</sup>10], and gate times are much faster. The major challenge with optical photon implementation is that photons interact weakly with each other, and efficient generation and detection of single photons is difficult. Nevertheless, photons are optimized for traveling long distances and are thus useful for realizing quantum communication. Integrated photonics for quantum computing applications remains an active area of research today and recent progress includes examples of quantum algorithms such as order-finding [LLL<sup>+</sup>11] and Shor's algorithm [PMO09] with 4 qubits.

Quantum computing systems based on solid-state technologies are appealing due to the ability to leverage fabrication techniques developed for silicon and transistor technologies in classical computing. Superconducting qubits are one example of such a system. Superconducting qubits consist of Josephson junctions made using conventional microfabrication techniques and are inherently small in physical size and easily scaled to large circuits. Charge, flux, or phase of a circuit consisting of Josephson junctions form the qubit and are controlled via voltage or current. In its early days, superconducting qubits were troubled by short coherence times; however, recent technical advances have improved coherence times to be orders of magnitude longer than typical gate times ( $\sim 10 \mu\text{s}$  and 10 ns respectively [PSB<sup>+</sup>11]), enabling high gate fidelities [CGT<sup>+</sup>09], demonstrations of various algorithms [DCG<sup>+</sup>09a] and quantum error correction [RDN<sup>+</sup>12] with up to 3 qubits.

Among the possible systems for realizing quantum computing, trapped ions were one of the first to make significant progress in satisfying the DiVincenzo criteria. Electronic or nuclear spin states of ions represent qubits, which are manipulated using laser pulses. Numerous advantages of ion traps include long lifetimes, natural reproducibility of the qubits, and long coherence times (ms-s depending on qubit type) compared to gate times ( $\sim \mu\text{s}$ ). These features have enabled many basic quantum algorithms, such as the Deutsch-Jozsa algorithm [GRL<sup>+</sup>03] and Grover search [BHL<sup>+</sup>05], to be realized with a few ions in the past decade. The largest quantum register used for algorithms so far is 5 qubits [BMS<sup>+</sup>11], but entangled states of up to 14 qubits have been achieved [MSB<sup>+</sup>11]. However, just like in other physical systems, there are also unique challenges in using trapped ions for quantum computing. The next two sections detail the current status of this field and some of the challenges that lie ahead.

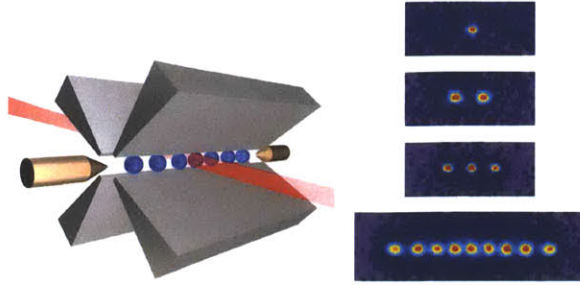


Figure 1-2: Left: schematic of a Paul trap consisting of 4 DC/RF electrodes and endcap electrodes to confine ions in a linear chain. Right: camera images of few ions in a Paul trap. Images courtesy of University of Innsbruck.

### 1.3 Trapped ion quantum computing

Ion traps, more specifically the Paul trap, employs oscillating electric fields to create a quasistatic, harmonic potential in which single or a few isolated ions can be confined (Figure 1-2). The strong interaction available for an electrostatic charge, combined with the isolation of the trapped ions from the environment, leads to large trap depths and therefore long lifetimes. Control over the ions' internal states is mediated by atom-laser interactions. The precise control techniques developed for manipulating the atomic and motional states of ions have led to applications such as spectroscopy [BSW82], precision measurements [RHS<sup>+</sup>08], and atomic clocks [MBH<sup>+</sup>04, DUB<sup>+</sup>01].

Cirac and Zoller were among the first to propose the use of trapped ions to implement quantum computing [CZ95]. Their proposal defined the internal states of the ions as qubits, and called upon the precise laser control of the ion's state developed over the past decades to manipulate these states and perform quantum operations. The common motional modes of more than one ion trapped in the same trap provides means to couple multiple ions and create entanglement. A number of similar proposals followed, and soon experimental realizations of basic components of quantum computing were achieved [HRB08, BW08], as summarized below.

#### 1.3.1 Current state-of-the-art

Since the Cirac-Zoller proposal, the basic DiVincenzo criteria have been satisfied and much progress has been made in implementing quantum operations and algorithms with trapped ions:

1. A scalable physical system with well characterized qubits.  
As stated in the previous section, trapped ions are well-isolated systems and in most cases their spin states are well-characterized. The motional states of ions trapped in a harmonic potential provide means to couple different ions in the same trap.
2. State initialization to  $|000\dots\rangle$ .  
State initialization is done by precisely manipulating the internal states of the ion using laser pulses. State preparation with fidelities exceeding 99.99% has been achieved [RCK<sup>+</sup>06].

3. Long coherence times.

The stability of the state for trapped ions in terms of spontaneous decay can be on the order of seconds for optical qubits and years for hyperfine qubits, so are usually not of concern. More relevant is the phase coherence of qubit states. For hyperfine qubits, coherence time of greater than 10 s, which is much longer than the time required for single- and two-qubit gates (typically 10-100  $\mu$ s), has been demonstrated [LOJ<sup>+</sup>05].

4. A “universal” set of quantum gates.

One possible universal gate set consists of arbitrary single-qubit gates and a two-qubit entangling gate [DiV95]. In trapped ions, single-qubit gates have been demonstrated with fidelities greater than 99.9% [RCK<sup>+</sup>06], and entangling gates with fidelities up to 99.3% [BKRB08].

5. A qubit-specific measurement capability.

Readout of trapped ion states can be performed to high fidelity by measuring population of states using the shelving technique: a laser tuned to address a transition that only couples to one of the two qubit states will scatter photons only if the qubit is in that state. Using this technique as a basis, readout fidelity of up to 99.99% has been achieved [MSW<sup>+</sup>08].

Beyond the DiVincenzo criteria, trapped ions have also implemented various basic algorithms, including the Deutsch-Josza algorithm [GRL<sup>+</sup>03], Grover search [BHL<sup>+</sup>05], teleportation [OMM<sup>+</sup>09], simulation of Ising models [KHC11], and quantum error correction [SBM<sup>+</sup>11].

### 1.3.2 Surface-electrode ion trap

One might naturally question what prevents ion traps from fully implementing large-scale quantum information processing. The answer relates to the word “scaling” in the first DiVincenzo criterion; the excellent fidelities achieved above were all performed on different physical experiments, and the demonstration of algorithms and protocols mostly involved only a few ions. Among the numerous approaches suggested to scaling up ion trap quantum processing, one is to employ well-established microfabrication techniques to make traps of higher density and complexity by arranging electrodes on a plane. These are known as surface-electrode ion traps.

The conventional (“3D”) Paul trap generates the harmonic trapping potential with 4 rod or blade electrodes. Limitations of such a geometry include the difficulty in trapping and controlling ion chains greater than tens of ions, and the challenge in precisely constructing a trap with more complicated geometries such as junctions. Meanwhile, quantum computing architectures based on ion traps have been proposed that involves trapping ions in separate zones and shuttling them into and out of computational/storage zones [KMW02].

Surface-electrode ion traps use the same type of RF and DC voltages to create the quadrupole potential for trapping, but the electrodes to generate the potentials form a 2D geometry. This allows the traps to be made by conventional microfabrication techniques such as laser machining and photolithography, greatly streamlining the development of new and complex trap designs. Early traps used substrates which were laser-machined and

hand-assembled [HOS<sup>+</sup>06, BOV<sup>+</sup>09, SPZSK08]. These traps were used to demonstrate the feasibility of this technique as well as a few key features such as shuttling without heating and transport through junctions. Scaling up these traps to hundreds of trapping zones and junctions would require new methods for constructing traps with high precision and repeatability.

Subsequent designs consisting of a single plane of conductors on a substrate were first proposed in [CBB<sup>+</sup>05]. Proof-of-concepts were realized in the form of printed circuit boards [PLB<sup>+</sup>06, BCL<sup>+</sup>07], while traps made with photolithography were demonstrated in Refs [SCR<sup>+</sup>06, SHO<sup>+</sup>06, LGA<sup>+</sup>08]. Subsequently a large number of groups have succeeded in designing, fabricating, and realizing surface-electrode traps with increasing precision and number of features as fabrication techniques improved [HLBH11]. Complex designs with junctions and tens of trapping zones have been fabricated with CMOS-compatible techniques [LLC<sup>+</sup>09, MHS<sup>+</sup>11]. Traps made of a variety of materials have been tested, including silver [LGA<sup>+</sup>08], gold [LGL<sup>+</sup>08, DNM<sup>+</sup>11], aluminum [LLC<sup>+</sup>09, MHS<sup>+</sup>11], and boron-doped silicon [Bri08]. Furthermore, a variety of traps were developed that have novel features, such as through-wafer loading slots, multi-layer constructions with vias, and integrated capacitors [MHS<sup>+</sup>11]. The design of surface-electrode traps with favourable characteristics and integrated optical/electrical elements remains an active area of research.

An early proposal of an ion trap quantum processor consisted of 300 qubits [Ste07], and another envisioned integrated CMOS electronic controllers and MEMs mirrors with a microfabricated ion trap [KPM<sup>+</sup>05]. Since then, integrating ion traps with interfaces to other control elements and systems have made significant progress. Particularly in the area of coupling ions to photons, traps have been made with integrated mirrors [MVL<sup>+</sup>11], optical fibers [VCA<sup>+</sup>10, KHC11], and lenses [JSN<sup>+</sup>11]. Others have proposed coupling trapped ions to other systems such as photons [Kim08] and superconducting qubits [TRBZ04, SBC<sup>+</sup>11], combining the advantages of each to create a more powerful hybrid system.

## 1.4 Challenges to scaling

Despite these successes in trap development, no one has yet succeeded in realizing a complex quantum algorithm using ions in a surface electrode trap. Many challenges lie ahead in realizing large-scale quantum information processing systems with trapped ions, both technical and fundamental. In this section, a few of these challenges will be described. This is by no means an exhaustive list, but serves to place the work described in this thesis in an appropriate context.

### 1.4.1 Gates in surface-electrode traps

At the beginning of this thesis work, surface-electrode traps were a relatively new technology. Although many successful trap designs and fabrication methods were demonstrated, there were few direct characterizations of quantum operations performed using such traps. It is possible that unknown sources of decoherence will adversely affect attempts to realize quantum operations in surface-electrode traps. Therefore, demonstrating a complete set of

quantum operations, including characterizing and understanding the effects of sources of decoherence on them, is necessary to prove the viability of this approach.

### 1.4.2 Fluctuating and static charges

In surface-electrode traps, the proximity of the trapped ion to its electrodes leads to increased sensitivity to noise that occurs on metal surfaces, in the form of fluctuating and static charges.

In most proposals of multi-ion entanglement gates, the motional state plays a key role in mediating interactions between different ions in a chain [CZ95, SM99]; thus motional state coherence is important to achieve high gate fidelity. Charges which oscillate at the secular frequency of the ions drives them out of the motional ground state, leading to motional state decoherence. The decoherence is several orders of magnitude higher than expected from Johnson noise coming from the trap electrodes [TKK<sup>+</sup>00], thus the term “anomalous heating” is used to describe the effect. In addition, the heating rate appears to scale rapidly with system size [DOS<sup>+</sup>06, TKK<sup>+</sup>00, ESL<sup>+</sup>07], which is highly detrimental to the goal of miniaturizing ion traps. The sources of these fluctuating charges at the secular frequency are not well-understood.

Stray static charges near the ion trap can lead to increased micromotion, which has a number of detrimental effects on the ion [BMB<sup>+</sup>98]. Drifting stray charges is a particular concern since they must be continuously monitored and corrected, adding another experimental overhead. There are many possible sources of stray charge, one of which is laser-induced charging of dielectrics on or near the trap electrodes. In microfabricated traps, the problem is compounded due to the proximity of the lasers to the trap surface. The nature of laser-induced charging needs to be better understood in order to minimize its effect on small ion traps.

### 1.4.3 System design and integration

In classical computing, SPICE is a simulation program that predicted the behavior of large systems from models of individual devices, which were physical objects susceptible to a large variety of physical effects causing noise and imperfections [NP73, MN73]. Simulating all of these physical effects is increasingly impractical as the system size grows. Therefore SPICE and similar programs must make judicious choices of simplifications in a physical model, in order to simultaneously produce results that accurately represent a physical device, and compute them in a practical amount of time [RNMW71, Qua89].

The development of quantum information processing systems will likely soon scale to a point where a comprehensive tool to aid in the system design will be needed. Physical quantum systems, just like physical classical devices, are subject to a great variety of noise sources leading to non-ideal behavior. With quantum systems, there is the additional challenge that the number of classical variables needed to specify the system scale exponentially and thus quickly exceed the capability of classical computers to model them. Nevertheless, even a small-scale simulation of today’s small-scale quantum computing systems would be

useful in identifying important noise sources and provide guidance to the design and development of next-generation systems.

In parallel with designing and simulating quantum systems, the models and tools need to be verified by physical data. Although many quantum algorithms have been demonstrated with a small number of ions and operations, more complex algorithms would demand the best of the currently available classical and quantum control techniques for their implementation. Practical realizations of quantum algorithms with trapped ions would serve to validate our physical models as well as provide important insight on areas of improvement for the next generation physical systems.

The next-generation quantum computing system will likely require coupling ions to other physical qubits in order to combine their advantages to enable further scaling. For example, qubits consisting of photons are able to travel long distances via optical fibers, and can potentially forming key elements of a quantum communication network. Thus it would be desirable to efficiently convert ion qubits to photonic qubits. A related and more technical issue is to maximize detection efficiency of photons scattered by ions, in order to achieve faster detection of quantum states. Both of these goals – efficient detection of ion fluorescence and converting ion qubits to photon qubits – require good coupling between ions and photons and therefore close proximity between optical elements and the trapped ion. Understanding and building ion-photon coupled systems with integrated optical components is a thus a key element in advancing future quantum processing systems.

## 1.5 Contributions of this work

This thesis addresses the three challenges relating to gates, sensors and systems presented in Section 1.4.

### Gates

To demonstrate that surface-electrode ion traps are a viable technology for future large-scale quantum computing, basic quantum operations must be realized and characterized. Using a single ion in a surface-electrode trap, we perform a two-qubit entangling gate, the Cirac-Zoller CNOT gate. Although this gate has been realized in ion traps previously [SKHR<sup>+</sup>03], the objective is to realize it in a surface-electrode trap and fully characterize it, in order to evaluate the feasibility of such a geometry for further development. In particular, we show that noise sources such as anomalous heating are not necessarily the limiting factor in our experiment.

### Sensors

The choice of a trap electrode material can be an important consideration in the context of minimizing anomalous heating and laser-induced charging. The temperature, distance, and frequency dependence of anomalous heating has been studied in the past [LGL<sup>+</sup>08, DOS<sup>+</sup>06], offering possible insights on the physical origin of heating. However, the effect of the trap material has received less attention. Studies of heating rate motivated by different

material properties may offer additional insight into the source of anomalous heating. We investigate heating rates in traps made of aluminum and superconductors, and compare the results to the other metals studied so far.

Laser-induced charging is of a practical concern in performing experiments with surface-electrode traps. Charging of the trap surface is observed in many experiments, but few systematic studies have been done. So far, laser-induced charging is hypothesized to be caused by the photoelectric effect, and thus its magnitude should depend on the trap electrode material and laser wavelength. We investigate this hypothesis by performing experiments on various trap materials and laser wavelengths.

## Systems

In performing quantum gates with trapped ions, numerous sources of decoherence are known and responsible for gate infidelity. These include both physical sources intrinsic to the system, such as dephasing due to magnetic field fluctuations and motional state decoherence, and technical sources resulting from imperfect control, such as laser frequency fluctuations and addressing error. Understanding and quantifying the effect of individual error sources on practical implementations of algorithms can help in designing larger systems and algorithms that minimize the effect of such errors. We develop a modeling system, dubbed TIQC-SPICE (TIQC = Trapped Ion Quantum Computer), which simulates quantum operations in practical ion trap experiments while accounting for important sources of noise.

We realize a few key quantum algorithms with a prototype ion trap quantum computer and explore the role of the design and simulation system in achieving this goal. The quantum Fourier transform is an algorithm component responsible for the speed-up of many quantum algorithms and we perform the algorithm in an ion trap experiment. Then we take advantage of the measurement and feedforward capabilities of ion trap systems to simplify its implementation, and use it to demonstrate order-finding using 3 qubits, and Shor's algorithm using 5 qubits. These experiments represent a small milestone for ion trap quantum computing. More importantly, such demonstrations contribute to our understanding of what is required to carry out quantum algorithms on the next larger scale.

Finally, we explore several possibilities for coupling ions to photons by integrating optical elements with surface-electrode ion traps. For improving the collection efficiency of ion fluorescence, we experiment with a transparent ion trap in order to collect ion fluorescence with an integrated photon detector. We also explore the possibility of integrating ion traps more directly with photon detectors by fabricating the detector and the trap on the same substrate, and point out some key technical challenges in one example of this approach. For coupling ions to photons, one way to realize a strong coupling is to construct an optical cavity around an ion. There are various challenges to this approach, such as the small distances required and the proximity of dielectric materials to the ion which may accumulate charge and cause stray electric fields. We present a first step toward realizing an ion-cavity system by integrating a surface-electrode ion trap directly with a high-finesse mirror.



## 1.6 Outline

This thesis is broadly organized as follows. First, **Chapters 2** and **3** provide some of the background information necessary to describe this work. Chapter 2 includes basic theoretical foundations of the operation of ion traps, atom-laser interactions, and quantum state manipulation. Chapter 3 describes the experimental apparatus that was used for all of the experimental work to follow, except that in Chapter 8. Little new content is found here which was not extensively covered in available references, so the discussion will be brief and focused on results and current status instead of derivations.

The remainder of the thesis is divided into three parts. Part I containing Chapter 4 describes an experiment to implement a quantum logic gate in surface-electrode traps. Part II containing Chapters 5-6 addresses the issue of fluctuating and static charges. Part III including Chapters 7-9 addresses the system-related issues of designing and simulating ion trap based quantum algorithms, the realization of a few example algorithms in an experiment, and integrating ions with optical elements.

### Part I: Gates

**Chapter 4** describes the demonstration of a standard two-ion quantum logic gate in a surface-electrode ion trap. We fully characterize the gate using quantum process tomography, and use simulation to deduce the limit on the gate fidelity in our experiment. In addition to developing the key technical tools for performing quantum operations, we believe that error sources are well-understood in these surface-electrode systems and no major unknowns exist to prevent their adoption for larger-scale experiments. The content of this chapter appears as Ref. [WLG<sup>+</sup>10].

### Part II: Sensors

**Chapter 5** looks at the issue of fluctuating charges, continuing previous work from this group in evaluating various materials for their electric field noise relevant for ion traps, including aluminum and superconducting traps. By studying the material dependence of heating rate, we aim to gain additional insight on possible physical origins of heating: in particular whether fluctuating electric fields causing it exist in the bulk or on the surface of electrodes. The content of Section 5.3 appears as Ref. [WGL<sup>+</sup>10].

**Chapter 6** investigates the issue of static charges, in the form of laser-induced charging. The chapter begins with a postulated physical model that can quantitatively relate the amount of charges to the observed variable, the micromotion of a single trapped ion. Experiments involving traps of different materials and various laser wavelengths are presented, indicating a wavelength and material dependence consistent with the physical model. The content of this chapter appears as Ref. [WLL<sup>+</sup>11].

### Part III: Systems

**Chapter 7** describes an attempt to address the issue of system design and evaluation, by the development of a simulation framework named TIQC-SPICE that includes realistic models

of experimental sources of error. Details of the simulation’s implementation are described, along with examples that demonstrates its applicability to typical ion trap experiments. The last section of the chapter describes another application of TIQC-SPICE, evaluating the in-circuit gate fidelity as an alternative to quantum process tomography for characterizing quantum gates.

**Chapter 8** describes the experimental realizations of several quantum algorithms using trapped ions. These include the 3-qubit quantum Fourier transform, the order-finding algorithm, and the smallest instance of Shor’s algorithm ( $15 = 3 \times 5$ ) using 5 qubits. We also simulate these experiments and the simulation results show good agreement with the experimental data, validating the models used to represent noise and imperfections in the system.

**Chapter 9** discusses two preliminary experiments and one proposal for coupling ions to photons. The first section describes an experiment to microfabricate an ion trap on top of a high-finesse mirror, eventually to become one half of an optical cavity. The second section describes ion traps made of indium-tin oxide, a transparent conductive material, in order to facilitate integrating a photodetector directly with the trap package. The final section describes a proposal to integrate ion traps with a superconducting photon detector.

**Chapter 10** concludes by summarizing the main results of this thesis and suggesting some implications and outlook for the future of trapped ion quantum computing.

## 1.7 Contributions of coworkers

Many co-workers contributed to the work described in this thesis.

The cryogenic ion trapping experiment was built by Waseem Bakr and Paul Antohi. Jaroslaw Labaziewicz built the software control system and laser systems, except the 461 nm laser system. Jeffrey Russom developed the 461 nm laser system and was responsible for moving the experiment from its original support structure to the optical table. Ruth Shewmon and Jaroslaw Labaziewicz constructed the reference cavity and all the supporting stabilization and locking systems for the qubit laser.

In Part I, the experiment to demonstrate the CNOT gate was carried out in close collaboration with Jaroslaw, and initial results appeared in his thesis. The state tomography pulse sequences were developed by Isaac Chuang. I worked on the state preparation pulse sequences, and simulated both pulse sequences in analytical form. The preparation and measurement sequences were implemented in the FPGA hardware by Jaroslaw. The initial numerical simulation of the CNOT gate implementation in Mathematica was done by Jaroslaw, while I extended the code to include tomography and data analysis. The trap used was fabricated by Yufei Ge and Eric Dauler.

In Part II, the aluminum, copper, and gold traps were fabricated by Yufei Ge and Nathan Lachenmyer. The initial superconducting traps were fabricated by Yufei and Eric; I fabricated the superconducting traps with the wire structure, with help on the sputtering step from Adam McCaughan. I carried out the heating rate and charging measurements. Guang Hao Low performed the calculations to calibrate the charging data from micromotion amplitude to ion displacement and electric field.

In Part III, Chapters 7, the TIQC simulator is a descendent of `quantumcomputer`, a simulation program in Matlab written by Hartmut Häffner. The simulator described here is a complete port and re-write from the original Matlab code to Python, though with many design elements preserved. The TIQC simulator is mostly my own work, except for the interface to Scipy's ODE solver and the process tomography module which were written by Thomas Monz, and the data analysis module `qpython` which was written by Philipp Schindler.

Chapter 8 relies heavily on work from colleagues at the University of Innsbruck. The  $^{40}\text{Ca}$  experiment was developed over several generations of postdocs and PhD students, the most recent including Thomas Monz, Philipp Schindler, Daniel Nigg, Julio Barriero, and Esteban Martinez, all of whom worked together to gather the data for the various algorithms presented. The pulse compiler used to generate the pulse sequences was the outcome of Volckmar Nebendahl's Masters thesis. Philipp worked with the pulse compiler to get the actual pulse sequences implemented in the experiment.

In Chapter 9, Peter Herskind conceived and developed the MirrorTrap experiment. The various MirrorTraps were fabricated by myself and Yufei. Measurement of mirror losses after fabrication was performed by Molu Shi. Yufei and Nathan made the early attempts for fabricating indium-tin oxide (ITO) ion traps. I fabricated the later ITO traps, except the sputtering step which was done by Gleb Akselrod. Amira Eltony performed the photodiode and trap characterization experiments. Amira also performed the calculations for estimating the amount of RF pickup on superconducting wires.

## 1.8 Publications

The following is a list of publications I have co-authored during my PhD studies.

- J. Labaziewicz, Y. Ge, D. R. Leibbrandt, S. X. Wang, R. Shewmon, I. L. Chuang, *Temperature Dependence of Electric Field Noise above Gold Surfaces*. Phys. Rev. Lett., vol. 101, 180602 (2008)
- S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, I. L. Chuang, *Individual Addressing of Ions Using Magnetic Field Gradients in a Surface-Electrode Ion Trap*. Appl. Phys. Lett., vol. 94, 094103 (2009)
- S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, I. L. Chuang, *Demonstration of a quantum logic gate in a cryogenic surface-electrode ion trap*. Phys. Rev. A, vol. 81, 062332 (2010)
- S. X. Wang, Y. Ge, J. Labaziewicz, E. Dauler, K. Berggren, I. L. Chuang, *Superconducting microfabricated ion traps*. Appl. Phys. Lett., vol. 97, 244102 (2010)
- P. F. Herskind, S. X. Wang, M. Shi, Y. Ge, M. Cetina, I. L. Chuang, *Microfabricated surface ion trap on a high-finesse optical mirror*. Opt. Lett., vol. 36, 3045-3047 (2011)
- S. X. Wang, G. H. Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, I. L. Chuang,

*Laser-induced charging of microfabricated ion traps.*

J. Appl. Phys., vol. 110, 104901 (2011)

- S. X. Wang, T. Monz, P. Schindler, R. Blatt, I. L. Chuang,  
*TIQC-SPICE: a simulation for trapped ion quantum computing.*  
in preparation (2012)
- T. Monz, S. X. Wang, P. Schindler, R. Blatt, I. L. Chuang,  
*In-Circuit Fidelity Estimation: Gate fidelity bounds from the implementation of complex quantum algorithms.*  
in preparation (2012)
- T. Monz, P. Schindler, D. Nigg, S. X. Wang, I. L. Chuang, R. Blatt,  
*Experimental demonstration of order-finding and Shor's algorithm with trapped ions.*  
in preparation (2012)

## Chapter 2

# Ion trapping fundamentals

In 1953, Wolfgang Paul developed the 3D quadrupole trap that uses oscillating AC electric fields to confine ions [PS53]. The Paul trap was initially used for mass spectrometry, but later was adapted in atomic physics by Hans Dehmelt and others to study and manipulate single or few ions in well-controlled, isolated environments. Recently, the surface-electrode ion trap was proposed [CBB<sup>+</sup>05] and subsequently demonstrated by many groups. Surface-electrode ion traps are amenable to microfabrication, which enables smaller feature sizes, arbitrary scaling of number of trapping sites, and flexibility in electrode layouts. The development of surface-electrode ion traps is an active area of research today.

Lasers are crucial components in ion trap experiments. They are used to cool the ions, and detect and manipulate the ion's quantum state. In 1995, Cirac and Zoller proposed an ion trap quantum computer based on single- and two-qubit operations [CZ95]. Since then, various alternative gate schemes have been proposed [MS99, LDM<sup>+</sup>03] and gate fidelities in excess of 99% have been realized with trapped ions.

This chapter describes the general theory of ion traps. Section 2.1 begins with the fundamental equations describing Paul traps, followed by recent treatment of planar ion traps. Section 2.2 covers concepts in laser-atom interactions, as well as the specific atomic structure of the  $^{88}\text{Sr}^+$  ion. Section 2.3 presents quantum operations and their physical realization in the laser+ion system. The theory described in this chapter have all been covered in depth by many references<sup>1</sup>; thus the discussion here will omit most derivations and focus on introducing key concepts and terminology used in the remainder of this thesis.

## 2.1 Ion traps

### 2.1.1 Paul trap

In a Paul trap, electric potentials are used to confine and trap charged particles at the potential minimum. By Earnshaw's theorem, a static electric field cannot be used to trap ions, but a combination of static and radio frequency (RF) electric field can form a dynamically

---

<sup>1</sup>ion trap: [Deh67, Gho95, WMI<sup>+</sup>98, Hou08, Lei09]; laser-ion interactions: [Roo00, MvdS99, Lab08]; quantum operations: [HRB08, Mon11]

stable trap. The total potential in a Paul trap with RF frequency  $\Omega$  is

$$\phi(x, y, t) = \phi_{dc}(x, y) + \phi_{rf}(x, y) \cos(\Omega t). \quad (2.1)$$

The standard Paul trap has hyperbolic electrodes defined as

$$\frac{x^2 - y^2}{r_0^2} = \pm 1 \quad (2.2)$$

where  $r_0$  is the distance from the origin (defined as the center of all the electrodes; see Figure 2-1 for coordinates). Thus the potential in the center of the electrode is

$$\phi(x, y, t) = \frac{x^2 - y^2}{2r_0^2} (U + V \cos(\Omega t)) \quad (2.3)$$

where  $U$  and  $V$  are RF and DC voltage amplitudes.

The equations of motion of an ion with charge  $Q$  and mass  $m$  is

$$\ddot{r}_i + \frac{2Q}{mr_0^2} (U + V \cos(\Omega t)) r_i = 0 \quad (2.4)$$

where  $r_i = x, y$ . This can be written in the form of Mathieu equations

$$\frac{d^2 x}{d\tau^2} + (a - 2q \cos(2\tau)) x = 0 \quad (2.5)$$

$$\frac{d^2 y}{d\tau^2} - (a - 2q \cos(2\tau)) y = 0. \quad (2.6)$$

The equation is characterized by the dimensionless quantities for time  $\tau$ , RF voltage  $q$ , and DC voltage  $a$ :

$$\tau = \Omega t / 2 \quad (2.7)$$

$$q = \frac{2QV}{mr_0^2 \Omega^2} \quad (2.8)$$

$$a = \frac{4QU}{mr_0^2 \Omega^2}. \quad (2.9)$$

The ion motion is stable for certain regions of the Mathieu parameter space defined by  $q$  and  $a$ . Particularly for  $a = 0$ , the motion is stable for  $0 < q < 0.9$ . In the case of  $a \ll q \ll 1$ , the solution describing the motion of the ion can be decomposed into a slow harmonic motion with frequency  $\omega$  called the secular motion, and a fast driven amplitude modulation with frequency near  $\Omega$  called micromotion. The amplitude of the micromotion is given by  $\frac{q}{2}|r|$  where  $|r|$  is the displacement from the center of the trap. Micromotion can cause unwanted Doppler shifts and adversely affect cooling [BMB<sup>+</sup>98]; thus most experiments seek to minimize the micromotion by adjusting DC potentials such that the ion's equilibrium position is located at the minimum of the RF potential (RF null), a procedure called compensation.

In the secular approximation, we ignore the micromotion and describe the ion's motion

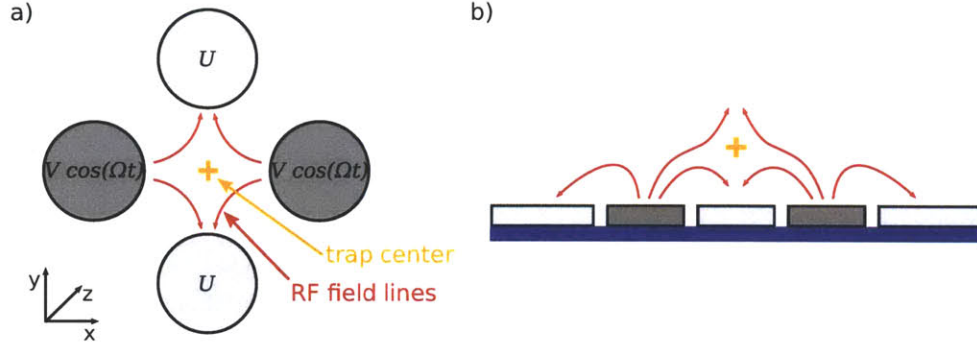


Figure 2-1: (a) Typical geometry of the four-rod Paul trap shown in cross section. Hyperbolic electrodes are commonly approximated by rods with circular cross section. DC and RF voltages share the same reference ground potential and are applied to pairs of electrodes as shown. Confinement in the axial direction is provided by static voltages on a pair of endcap electrodes (not shown). (b) Typical geometry of a surface-electrode trap consisting of 2 RF (dark gray) and 3 DC (light gray) electrodes. The outer DC electrodes can be segmented to provide axial confinement.

as a harmonic oscillator within a pseudopotential

$$Q\phi_p = \frac{Q^2}{4m\Omega^2} |\nabla\phi_{rf}|^2 \quad (2.10)$$

$$= \frac{1}{2} m\omega^2 r_0^2 \quad (2.11)$$

where the secular frequency  $\omega$  is

$$\omega \simeq \frac{q\Omega}{2\sqrt{2}} = \frac{QV}{mr_0^2\Omega\sqrt{2}}. \quad (2.12)$$

Confinement along the  $z$  axis is provided by a pair of electrodes held at a constant DC voltage. This superimposes a weak static confining potential such that the total DC potential is

$$\phi_{dc}(x, y, z) = U \left( \frac{x^2 - y^2}{2r_0^2} + \frac{z^2}{r_0^2} \right). \quad (2.13)$$

### 2.1.2 Surface-electrode ion trap

The required quadrupole trapping potential can also be produced by a 2D set of electrodes [CBB<sup>+</sup>05]. This is commonly referred to as the planar or surface-electrode ion trap. The equations of motion describing charged particles in a surface-electrode trap are similar to those for the 4-rod trap. The electrode geometry is chosen such that the desirable trapping field and trapping potential is created. For any given electrode geometry, this is a boundary-value problem of finding the potential  $\phi$  in the space  $y > 0$  when the potential is held constant on each electrode. Typically this is done numerically with finite element methods [CPO, SPM<sup>+</sup>10]. For geometries consisting of simple rectangular electrodes, an analytical method is also available [Hou08].

The pseudopotential in a surface-electrode ion trap can be described as

$$Q\phi_p = \frac{Q^2}{4m\Omega^2} |\nabla\phi_{rf}|^2 \quad (2.14)$$

$$= \frac{Q^2 V^2}{m\Omega^2} (1/r)^2. \quad (2.15)$$

Here the parameter  $r_0$  representing the electrode spacing is replaced by a general geometric factor  $r$ , which in some cases can be expressed as a closed function of the electrode geometry [Hou08].

## 2.2 Sr<sup>+</sup> qubit

Single trapped ions make ideal qubits due to their long trapping lifetimes, easily controlled internal states with long coherence times, and reproducibility. This section describes the general theory of laser-atom interactions with application to the <sup>88</sup>Sr<sup>+</sup> ion. Section 2.2.1 presents the formalism of the atom-laser interaction Hamiltonian, introducing the idea of using coherent light sources to control the internal state of ions. Section 2.2.2 discusses the specific atomic structure of the <sup>88</sup>Sr<sup>+</sup> ion. Section 2.2.3 describes the motional state of trapped ions, which is a key element in enabling quantum gates with ions.

### 2.2.1 Laser-ion interactions

A two-level atom interacting with a single-mode laser is described by the Hamiltonian

$$H = H_0 + H_1 \quad (2.16)$$

$$H_0 = \omega_0 \sigma_z + \omega_z a^\dagger a \quad (2.17)$$

$$H_1 = \Omega(\sigma_+ + \sigma_-) \left( e^{i(kx - \omega t + \phi)} + e^{-i(kx - \omega t + \phi)} \right) \quad (2.18)$$

where

$\sigma_z, \sigma_+, \sigma_-$  = Pauli spin matrices

$\omega_0$  = energy splitting of the two-level atom

$\omega_z$  = frequency of the harmonic trap; secular frequency

$\Omega = \frac{eE_0}{\hbar} |\langle e|\vec{r} \cdot \hat{\epsilon}|g\rangle|$  Rabi frequency

$\eta$  = Lamb-dicke factor

$\omega = \omega_0 + \delta$  where  $\delta$  is the laser detuning from  $\omega_0$

$k$  = wavenumber of laser

$\phi$  = laser phase.



Here  $\hbar = 1$ . Defining the Lamb-Dicke parameter

$$\eta = k\sqrt{\frac{\hbar}{2m\omega}}, \quad (2.19)$$

the interaction term can be written in terms of creation and annihilation operators:

$$H_1 = \Omega(\sigma_+ + \sigma_-) \left( e^{i(\eta(a+a^\dagger)-\omega t+\phi)} \right) + h.a.. \quad (2.20)$$

The Lamb-Dicke parameter  $\eta$  relates the spatial extent of the lowest harmonic oscillator state to the wavelength of the atomic transition. In the limit that  $\eta \ll 1$  (the Lamb-Dicke regime), the exponent  $e^{i(\eta(a+a^\dagger))}$  can be expanded to get the lowest order terms

$$H_c = \frac{1}{2}\Omega_{n,n}(\sigma^+ + \sigma_-) \quad \text{carrier} \quad (2.21)$$

$$H_r = \frac{1}{2}i\Omega_{n-1,n}(a\sigma^+ - a^\dagger\sigma^-) \quad \text{red sideband} \quad (2.22)$$

$$H_b = \frac{1}{2}i\Omega_{n-1,n}(a^\dagger\sigma^+ - a\sigma^-) \quad \text{blue sideband} \quad (2.23)$$

with Rabi frequencies

$$\Omega_{n,n} = \Omega(1 - \eta^2 n) \quad (2.24)$$

$$\Omega_{n-1,n} = \eta\sqrt{n}\Omega \quad (2.25)$$

$$\Omega_{n+1,n} = \eta\sqrt{n+1}\Omega. \quad (2.26)$$

### Frame transformations

The interaction Hamiltonian is time-dependent and has terms rotating at  $\omega_0$ , which can be removed by transformation into a frame defined by  $H_0$ . In general, to perform a frame transformation from a given laboratory frame  $H$  to a desired frame  $H_f$ , the Hamiltonian in the new frame is given by

$$V = e^{iH_f t}(H - H_f)e^{-iH_f t} \quad (2.27)$$

and the quantum state in the new frame,  $|\phi_f\rangle$  as a function of the original quantum state  $|\phi\rangle$  is given by

$$|\phi_f\rangle = e^{iH_f t}|\phi\rangle. \quad (2.28)$$

We wish to define two frames of reference as follows.  $H_{qc}$ , referred to as the ‘‘quantum computing’’ frame, has the property that states are stationary in the absence of lasers.  $H_{laser}$ , the laser frame, is chosen such that it is time-independent when the laser is on, simplifying evaluation of the state evolution during a laser pulse.

The QC frame is defined by  $H_0$  and the resulting Hamiltonian (making the Lamb-Dicke approximation) is

$$H_{qc} = \Omega \left( (1 + i\eta)(ae^{i\omega_z t} + a^\dagger e^{-i\omega_z t})\sigma^+ e^{-i\delta t} + h.a. \right). \quad (2.29)$$

Thus the atom’s states are stationary in the absence of a laser ( $\Omega = 0$ ).

The laser frame is defined by  $H_L = (\omega_0 + \delta)\sigma_z$ . In this frame, the Hamiltonian is

$$H_{laser} = -\delta\sigma_z + \omega_z a^\dagger a + \Omega \left( (1 + i\eta)(a + a^\dagger)\sigma^+ + h.a. \right). \quad (2.30)$$

This Hamiltonian is time-independent and is uniquely defined for each laser detuning  $\delta$ .

In some simulations (see Chapters 4 and 7), we are interested in calculating the state evolution in a pulse sequence: ions are initialized to some state, then a set of laser pulses (gates) with various durations, phases, and detunings are applied. The ion states are expressed in the QC frame where they are stationary, while computations are carried out in the laser frame where the Hamiltonian is time-independent. Thus every new laser pulse requires switching between the QC frame and the laser frames. To do this switching, we need a unitary transformation which combines the two frame transformations:

$$e^{-iH_0 t} |\phi_{qc}\rangle = e^{-iH_L t} |\phi_{laser}\rangle \quad (2.31)$$

$$|\phi_{laser}\rangle = e^{i(H_L - H_0)t} |\phi_{qc}\rangle \quad (2.32)$$

$$= e^{i(\delta\sigma_z - \omega_z a^\dagger a)t} |\phi_{qc}\rangle. \quad (2.33)$$

## Stark shifts

In the case of on-resonant interaction where the laser is tuned near the atomic transition frequency ( $\Omega \gg \delta$ ), the effect of off-resonant interactions can be neglected. However, in the case where this condition does not hold, as for example when a laser is detuned to excite a sideband transition, off-resonant coupling to the carrier changes the resonance frequency of the sideband transition. In the 2-level system, the Hamiltonian in the laser frame is:

$$H_{laser} = -\delta\sigma_z + \frac{1}{2}(\Omega e^{i\phi}\sigma^+ + \Omega e^{-i\phi}\sigma^-) \quad (2.34)$$

$$= \begin{pmatrix} -\delta & \Omega e^{i\phi} \\ \Omega e^{-i\phi} & \delta \end{pmatrix}. \quad (2.35)$$

Solving for the eigenvalues of the Hamiltonian gives  $\pm \frac{1}{2}\sqrt{\delta^2 + \Omega^2}$ . The shifted energies of the states are:

$$\Delta E = \pm \frac{\delta}{2|\delta|} \left( \sqrt{\delta^2 + \Omega^2} \pm |\delta| \right). \quad (2.36)$$

In the Lamb-Dicke regime when the laser excites the sidebands, the approximation  $\Omega \ll \delta$  usually holds, and leads to the approximate Stark shift  $\Delta E = \pm \frac{\Omega^2}{2\delta}$ .

## 2.2.2 Atomic structure

$^{88}\text{Sr}^+$  is used for all experiments in this work. This species is a good candidate for experiments in quantum computing, due to its simple level structure and the availability of laser diodes for all relevant atomic transitions.  $^{88}\text{Sr}^+$  is a hydrogen-like atom with no nuclear spin, and its atomic structure is shown in Figure 2-2. The S-P transition is used for efficient doppler cooling and detection. The S-D quadrupole transition has a lifetime of  $\sim 350$  ms and is used to store quantum information.

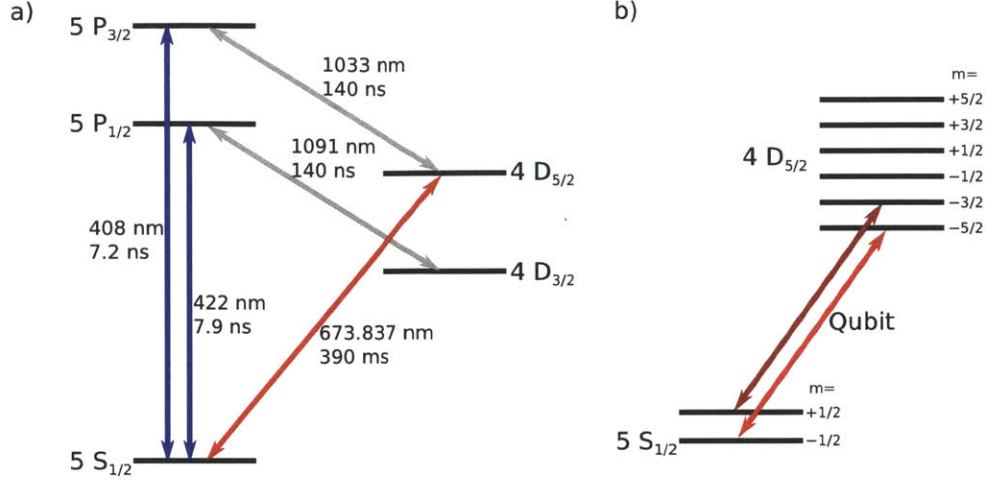


Figure 2-2: Atomic structure of the  $^{88}\text{Sr}^+$  ion. (a) Electric dipole transitions are shown in blue (Doppler cooling and detection) and gray (repumping). The quadrupole transition at 674 nm is used as the qubit transition. All wavelengths are in air. (b) Our chosen qubit transition is the between the lowest Zeeman substates of the S and D levels. The  $S_{1/2}(m=+1/2) \rightarrow D_{5/2}(m=-3/2)$  transition is pumped to empty the ion from the  $S_{1/2}(m=+1/2)$  substate.

In the presence of an external applied magnetic field, the atomic structure is modified by Zeeman splitting to avoid energy degeneracy. The energy shift of a state with Landé  $g$ -factor  $g_x$  is  $\Delta E = mg_x \mu_B B$ , where  $m$  is the angular momentum,  $\mu_B$  is the Bohr magneton  $\simeq 1.4 \text{ MHz/G}$ , and  $B$  is the applied magnetic field.

The various atomic transitions only couple Zeeman sublevels; dipole transitions such as S-P allow  $\Delta m = 0, \pm 1$ , while quadrupole transitions such as S-D allow  $\Delta m = 0, \pm 1, \pm 2$ . In addition, depending on the magnetic field direction, laser  $\vec{k}$  vector and polarization  $\hat{e}$ , only specific states are coupled. An example of a detailed derivation can be found in [Roo00]. In the case of  $\vec{k} \perp \hat{e} \perp \vec{B}$  as in the experiment in this thesis, the allowed transitions are  $\Delta m = \pm 2$ . The  $S_{1/2}(m = -1/2) \leftrightarrow D_{5/2}(m = -5/2)$  transition is chosen as the qubit transition, as indicated in Figure 2-2(b).

### 2.2.3 Motional state

A laser-cooled ion typically occupies the bottom of the trap potential with secular frequency  $\omega_z$ . The trap potential can thus be approximated as a harmonic potential with discretized energy levels corresponding to typically  $\sim 1 \text{ MHz} \approx 50 \text{ } \mu\text{K}$ . The Hamiltonian for the harmonic oscillator with frequency  $\omega_z$  is

$$H = \omega_z \left( a^\dagger a + \frac{1}{2} \right). \quad (2.37)$$

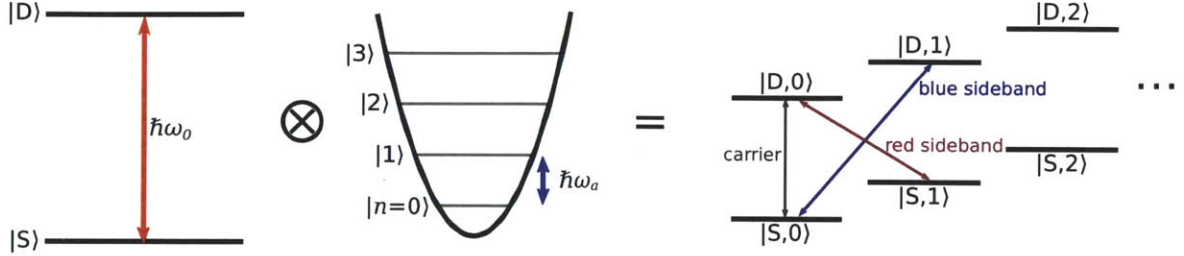


Figure 2-3: Tensor product of atomic and motional states describing the complete Hilbert space for quantum computation with a single ion.

The eigenstates of  $H$  is the set of states  $|n\rangle$  which satisfies

$$a |n\rangle = \sqrt{n} |n-1\rangle \quad (2.38)$$

$$H |n\rangle = \omega_z \left( n + \frac{1}{2} \right) |n\rangle. \quad (2.39)$$

The complete level structure of a single trapped ion, in the form most useful for quantum information, is described as a tensor product of the atomic (2-level) states and the motional states, as illustrated in Figure 2-3.

## 2.3 Quantum state manipulation

The  $S_{1/2}(m = -1/2)$  and  $D_{5/2}(m = -5/2)$  states form the atomic qubit. The motional qubits are used to mediate coupling between ions trapped in the same chain, and to define auxiliary levels needed to perform certain types of entangling gates such as the Cirac-Zoller CNOT gate. This section describes the physical realization of quantum operations including state initialization, measurement, and quantum gates, via the laser-atom interactions presented in the preceding sections.

### 2.3.1 State initialization

When an ion is initially trapped, its temperature is on the order of hundreds of quanta. Doppler cooling and sideband cooling are necessary to reduce the motional state to a well-defined state (usually the ground state,  $n = 0$ ) to proceed with quantum operations.

Doppler cooling on the  $^{88}\text{Sr}^+$  ion is performed on the 422 nm  $S_{1/2} \leftrightarrow P_{1/2}$  transition. The decay rate of this transition is about 20 times higher than the decay rate of the 1091 nm  $P_{1/2} \leftrightarrow D_{3/2}$  transition; the  $D_{3/2}$  state needs to be continuously repumped to maintain cooling. Dark states, or coherent superpositions of the S and D states, are avoided by appropriately detuning both the 1091 nm and 422 nm lasers. Doppler cooling is insufficient to reach the motional ground state because the scattered photons induce a velocity (photon recoil). The Doppler cooling limit is typically  $\sim 10$  quanta for the  $\text{Sr}^+$  ion with an axial trap frequency of 1 MHz. In experiment, the actual cooling limit may be higher due to power broadening or improperly set detuning.



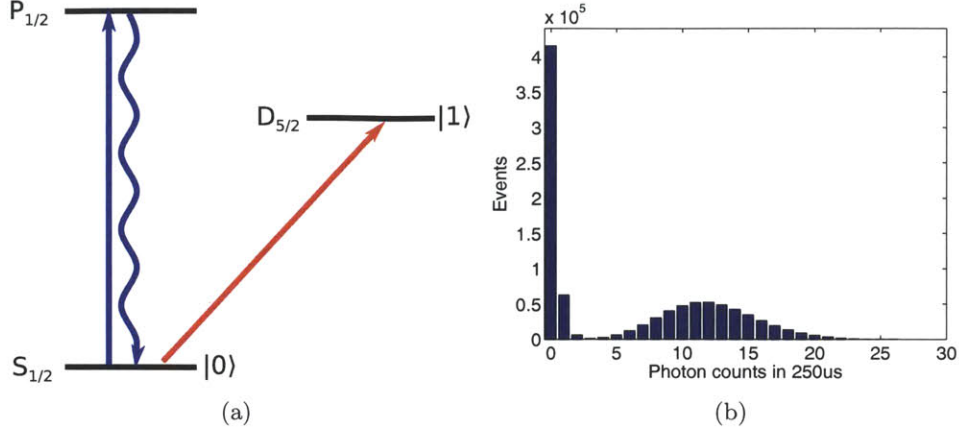


Figure 2-4: (a) Atomic levels for state detection using the electron shelving technique. Scattered light on the S-P transition distinguishes the S and D states. (b) Distribution of count rates in the case of 1 ions; the two peaks corresponding to the ion being in the S (bright) or D (dark) states.

Sideband cooling on the S-D transition is performed with the qubit laser at 674 nm to reach the motional ground state. The ion is pumped from the  $|S_{1/2}, n\rangle$  state to the  $|D_{5/2}, n-1\rangle$  state by driving a  $\pi$  pulse on the red sideband transition. The ion preferentially decays to the  $|S_{1/2}, n-1\rangle$  state via the carrier transition, thus removing one phonon in the process. In practice since the S-D transition has a long lifetime, the decay is implemented by incoherently driving the 1033 nm  $D_{5/2} \leftrightarrow P_{3/2}$  transition, where the spontaneous decay rate from the  $P_{3/2}$  transition is much faster. The intensity and pulse duration of the 1033 nm pulse is adjusted to optimize the cooling rate. The  $|S_{1/2}, n=0\rangle$  ground state does not couple to the red-sideband pulse and thus remains in the state even with further  $\pi$  pulses. The cycle consisting of a  $\pi$  pulse on the red sideband followed by a pulse on the 1033 nm transition is typically repeated 90-150 times to reach the motional ground state with 90% probability. During sideband cooling, population on other Zeeman substates which are outside of the desired Hilbert space, such as the  $S_{1/2}(m=+1/2)$  state, is also pumped out by applying multiple cycles of  $\pi$  pulses on the relevant transition.

### 2.3.2 State detection

The electron shelving technique is used to distinguish the S and D states with high efficiency. The simplified level scheme in Figure 2-4 illustrates the method. The S state is coupled to the P state with a strong transition, which is continuously driven with the Doppler cooling laser at 422 nm. If the ion is in the S state, photons are scattered and the ion fluorescence can be detected. If the ion is in the D state, no fluorescence is detected.

The state discrimination is a function of the detection time and background signal. Photon count distribution follows two Poisson distributions: one is centered around the ion being in the S state, and the other centered around the ion being in the D state but likely not 0 due to stray light or detector dark counts. Figure 2-4 shows an example of the histogram of photon counts. Here the detection time is set to be 250  $\mu$ s, so that for an

average ion fluorescence of 48000 photons/s, 12 photons are detected during this time.

The motional state of the ion cannot be measured directly, and must be determined in conjunction with the atomic state by the use of the sideband transitions. Measuring the motional state roughly corresponds to determining the temperature of the ion. Several methods exist to determine the motional state.

In the Lamb-Dicke regime, the Rabi frequency on the blue sideband is dependent on  $n$ :  $\Omega_{n+1,n} = \Omega_{n,n}\eta\sqrt{n+1}$ . If a pulse on the blue/red sideband ( $\delta = \pm\omega_z$ ) with a fixed pulse length  $t$  is applied to an ion in the  $|S, n\rangle$  state, then the transition probability is the sum of the transition probabilities on each sideband with  $n = 0, 1, 2, \dots$  and can be expressed as:

$$P_{blue} = \sum_{n=0}^{\infty} P_n \sin(\Omega_{n+1,n}t/2)^2 \quad (2.40)$$

$$P_{red} = \sum_{n=0}^{\infty} P_n \sin(\Omega_{n-1,n}t/2)^2. \quad (2.41)$$

Thus by observing Rabi oscillations on the blue sideband (fix detuning while measuring the transition probability as a function of  $t$ ), the resulting curve can be decomposed into components with frequencies  $\Omega_{n,n}\eta\sqrt{n+1}$ . The occupation of the components corresponds to the population distribution of the motional states.

Alternatively, if the ion is assumed to be in the thermal state, a simpler measurement can be made to determine the average number of quanta  $\langle n \rangle$ . The thermal state occupation probability is given by  $P_n = \frac{\langle n \rangle^n}{(\langle n \rangle + 1)^{n+1}}$ . Substituting to the above equation yields

$$\begin{aligned} P_{blue} &= \sum_{n=0}^{\infty} P_{n+1} \sin(\Omega_{n,n+1}t/2)^2 \\ &= \frac{\langle n \rangle}{1 + \langle n \rangle} \sum_{n=0}^{\infty} P_n \sin(\Omega_{n,n+1}t/2)^2 \\ &= \frac{\langle n \rangle}{1 + \langle n \rangle} P_{red}. \end{aligned} \quad (2.42)$$

Thus  $\langle n \rangle$  can be determined by measuring  $P_{blue}$  and  $P_{red}$ , the ratio of which is independent of  $\Omega$  and  $t$ .

### 2.3.3 Quantum operations

A set of quantum operations consisting of single-qubit gates and a two-qubit entangling gate is universal: any arbitrary quantum operation can be constructed from combinations of those gates [DiV95]. Here we list a particular set of single-qubit gates, and two example of two-qubit gates, and show how they are mapped to appropriate laser pulses on trapped ions.

### Single-qubit, carrier rotations

A single-qubit rotation,  $R(\theta, \phi)$ , describes a rotation around the Bloch sphere with rotation angle  $\theta$  and phase  $\phi$ . In particular, it is sufficient to specify the angle to be one of  $0, \pi/2, \pi$ , or  $3\pi/2$ , where  $R(\theta, 0) = R_x(\theta)$ ,  $R(\theta, \pi/2) = R_y(\theta)$ , etc. The operations are:

$$R_x(\theta) = \exp\left(i\frac{\theta}{2}\sigma^x\right) \quad (2.43)$$

$$R_y(\theta) = \exp\left(i\frac{\theta}{2}\sigma^y\right). \quad (2.44)$$

In ions, these rotations can be implemented by a carrier pulse. The Hamiltonian describing a carrier pulse is

$$H_{qc} = \frac{1}{2}\Omega \left( e^{i\phi}\sigma^+ + e^{-i\phi}\sigma^- \right). \quad (2.45)$$

The unitary matrix is

$$\exp(-iH_{qc}t) = \exp\left(-i\frac{\Omega t}{2}\sigma^x\right) \quad (2.46)$$

$$= R_x(\Omega t). \quad (2.47)$$

The  $Z$  rotation  $R_z(\theta)$  can be implemented in several ways. If only one ion is present and all operations are single-qubit in a given gate sequence, a single  $R_z(\theta)$  operation is equivalent to a phase shift of all subsequent gates by  $\Delta\phi = \theta$ . It can also be explicitly implemented by the identity

$$R_z(\theta) = R_x(-\pi/2)R_y(\theta)R_x(\pi/2). \quad (2.48)$$

Finally, such a phase rotation can be realized by far detuning the laser from the carrier transition and using the AC Stark shift effect to impart a phase shift with minimal effect on the population.

### Single-qubit, sideband rotations

Certain gate schemes requires transferring quantum information between ions in a chain, mediated by the shared motional state. For example, the Cirac-Zoller Controlled NOT gate between two ions need to map one of the ion's atomic state to the shared motional state. This mapping is accomplished with the blue sideband pulse.

The Hamiltonian for a blue-sideband pulse on resonance is

$$H_{qc} = \frac{1}{2}\Omega_{n+1,n} \left( e^{i\phi}\sigma^+ a^\dagger + e^{-i\phi}\sigma^- a \right). \quad (2.49)$$

The resulting unitary operator acts similarly to the  $R(\theta)$  operation on the carrier transition, except that the two coupled states are  $|S, n\rangle$  and  $|D, n+1\rangle$ .

## Two-qubit CNOT gate

The controlled-NOT gate between two ions is an example of a two-qubit entangling operation that can form a universal gate set together with the above single-qubit operations. The action of the gate is to apply a bit flip to a target qubit depending on the state of a control qubit. For example, in a Hilbert space described by the basis set  $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ , the unitary operator is

$$U_{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \quad (2.50)$$

The originally proposed Cirac-Zoller CNOT gate requires an auxiliary state to perform a conditional rotation; the  $|S, n = 1\rangle$  state and the red sideband pulse, which does not couple to the  $|S, n = 0\rangle$  ground state, accomplished this. This gate was first implemented in a single Beryllium ion [MMK<sup>+</sup>95]. A subsequent implementation used an alternate, composite pulse sequence that avoids the need for an auxiliary level [SKHR<sup>+</sup>03]:

$$U_{CNOT} = Z(\pi)R\left(\frac{\pi}{2}, -\frac{\pi}{2}\right)R^+\left(\frac{\pi}{\sqrt{2}}, \frac{\pi}{2}\right)R^+\left(\frac{\pi}{2}, 0\right)R^+\left(\frac{\pi}{\sqrt{2}}, \frac{\pi}{2}\right)R^+\left(\frac{\pi}{2}, 0\right)R\left(\frac{\pi}{2}, -\frac{\pi}{2}\right) \quad (2.51)$$

where  $R$  is the carrier rotation and  $R^+$  is the sideband rotation. The Cirac-Zoller CNOT gate has the disadvantage that because the pulse duration of the sideband depends on the phonon occupation number  $n$ , the gate requires the ion to initially be in the motional ground state, making it potentially more sensitive to motional state heating. In Chapter 4 we implement this gate on the atomic and motional states of a single  $\text{Sr}^+$  ion where motional state heating was not found to be a limiting factor on gate fidelity.

## Mølmer-Sørensen gate

The Mølmer-Sørensen (MS) gate is an entangling operation on an ion chain with arbitrary number of ions. Starting with all ions in the ground state, the operation creates Greenberger-Horne-Zeilinger (GHZ) states  $\frac{|00..0\rangle + |11..1\rangle}{\sqrt{2}}$ . The ion chain is globally illuminated with a bichromatic light field, with detunings  $\delta = \pm\omega_z \mp \Delta$ , where  $\Delta$  is a small detuning away from the sidebands. The detuning means that intermediate states of a two-photon transition are only virtually populated. Transition paths between the states  $|00..0\rangle$  and  $|11..1\rangle$  interfere in such a way that the effective Rabi frequency of the two-photon transition is independent of the motional state (see Figure 2-5). More details of the MS gate can be found in Refs [MS99, SM99].

The Hamiltonian is:

$$H_{MS} = \tilde{\Omega} \left( a e^{i\Delta t} + a^\dagger e^{-i\Delta t} \right) S_\phi \quad (2.52)$$



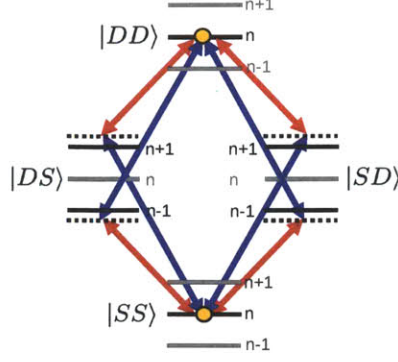


Figure 2-5: Schematic of the Mølmer-Sørensen gate. A bichromatic laser field is applied on two ions, with a small detuning from the blue and red sidebands. The paths interfere such that the effective Rabi frequency of the two-photon transition is independent of  $n$ .

where  $S_\phi$  is the collective rotation on all ions  $k = 1..N$ :

$$S_\phi = \sum_{k=1}^N \sigma_\phi^k \quad (2.53)$$

$$\sigma_\phi = \sigma_x \cos(\phi) + \sigma_y \sin(\phi) \quad (2.54)$$

and  $\tilde{\Omega}$  is the effective Rabi frequency of the two-photon transition:

$$\tilde{\Omega} = \frac{(\Omega\eta)^2}{2\Delta}. \quad (2.55)$$

At a gate time  $\tau = 1/\Delta$ , the unitary operation is

$$U_{MS}(\theta, \phi) = \exp(-i\theta S_\phi^2). \quad (2.56)$$

The Mølmer-Sørensen gate has the advantage of being insensitive to the phonon state. It has been demonstrated on thermal states cooled only with Doppler cooling [KBZ<sup>+</sup>09]. The bichromatic laser pulse is naturally insensitive to common-mode frequency fluctuations and naturally has no Stark shift as would be associated with a monochromatic sideband pulse. With the addition of amplitude shaping, the MS gate has been demonstrated with better than 99% fidelity [BKRB08].



## Chapter 3

# Experimental Apparatus

This chapter describes the experimental system for microfabricated ion traps used to perform most of the experiments described in this thesis. The cryostat, lasers, and control systems have been described extensively in Refs [Lab08, ASA<sup>+</sup>09, LRB<sup>+</sup>07], and the microfabricated traps have been described in [Lab08]. Only major features and changes will be described here.

### 3.1 Cryostat

Cryogenic cooling of microfabricated ion traps confer a number of advantages. The motional heating rate of trapped ions, originating from electric field noise, has been shown to be reduced by orders of magnitude at cryogenic temperatures [DOS<sup>+</sup>06, LGA<sup>+</sup>08]. A practical advantage is that ultra-high vacuum conditions, a necessary environment for trapping ions with long lifetimes by minimizing collision with background gases, can be achieved within 24 hours of installing a new trap, compared to usually 1-2 weeks for a conventional room-temperature vacuum system. Finally, a cryogenic environment offers the probability of integrating with other low-temperature devices such as superconducting resonators or photon detectors.

A variety of cryostats are commercially available, including bath, closed-cycle, flow, and others. The bath cryostat is the simplest one and consists of insulated tanks holding cryogens. Its advantages include vibration-free operation and minimal maintenance; the main disadvantage is the continual consumption of cryogens. The liquid helium bath cryostat (TK1813, QMC Instruments) used in this work is shown in Figure 3-1. It consists of a 1.4 L liquid helium tank, on which the trap mount and DC filtering electronics are mounted. Above it is a 1.75 L liquid nitrogen tank, on which is attached a metal radiation shield that surrounds the 4K region to protect it from blackbody radiation.

The cryostat is sealed with O-rings and pumped with a turbomolecular pump (20 L/s) to a pressure of  $\sim 1 \times 10^{-6}$  torr, as measured with an ion gauge close to the pump, before cooling. After cooling to liquid helium, the vacuum inside the trap region is expected to be below  $1 \times 10^{-10}$  torr based on achieved ion lifetimes of hours. Two activated charcoal getter pumps, one each heat sunk to the liquid helium and liquid nitrogen tanks, provide cryopumping of residual gas.

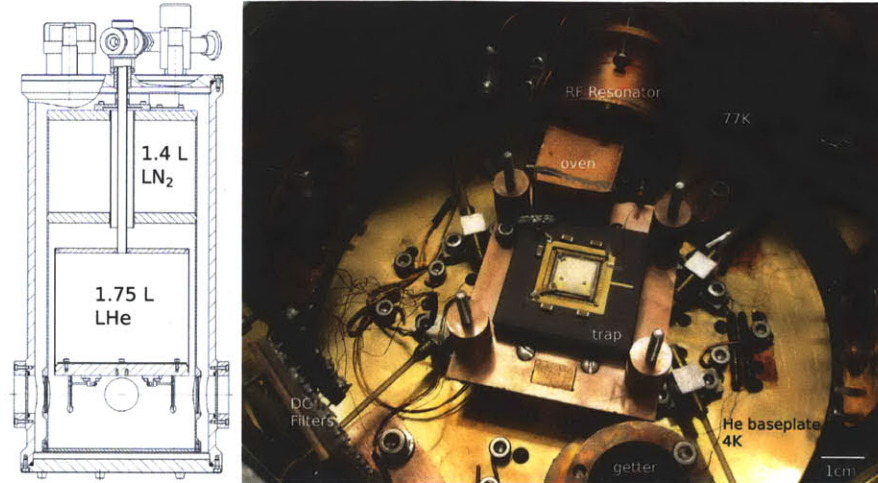


Figure 3-1: (a) Cross-sectional schematic of the cryostat showing the liquid helium and liquid nitrogen tanks. The 77K shield attached to the liquid nitrogen tank encloses the 4K tank and work area. (b) Photo of the internals of the cryostat showing an installed trap.

Optical access to the cryostat is provided by three openings for windows (BK7 on the 77K shield and Melles-Griot 02 WBK 226 for the room-temperature shield) around the sides of the cryostat for lasers, and one on the bottom for imaging. A 55-pin feedthrough provides all electrical connections. Wiring for the DC connections are provided by 36 AWG phosphor bronze wires, chosen for their low thermal conductivity. Two of the wires are joined to NbTi superconducting wires which are thermally sunk to the cryostat base before reaching the trap, for applications requiring high current.

A helical resonator mounted on the 77 K shield is used to step up the RF voltage to the trap. The resonator is made of 22 turns of 1 mm copper wire, and provides a step-up of 20. With a load of 3 pF typical for a trap, the resonance frequency is typically between 33-40 MHz. The step-up and resonance frequency do not change significantly between room temperature and 77 K. The helical resonator is driven with a frequency generator (Agilent 33250A) and a 4W power amplifier (Mini-Circuits T1A-1000-1R8). A bi-directional coupler is used to measure the reflected power from the resonator and trap in order to determine the resonant frequency, which varies slightly between traps due to variations in trap capacitance. Figure 3-2 illustrates the RF circuitry used to supply the trap.

The cryostat is rigidly mounted on a 80/20 support structure on an optical table, to isolate it from vibrations coupling via the floor. A square hole in the optical table allows the camera, photomultiplier tube (PMT), and imaging optics to be mounted below the table. Figure 3-3 shows a photograph of the apparatus including the cryostat, optical table, and vacuum pumps.

### 3.2 Laser system

Quantum operations on the  $^{88}\text{Sr}^+$  ion requires 4 lasers: 422 nm for Doppler cooling and detection, 674 nm for manipulating the qubit, and 1091 nm and 1033 nm for repumping the

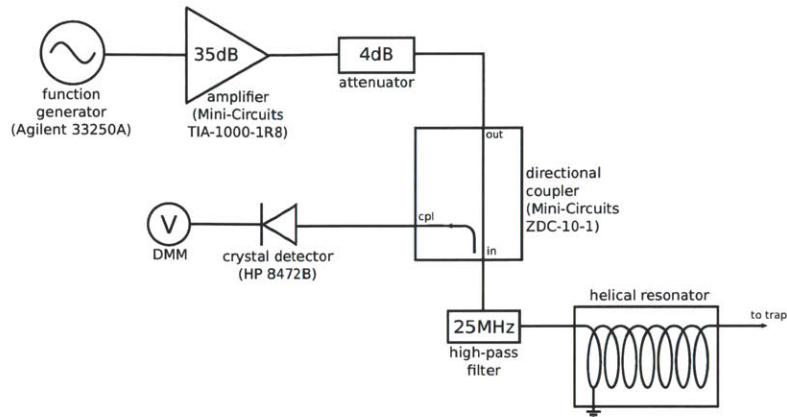


Figure 3-2: Block diagram of RF chain for the trap RF electrode. Input signal of  $\sim 100$ - $400$  mVpp from a function generator is amplified and passes through a directional coupler before entering the cryostat. The directional coupler, combined with the crystal detector, is used to monitor the reflected power from the trap. A helical resonator inside the cryostat provides a step-up of 20 to result in RF voltages of typically 100-300 V on the trap.

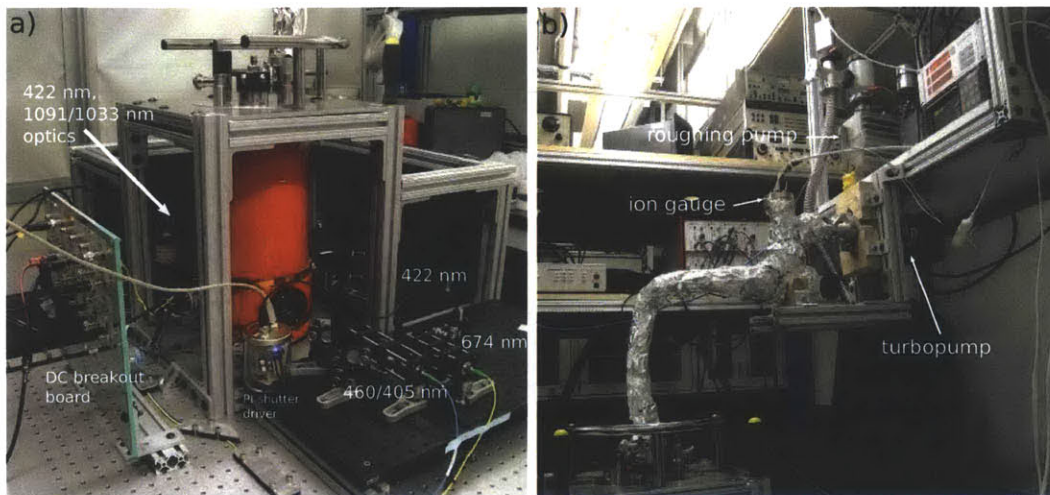


Figure 3-3: Cryogenic ion trap experimental setup. (a) The cryostat is mounted on an optical table which has a hole through the center for the imaging system mounted below the cryostat. Three breadboards next to the cryostat holds the laser delivery optics. (b) The cryostat is pumped by a small turbopump and a roughing pump mounted from the ceiling.



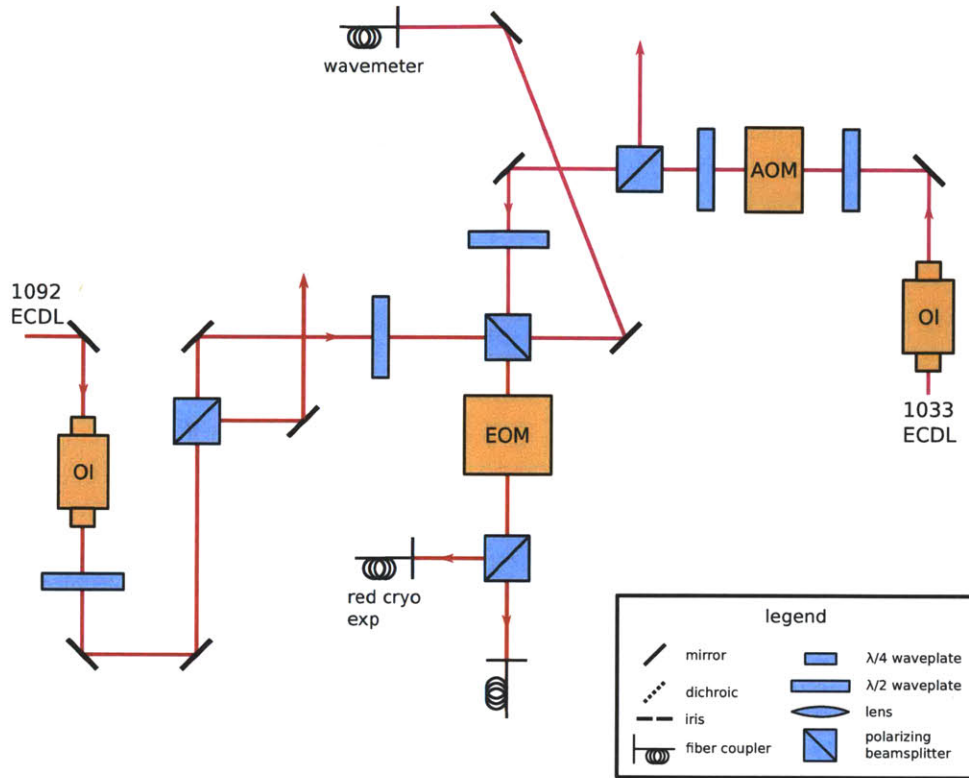


Figure 3-4: Beam distribution optics for the 1091 and 1033 nm lasers. (OI = optical isolator, AOM = acousto-optical modulator). Arrows lead to beam paths for other experiments. The legend for components apply to the remaining optical diagrams in this chapter.

D states. These lasers are provided by external cavity diode lasers (ECDL) with passive optical feedback and an additional filter cavity. The systems consist of a conventional ECDL, coupled to a 20 cm filter cavity, all mounted on a monolithic baseplate and pumped to vacuum using a turbopump. These laser systems reduce the linewidth from typically 1 MHz for bare ECDLs to  $\sim 20$  kHz with optical feedback and filtering. Their outputs are sent through acousto-optical modulators (AOMs) which performs switching and frequency tuning at the MHz level, before being distributed to multiple experiments. Figures 3-4 and 3-5 illustrates the beam-shaping and distribution optics for the 422, 1092, and 1033 lasers. More details of the laser systems can be found in Refs [Lab08, LRB<sup>+</sup>07].

### Qubit laser

The 674 nm qubit laser requires smaller linewidth to address the narrow ( $\sim 1$  Hz) qubit transition. This is achieved by locking the 674 nm laser output to a high-finesse, stable cavity. The cavity consists of two high-reflectivity mirrors optically contacted to a tube made of ultra-low expansion (ULE) glass. The cavity is mounted inside a thin copper can to provide a uniform temperature environment, and further contained inside a vacuum chamber which is pumped down to  $10^{-7}$  torr for vibration and temperature isolation. A temperature sensor circuit based on a quartz crystal oscillator is attached to the surface of the copper can; the copper can is wrapped with a heater wire. The temperature of

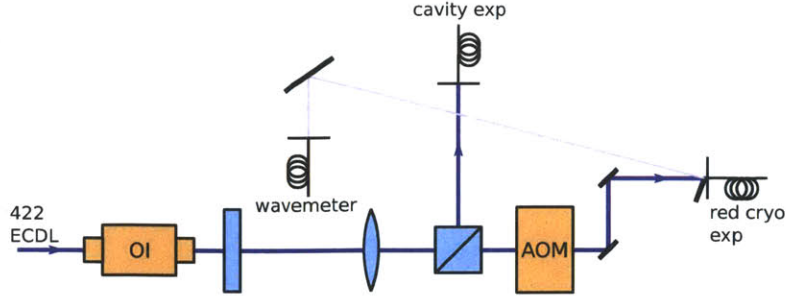


Figure 3-5: Beam distribution optics for the 422 nm laser.

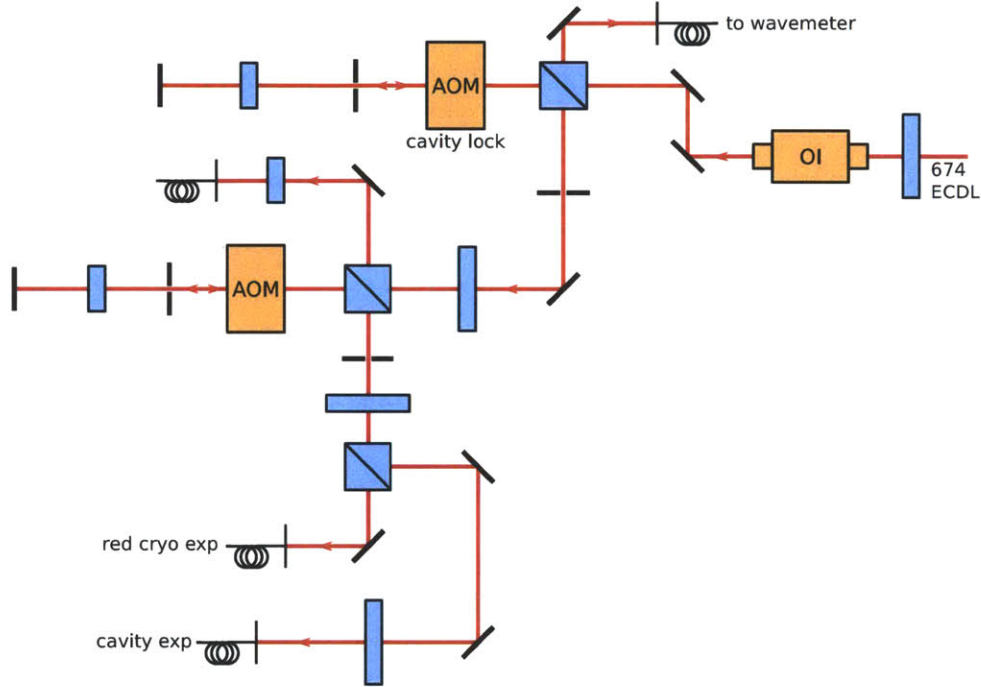


Figure 3-6: Beam shaping, cavity lock, and distribution optics for the 674 nm qubit laser. The two AOMs are in a double-pass configuration.

the cavity is stabilized to  $30^{\circ}\text{C}$  via an FPGA-based PID loop to better than 1 mK. The resulting linewidth of the qubit laser is on the order of 100 Hz, sufficient to achieve single-qubit coherence times of up to  $800\ \mu\text{s}$ . More details of the ULE cavity and associated control systems can be found in Refs [Lab08, She08].

To bridge the frequency difference between the ECDL locked to the fixed ULE cavity and the ion's qubit frequency, a double-pass 200 MHz AOM is used. The optics for locking the 674 diode laser to the cavity and distributed to the experiments is shown in Figure 3-6. The frequency applied to the AOM to address the S-D carrier transition thus gives an indication of the changes in the cavity frequency over time. Figure 3-7 shows an initial sudden change in the cavity frequency due to temperature cycling of the cavity to  $32.7^{\circ}\text{C}$ , and subsequent relaxation to near its original value over several years.

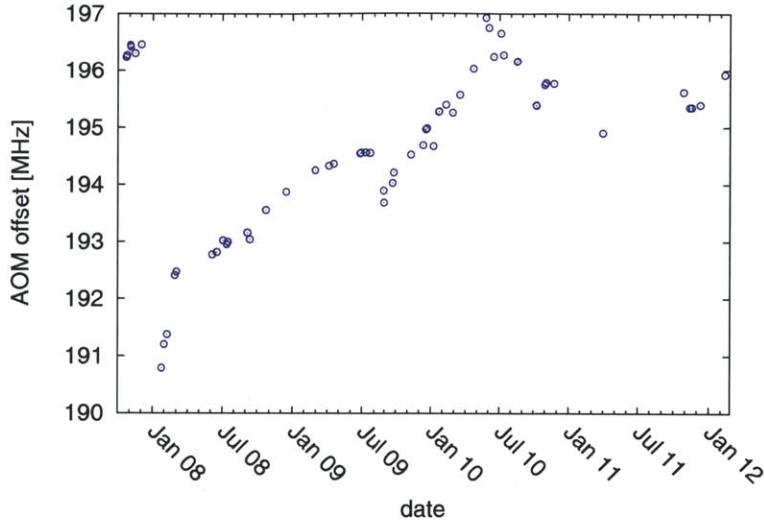


Figure 3-7: Drift of the 674 nm cavity frequency over several years. Y-axis is the frequency applied to the experiment AOM to address a single ion’s carrier qubit transition.

### Photoionization lasers

Two additional lasers, at 460 nm and 405 nm, are necessary to ionize neutral Sr atoms from a thermal vapor to load the trap. Photoionization is a two-photon process; the state  $S_0$  is coupled to the  $P_1^0$  state via an electric dipole transition at 460 nm with a linewidth of  $\sim 30$  MHz. The  $P_1^0$  state is coupled to an auto-ionizing  $D_2$  state using a laser at 405 nm [MBK95]. This transition is broad, with a linewidth of 1 nm [Gup04].

The 405 nm laser is supplied by an ECDL. No further stabilization is necessary. The power at the ion is set to be  $140 \mu\text{W}$ . More power should lead to faster ionization rate and higher loading rate, but at the expense of increased risk of causing laser-induced charging on the trap due to the short wavelength (see Chapter 6.)

The 460 nm laser is not commercially available as a diode, and is generated by a frequency-doubled diode at 921 nm. The 921 nm diode is rated for 300 mW and is configured in a standard ECDL package (Toptica DL100). The output of the diode is mode matched to a bow-tie cavity (Toptica SHG110), where a temperature-controlled PPKTP crystal (Raicol, 15 mm, uncoated) is used for the frequency conversion. The temperature of the crystal is stabilized to 0.1 K via PID control to the phasematching temperature at 46.8 C. The output power at 461 nm is  $\sim 2\text{-}3$  mW, shared among several experiments. Typical power at the ion is set to be  $90 \mu\text{W}$ . The beam shaping and distribution optics are shown in Figure 3-8.

### Laser delivery

The six lasers for cooling and manipulating the ion are transported to the optical table of the cryostat via four single-mode optical fibers. The 461 and 405 nm lasers are combined into a single fiber before the experiment table, and similarly for the 1033 and 1092 nm lasers. Figure 3-9 shows the optics used to deliver light from the fibers to the trap. Table



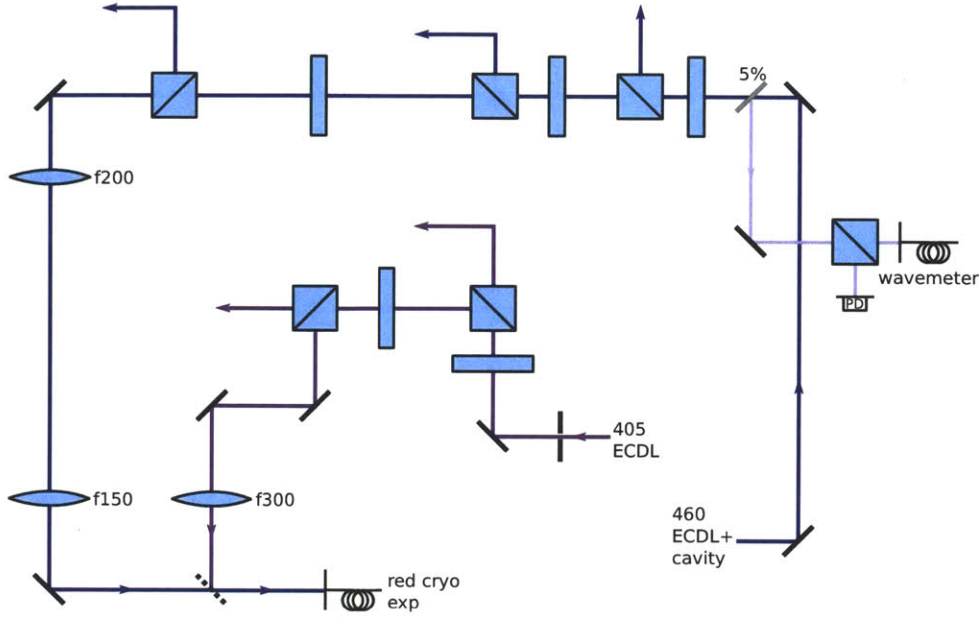


Figure 3-8: Beam shaping and distribution optics for the 405 nm and 460 nm photoionization lasers. Arrows lead to beam paths for other experiments.

Wavelength (nm)	beam waist ( $\mu\text{m}$ )	viewport transmission (%)	power at ion ( $\mu\text{W}$ )
405	100	86	140
460	100	85	90
422	40	90	7
674	35	91	600
1033	200	78	7
1092	200	77	13

Table 3.1: Typical beam waists, cryostat viewport transmission, and power at the ion for all lasers used in the experiment. Transmission refers to the total transmission through outer and inner cryostat viewports.

3.1 shows the typical powers, beam waist, and optical transmission of the cryostat windows for each of the wavelengths used.

### 3.3 Surface-electrode trap

The surface-electrode trap used in this work is a single-zone trap with a five-electrode design, arranged as DC-RF-GND-RF-DC along the  $x$  axis. The layout of the electrodes is shown in Figure 3-10. The RF electrode width is asymmetric in the trapping regions in order to tilt the principle axes of the trapping potential. This allows the cooling lasers, confined to the plane parallel to the trap, to have projections along all principle axes and cool all degrees of motion of the ion. In addition, the notch in the ground electrode compensates for the asymmetry of the GND/RF electrode along the trap axis, such that there exists a well-defined RF null.

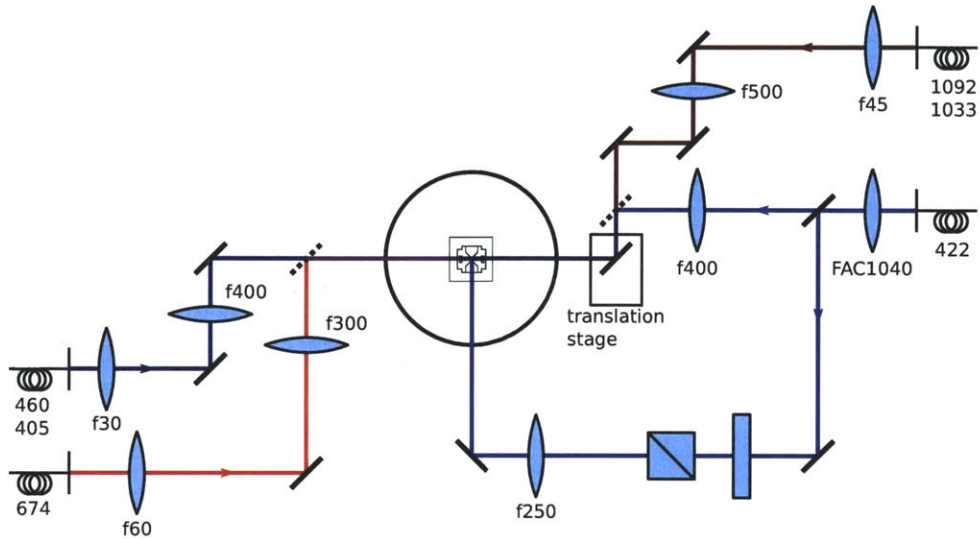


Figure 3-9: Laser delivery optics to the cryostat. For each fiber path, the first lens collimates the beam and the second lens focuses the beam onto the ion.

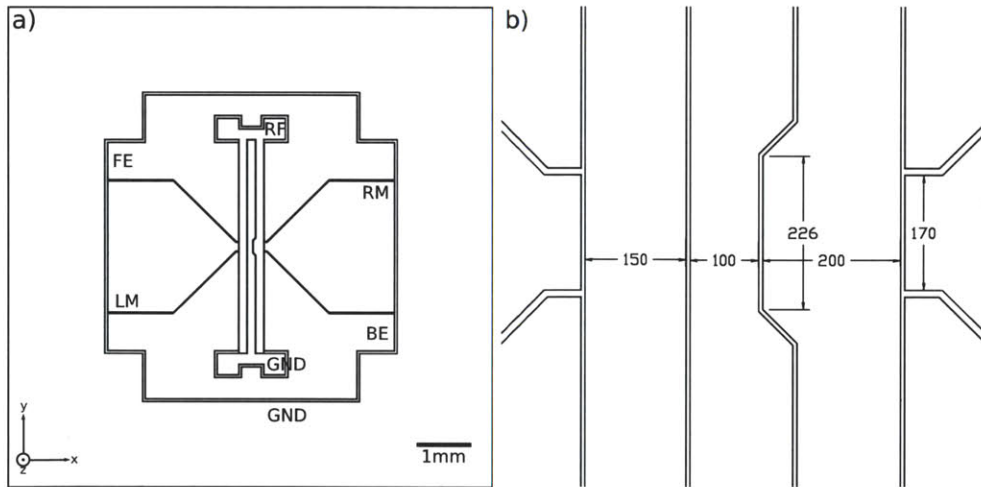


Figure 3-10: (a) Schematic of the trap electrodes. LM, RM, FE, BE are names given to the four DC compensation electrodes. (b) Key dimensions of the central trapping region for a  $100 \mu\text{m}$  trap, in  $\mu\text{m}$ .

Trap size ( $\mu\text{m}$ )	$V_{BE,FE}$ (V)	RF freq ( $2\pi$ MHz)	RF amp (V)	axial $\omega_z$ ( $2\pi$ MHz)	radial $\omega_z$ ( $2\pi$ MHz)	measured axial $\omega_z$ ( $2\pi$ MHz)
75	5	35.7	120	0.95	2.1, 2.8	0.90
100	5	33.4	120	0.6	1.5, 2.0	0.58
	10	33.4	160	0.9	1.8, 2.7	0.8
	25	33.4	240	1.45	2.2, 4.0	1.33
150	10	38.0	340	0.67	1.3, 1.9	0.60
	15	38.0	400	0.84	1.4, 2.2	0.76

Table 3.2: Typical operating parameters for traps of size 150, 100, 75  $\mu\text{m}$ , as predicted by CPO. RF frequencies are those actually used in at least one instance of a trap. RF amplitude is chosen such that the trap depth is 20-30 meV. The ion height roughly corresponds to the trap size (width of the narrow part of the ground electrode).

We fabricated and tested traps of this design with the narrowest part of the GND electrode having widths of 150, 100, and 75  $\mu\text{m}$ . Trap parameters including ion height, DC compensation voltages, secular frequencies, and trap depths are found via numerical simulation using Charged Particle Optics, a software package from Electron Optics [CPO], and are typically within 20% of the experimentally measured values. The theoretical compensated DC electrode voltages satisfy  $V_{BE} = V_{FE} = -V_{LM} = -\frac{5}{4}V_{RM}$ . Table 3.2 lists the trap sizes, each with typical DC operating voltages, RF source amplitude, and measured axial secular frequencies. Due to the helical resonator being located inside the cryostat and the design of the resonator where no connection was made for a monitoring output, the RF voltage on the trap cannot be directly measured, but can only be indirectly inferred from secular frequency measurements and trap modeling. The table lists the typical amplitude on the signal generator used to drive the helical resonator.

A list of all traps tested in this thesis can be found in Appendix A.

### 3.4 Experiment control

A variety of control electronics and software systems are necessary to run the experiment apparatus and perform data acquisition. These consists of a combination of commercial electronics (power supplies, function generators etc) and FPGAs developed in-house, communicating with a master control computer via GPIB and network interfaces.

The four trap electrodes are driven by power supplies (Keithley 2400, Agilent E3631A) and are controlled via GPIB. The Keithley power supplies are able to drive up to 200V output. The Agilent supplies are connected in series such that their  $\pm 25\text{V}$  outputs drive the differential voltage on a pair of electrodes. All power supply outputs are floating, so that the electric ground is referenced to the cryostat baseplate. Power supplies in principle give stable voltages, but expansion to a multi-zone trap with more than a few electrodes would likely require multi-channel DAC systems as a more flexible and economical solution.

Ion signal from the PMT is counted via an FPGA-based photon counting module. The FPGA clock runs at 62.5 MHz and allows for high resolution counting (16 ns). Most of the experiments in this thesis rely on precise control over the amplitude, frequency, and phase

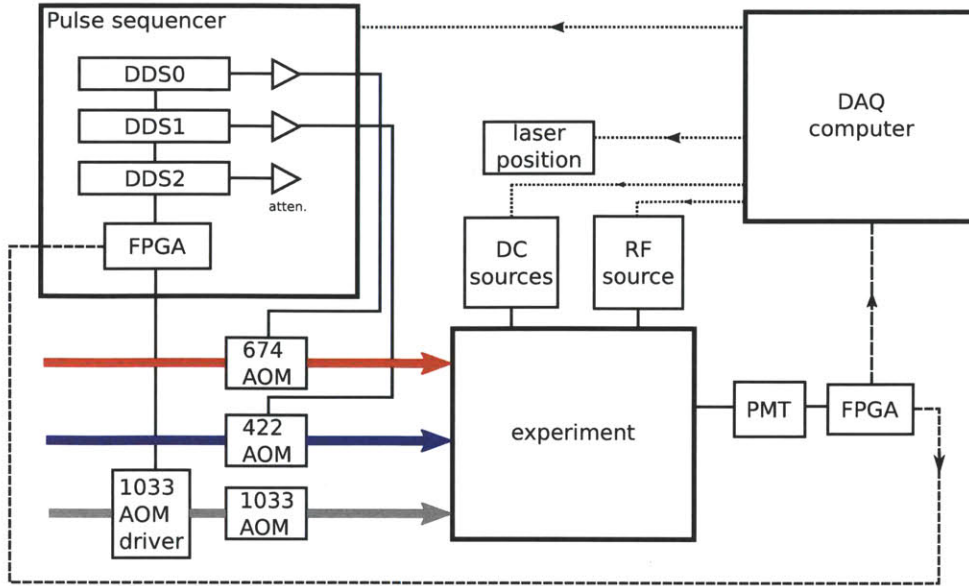


Figure 3-11: Block diagram for the experiment and control systems.

of the lasers pulses, and this is accomplished by the pulse sequencer. The pulse sequencer consists of 3 direct-digital synthesizers (DDSs, Analog Devices AD 9858) and an FPGA, used to switch laser pulses via AOMs (IntraAction and Neos). The FPGA for the pulse sequencer implements a von Neumann state machine, with a program counter, 32-bit  $\times$  4000 memory, two global memory registers, and an instruction set containing operations for the DDS as well as basic logic, timing, and flow control operators. Much more detail of the pulse sequencer can be found in the appendix of [Lab08]. The instruction set and basic pulse sequences used in this thesis are in Appendix B.

Figure 3-11 shows a block diagram with key elements of the entire experiment control system.

**Part I**

**Gates**



## Chapter 4

# Controlled-NOT gate in a surface trap

In recent years there has been increasing interest in microfabricated surface-electrode traps, owing to their inherent scalability [SCR<sup>+</sup>06, SHO<sup>+</sup>06]. But until recently, quantum gates have not been demonstrated in such traps. An issue with miniaturization of traps is that anomalous heating of the ion's motional state scales unfavorably with trap size [ESL<sup>+</sup>07], potentially limiting gate fidelity in traps of dimensions suitable for scalability [Ste07]. However, it has been shown that by cooling to cryogenic temperatures, the heating rate can be reduced by several orders of magnitude from room-temperature values [LGA<sup>+</sup>08], thus providing one potential solution to this problem. In this chapter, we demonstrate a quantum gate in a microfabricated surface-electrode ion trap that is operated in a cryogenic environment, and present some control techniques developed for this experiment.

We implement a Cirac-Zoller controlled-NOT (CNOT) gate using qubits represented by the atomic and motional states of a single ion. The  $S \leftrightarrow D$  optical transition in  $^{88}\text{Sr}^+$  is used as one of the qubits. The motional ground state and first excited state of the ion in the harmonic trap potential form the second qubit. The optical transition has the advantage of a long lifetime while requiring only a single laser (unlike hyperfine qubits), but the qubit is first-order Zeeman sensitive, which makes it susceptible to magnetic field noise. Taking advantage of the cryogenic environment, we stabilize the magnetic field using a pair of superconducting rings [GTC<sup>+</sup>91]. For readout, the qubit encoded in the motional state of the ion normally cannot be measured directly, but conditional pulse sequences allow full state tomography of the qubit system. We perform full process tomography on the atomic and motional qubits of a single ion to evaluate the performance of the gate.

The control techniques developed here may be applicable to the use of a single ion to probe and manipulate other systems, even though they focus on a single ion and do not necessarily imply scalability. Some such systems include the coupling of ions to superconducting qubits [TRBZ04], micromechanical cantilevers [HUG<sup>+</sup>05], cavities [KK09], and wires [DLC<sup>+</sup>09]. In many of these experiments, maximizing the coupling requires proximity of the ion to a surface, and coherence of the motional state is also desired.

This chapter is organized as follows. Section 4.1 briefly introduces the gate with references to prior work. Section 4.2 discusses motional state decoherence (heating rate) and

magnetic field stabilization, two elements of a quantum gate experiment which are unique to our surface-electrode trap implementation. Section 4.3 describes the state preparation and measurement sequences that allowed us to implement quantum process tomography on the single-ion system. Section 4.4 describes the realization of the CNOT gate, along with gate performance and error sources. Section 4.5 concludes with a perspective on applications, disadvantages, and alternatives to realizing gates on surface-electrode traps. Sections 4.2-4.4 have been published in [WLG<sup>+</sup>10].

## 4.1 CNOT gate

The controlled-NOT gate is a type of entangling gate in which the state of a target qubit depends on the state of a control qubit. Specifically, the state of the target qubit remains constant if the control qubit is in state  $|0\rangle$ , and the target qubit is flipped if the control qubit is in state  $|1\rangle$ . We use the atomic state of a single ion as the control qubit and the motional state as the target qubit. Successful implementation of this gate is conditional on the phase and amplitude coherence of the atomic and motional qubits during the gate operation.

This CNOT gate, as implemented between the atomic and motional states of a single ion, was first demonstrated in Ref. [MMK<sup>+</sup>95]. Based on population measurements only, the obtained similarity was 80%. Subsequently the gate was realized on two  $\text{Ca}^+$  ions [SKHR<sup>+</sup>03]. The most recent implementation of this gate was in Ref. [RKS<sup>+</sup>06] in which full process tomography was carried out on two ions and a fidelity of 93% was obtained.

The implementation of the CNOT gate described here uses the composite pulse sequence as described in Ref. [SKHR<sup>+</sup>03]. Stark shift corrections are implemented in the experiment control scheme by shifting reference frames as is done in NMR [VC05]. We perform truth table measurements as in [MMK<sup>+</sup>95] and obtain similarities (which can be considered as a classical fidelity) of 95%. Stark shifts and truth table measurements are described in [Lab08]. The following sections describe the extension of that work, in the implementation of full process tomography on the atomic and motional qubits.

## 4.2 Coherence

### 4.2.1 Motional state coherence

The Cirac-Zoller CNOT gate employs superpositions of ion motional states as intermediate states during the gate, and thus is sensitive to motional decoherence. In particular, a high ion heating rate will reduce the gate fidelity. An upper bound on the maximum heating rate tolerable,  $\dot{n}_{\text{max}}$ , can be given by consideration of the total time  $T_{\text{gate}}$  required for the pulse sequence implementing the CNOT gate, together with a design goal for the gate error probability  $p_{\text{gate}}$  desired. Assuming that a single quantum of change due to heating will cause a gate error, then  $\dot{n}_{\text{max}} < T_{\text{gate}}/p_{\text{gate}}$ . For  $T_{\text{gate}} \sim 230 \mu\text{s}$  (for our experiment), a heating rate of  $\dot{n}_{\text{max}} < 40$  quanta/s is needed to get  $p_{\text{gate}} \sim 0.01$ .

We measured the heating rate of the trap at a secular frequency of  $2\pi \times 1.32$  MHz. The number of motional quanta is measured by probing the blue and red sidebands of



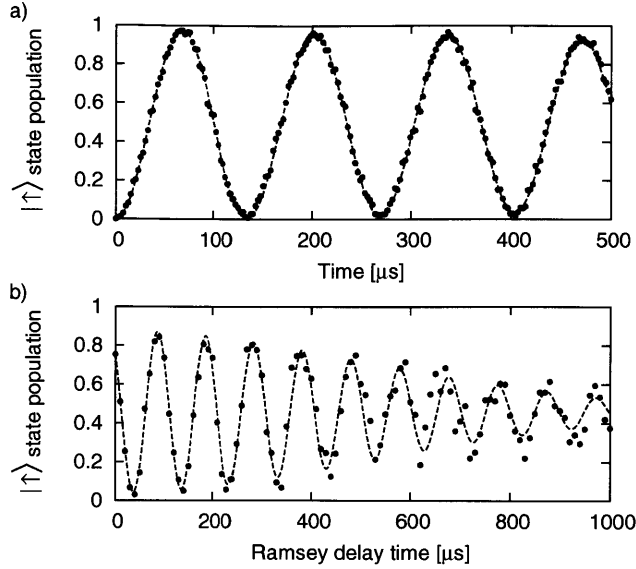


Figure 4-1: (a) Rabi oscillations on the blue sideband. The fitted initial contrast is 97.6(3)% and the frequency is 46.7 kHz. (b) Ramsey spectroscopy on the blue sideband. The fitted Gaussian envelope of the decay has time constant  $T_2^* = 622(37)\mu\text{s}$ .

the  $S \leftrightarrow D$  transition using the shelving technique, and comparing the ratio of shelving probability on each sideband, as described in Section 2.2.3. The heating rate is determined by varying the delay before readout and comparing the number of quanta versus delay time. The measured heating rate is weakly dependent on the RF voltage and DC compensation voltages. Noise on the RF pseudopotential can cause heating [WMI<sup>+</sup>98, BOV<sup>+</sup>09], so the ion micromotion is minimized using the photon correlation method [BMB<sup>+</sup>98]. For more details about the measurements, see Ref. [Lab08]. The heating rate can also depend on the trap's processing history and may vary between temperature cycles [LGL<sup>+</sup>08]; for this trap, the variation is small. In a typical experimental run, the RF voltage and DC compensation values are adjusted to minimize the heating rate before the coherence time and quantum gate data are taken. Typical heating rates obtained in this trap are 4-6 quanta/s, while the lowest heating rate measured is 2.1(3) quanta/s. Figure 4-1(a) shows Rabi flops on the blue sideband after the ion is initialized to the motional ground state with average number of quanta  $\bar{n} < 0.01$ . The fitted initial contrast is 97.6(3)% and the frequency is 46.7 kHz. Motional state coherence is demonstrated by performing Ramsey spectroscopy on the blue sideband (Figure 4-1(b)). The coherence time  $T_2^*$  is 622(37)  $\mu\text{s}$ . This is comparable to the coherence time of 660(12)  $\mu\text{s}$  of the atomic qubit as measured by the same method on the carrier transition.

#### 4.2.2 Magnetic field stabilization

When the optical qubit is encoded in a pair of levels that are first-order sensitive to magnetic fields, field fluctuations on the time scale of gate operations will decrease gate fidelity. One way of passively stabilizing the field is by use of a  $\mu$ -metal shield, which is expensive and inconvenient for optical access, and also mainly effective for low-frequency noise. Active

stabilization of the magnetic field using a flux gate sensor and coils has been implemented in another experiment [SKGR<sup>+</sup>03], at the cost of higher complexity.

Superconducting solenoids have been employed for passively stabilizing ambient magnetic field fluctuations in NMR experiments, with field suppression by a factor of 156 [GTC<sup>+</sup>91]. A similar method for ion traps which would permit good optical access is desired. In the NMR implementation, the field needs to be stabilized over a region 1 cm in length, whereas in an ion trap the region of interest is much smaller. Our method uses the same principle of superconductive shielding, but the small region and requirement for optical access suggest a more compact approach.

We stabilize the magnetic field by employing the persistent current in two superconducting rings, placed closely adjacent to the ion trap chip. This is a very compact and experimentally convenient arrangement, with high passive field stability and little barrier to optical access. Below the trap is a  $1 \times 1$  cm<sup>2</sup> square Nb plate with a 1.5 mm diameter hole, located 0.5 mm below the trap center. Above the trap is a 50 cm<sup>2</sup> square plate with an 11 mm diameter hole, located 7 mm above the trap center (see Figure 4-2(a)). Both rings are 0.5 mm thick. This geometry was chosen to optimize the field suppression at the trap location using the method to calculate magnetic fields in superconducting rings described in Ref. [BC03].

With a single trapped ion we measured the field suppression by applying a constant field with external coils, cooling the trap and Nb rings to below  $T_c$ , and reducing the field while measuring the  $S \leftrightarrow D$  transition frequency. The magnetic field is calculated from the Zeeman splitting between the  $m = -1/2 \leftrightarrow m = -5/2$  transition and the  $m = +1/2 \leftrightarrow m = -3/2$  transitions. A 50-fold reduction in field sensitivity was observed (Figure 4-2), in agreement with the numerical calculation. To determine the effectiveness of the noise suppression on coherence of the atomic qubit, we measured the decay of Ramsey fringes as a function of the separation of the Ramsey  $\pi/2$  rotations on the carrier  $S \leftrightarrow D$  transition. Such a measurement also includes effects caused by laser linewidth and the drift in laser-ion distance. We found that reducing the magnetic field noise by a factor of 50 did not improve the coherence time by more than a factor of 2, from  $T_2^* \sim 350 \mu\text{s}$  to  $\sim 660 \mu\text{s}$ . This suggests that magnetic field noise is no longer a dominant source of decoherence when compared to laser linewidth. Although this measurement was done under DC and the dominant source of magnetic field fluctuations is frequencies near 60 Hz and its harmonics, we can estimate the bandwidth of this compensation scheme by relating it to material properties of niobium as a type-II superconductor. The field suppression factor is determined by how fast the induced currents in the superconducting rings respond to changes in the external field, which depends on the ring's inductance (a geometric factor independent of frequency) and resistance. Above the first critical field, type-II superconductors exhibit flux pinning, which leads to ac resistance, but the critical field for niobium is on the order of 1000 G [PFCP07, BM69]. Below the critical field, superconductors can still exhibit a frequency-dependent AC resistance as described in Ref. [Kax03]. However, for niobium the effect is not significant until frequencies up to  $\sim 10^{12}$  Hz. Therefore at typical bias fields (4 G) and frequencies relevant to our qubit ( $<1$  kHz), niobium behaves as a perfect superconductor and we expect the field suppression factor to be the same as that measured under DC.

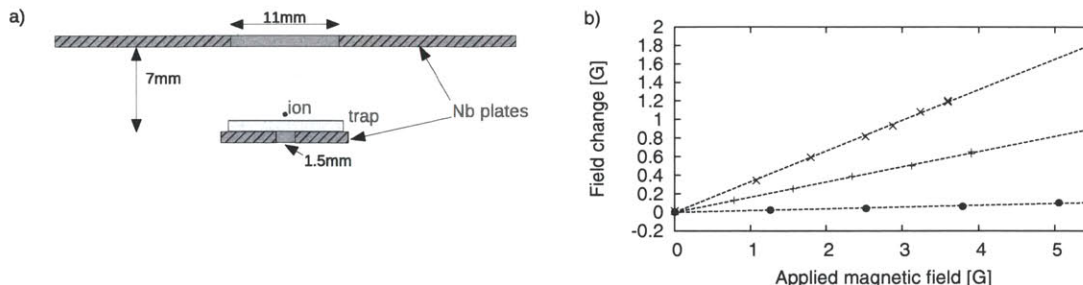


Figure 4-2: (a) Two superconducting disks, one below and one above the trapped ion, stabilize the magnetic field in the  $\hat{z}$  direction. (Not to scale.) (b) Magnetic field fluctuation suppression due to the top disk only ( $\times$ ), the bottom disk only ( $+$ ), and both disks ( $\bullet$ ). When both disks are used, field changes are suppressed 50-fold.

Greater reduction can be obtained by optimizing the geometry further, for example, by decreasing the distance between the plates to 4 mm, but is not implemented because of physical constraints in the apparatus. This method stabilizes the magnetic field only along the axis of the superconducting rings, but since the 4-G bias field defining the quantization axis is applied in the same direction, field noise in the x or y direction contributes only quadratically to the change in the total field [GTC<sup>+</sup>91].

### 4.3 Process tomography on a single ion

With  $N = 2$  qubits encoded in a single ion and methods of coupling and controlling these states, a Cirac-Zoller CNOT gate can be implemented [CZ95]. The CNOT gate is universal in that all quantum operations can be decomposed into single-qubit operations and the CNOT gate, and is thus of interest for implementing quantum information processing in ion traps. To evaluate the performance of such a gate, we prepare the system in a set of basis states that spans the space of  $2^N \times 2^N$  density matrices and perform a set of measurements that completely specifies the resulting state (state tomography). Quantum process tomography (QPT) is performed on the two qubits to construct the process matrix, allowing a full characterization of the gate. Section 4.3.1 gives a brief summary of state tomography using conditional measurements. Section 4.3.2 describes a minimal set of available measurements and operations in this two-qubit single-ion system necessary for QPT. Sections 4.3.3 and 4.3.4 list the pulse sequences for preparing all basis states and measuring the outcome. Section 4.3.5 briefly describes the construction of the process matrix that fully characterizes the gate from these measurements.

#### 4.3.1 Two-qubit state tomography for one ion

State tomography on the single-ion system of atomic and motional qubits requires a nontrivial set of operations, since a single qubit rotation on the motional qubit cannot be realized directly except by first swapping it with the internal state, performing the desired gate, then swapping back. The swap operation is complicated since the most straightforward

physical operations, red- and blue-sideband pulses, generally take the system out of the computational space, and into higher-order motional states such as  $|2\rangle$  [CC00]. For the CNOT gate, a set of composite pulse sequences can keep the system in the computational space. But if the goal is measurement of the two-qubit state space rather than the realization of a coherent operation, an alternative approach can be employed. A sequence of measurements, with the second conditioned on the results of the first, can suffice to allow full state tomography on the two-qubit atomic+motional state space. This is an extension of the single-ion tomography technique described in [MMK<sup>+</sup>95].

The conditional measurement sequence is as follows. First we apply an optional  $\pi$  pulse on the carrier transition; then the internal atomic state is measured by fluorescence detection. When this measurement scatters photons, it provides information about the internal state only and the motional state information is lost. When this first measurement does not scatter photons, a  $\pi$  pulse is applied on the blue-sideband transition, which allows measurement of the population in the state pairs  $\{|S0\rangle, |S1\rangle\}$  or  $\{|D0\rangle, |D1\rangle\}$ , depending on whether the initial carrier  $\pi$  pulse was applied or not. Two measurements, with and without the carrier pulse, are sufficient to determine the population in all four states.

### 4.3.2 Process tomography: operator definitions

The state tomographic measurement just described measures state populations only, which are the diagonal elements of the full density matrix. Relative phases between qubit states, which determine coherence properties of the state, are also needed in order to perform complete process tomography. The phases can be obtained by appropriate rotations of the qubits prior to measurement. Here we define the measurement and rotation operators for the following sections.

The single available measurement is the usual fluorescence detection, which is a projective measurement into the  $|S\rangle$  state, denoted  $P_S$ . Let  $P_D$  denote a projection into the  $|D\rangle$  state. The matrices for  $P_S$  and  $P_D$  in the basis  $|D0, D1, D2, S0, S1, S2\rangle$  are

$$P_S = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \otimes I_3. \quad P_D = I_6 - P_S. \quad (4.1)$$

The available unitary operations are as follows:

- $R_x(\theta), R_y(\theta)$ : Single qubit (carrier) rotations on the  $\{|S\rangle, |D\rangle\}$  qubit.
- $R_x^+(\theta), R_y^+(\theta)$ : Blue-sideband rotations, connecting  $\{|S0\rangle, |D1\rangle\}$  and  $\{|S1\rangle, |D2\rangle\}$  (neglecting higher-order vibrational modes).  $\theta$  is the rotation angle on the  $\{|S0\rangle, |D1\rangle\}$  manifold.
- Red-sideband rotations can be defined similarly, but are actually not necessary for construction of a complete measurement set.

$\Psi(i)$	Operations applied to $\Psi_0$	State
$\Psi(1)$	$R_y(-\pi)$	$ D0\rangle$
$\Psi(2)$	$R_x^+(-\pi)$	$ D1\rangle$
$\Psi(3)$	$I$	$ S0\rangle$
$\Psi(4)$	$R_y(\pi)R_x^+(\pi)$	$ S1\rangle$
$\Psi(5)$	$R_x^+(\pi)R_y(-\pi/2)$	$( D0\rangle +  D1\rangle)/\sqrt{2}$
$\Psi(6)$	$R_y^+(-\pi)R_y(-\pi/2)$	$( D0\rangle + i D1\rangle)/\sqrt{2}$
$\Psi(7)$	$R_y(-\pi/2)$	$( D0\rangle +  S0\rangle)/\sqrt{2}$
$\Psi(8)$	$R_x(\pi/2)$	$( D0\rangle + i S0\rangle)/\sqrt{2}$
$\Psi(9)$	$R_y(-\pi)R_x^+(-\pi/2)$	$( D0\rangle +  S1\rangle)/\sqrt{2}$
$\Psi(10)$	$R_y(-\pi)R_x^+(\pi/2)$	$( D0\rangle + i S1\rangle)/\sqrt{2}$
$\Psi(11)$	$R_x^+(\pi/2)$	$( D1\rangle +  S0\rangle)/\sqrt{2}$
$\Psi(12)$	$R_y^+(\pi/2)$	$( D1\rangle + i S0\rangle)/\sqrt{2}$
$\Psi(13)$	$R_y(\pi/2)R_x^+(\pi)$	$( D1\rangle +  S1\rangle)/\sqrt{2}$
$\Psi(14)$	$R_x(-\pi/2)R_x^+(-\pi)$	$( D1\rangle + i S1\rangle)/\sqrt{2}$
$\Psi(15)$	$R_y(-\pi)R_x^+(-\pi)R_y(\pi/2)$	$( S0\rangle +  S1\rangle)/\sqrt{2}$
$\Psi(16)$	$R_y(-\pi)R_y^+(\pi)R_y(\pi/2)$	$( S0\rangle + i S1\rangle)/\sqrt{2}$

Table 4.1: State preparation operations.

Explicitly, these rotation matrices are defined as follows:

$$\begin{aligned}
R_x(\theta) &= \exp[-i\theta(\sigma_x \otimes I_3)] \\
R_y(\theta) &= \exp[-i\theta(\sigma_y \otimes I_3)] \\
R_x^+(\theta) &= \exp[\theta(\sigma_+ \otimes a^\dagger - \sigma_- \otimes a)/2] \\
R_y^+(\theta) &= \exp[-i\theta(\sigma_+ \otimes a^\dagger + \sigma_- \otimes a)/2]
\end{aligned} \tag{4.2}$$

where  $a^+$  and  $a$  are the creation and annihilation operators in the Jaynes-Cummings Hamiltonian.

### 4.3.3 State preparation

For every measurement sequence, the ion is initialized to the state  $\Psi_0 \equiv |S0\rangle$ . The sequences of operations listed in Table 4.1 generates the 16 input states that span the space of  $4 \times 4$  density matrices created from the product states  $|D0, D1, S0, S1\rangle$ .

### 4.3.4 Complete basis of measurements

The following is a procedure for performing complete state tomography of the two-qubit  $\{|S\rangle, |D\rangle\} \otimes \{|0\rangle, |1\rangle\}$  state of a single ion, using the measurements and operations in Section 4.3.2. This is a generalization of the method used to measure just the diagonal elements of the density matrix. There are two kinds of measurement used; we call them  $M_U$  and  $M_{UV}$ .

$M_U$  involves performing a unitary operation  $U$  on the input state and then projecting

$M_j$	Measurement functions
$M_1$	$M_U(I)$
$M_2$	$M_{UV}(I, R_y^+(\pi))$
$M_3$	$M_{UV}(R_y(\pi), R_y^+(\pi))$
$M_4$	$M_U(R_y(\pi/2))$
$M_5$	$M_U(R_x(\pi/2))$
$M_6$	$M_{UV}(I, R_y(\pi/2)R_y^+(\pi/2))$
$M_7$	$M_{UV}(R_y(\pi), R_y(\pi/2)R_y^+(\pi/2))$
$M_8$	$M_{UV}(R_y(\pi/2), R_y(\pi/2)R_y^+(\pi/2))$
$M_9$	$M_{UV}(R_y(\pi/2), R_x(\pi/2)R_y^+(\pi/2))$
$M_{10}$	$M_{UV}(I, R_x(\pi/2)R_y^+(\pi/2))$
$M_{11}$	$M_{UV}(R_x(\pi), R_x(\pi/2)R_y^+(\pi/2))$
$M_{12}$	$M_{UV}(R_x(\pi/2), R_x(\pi/2)R_y^+(\pi/2))$
$M_{13}$	$M_{UV}(R_x(\pi/2), R_y(\pi/2)R_y^+(\pi/2))$
$M_{14}$	$M_{UV}(R_y(\pi/2), R_y^+(\pi/2))$
$M_{15}$	$M_{UV}(R_x(\pi/2), R_y^+(\pi/2))$

Table 4.2: State measurement functions.

into the  $|S\rangle$  subspace  $P_S$ . This is described by the measurement operator

$$M_U(U) = U^\dagger P_S U. \quad (4.3)$$

Typically,  $U$  will be a rotation in the  $\{|S\rangle, |D\rangle\}$  subspace, implemented by a carrier transition pulse.

$M_{UV}$  involves first performing a unitary operation  $U$  on the input state and making a measurement to detect fluorescence, which is equivalent to projecting to the  $|S\rangle$  subspace. Since  $|D\rangle$  is long lived, this projection leaves the  $\{|D0\rangle, |D1\rangle, \dots\}$  subspace undisturbed, but motional state information is lost if the ion is in state  $|S, n\rangle$ . If no fluorescence is detected, the post-measurement state is  $P_D \rho P_D$ . Conditioned on the first measurement returning  $|D\rangle$  (no fluorescence), a unitary transform  $V$  is performed, and finally another into the  $|S\rangle$  subspace  $P_S$ . If the first measurement returns fluorescence, the measurement sequence stops, in which case only information about the atomic state is obtained.  $M_{UV}$  is described by the measurement operator

$$M_{UV}(U, V) = U^\dagger P_D V^\dagger P_S V P_D U. \quad (4.4)$$

Typically,  $U$  will be a rotation in the  $\{|S\rangle, |D\rangle\}$  subspace, while  $V$  will be one or more rotations on the carrier and the red or blue sideband.

The measurements listed in Table 4.2 provide a complete basis of observables from which the full density matrix  $\rho$  can be reconstructed, assuming that  $\rho$  is initially in only the two-qubit computational subspace. These measurement observables are linearly independent.

The relationship between measurements and the density matrix can be expressed by a matrix  $A$  with elements

$$A_{ij} = M_j(\Psi(i)). \quad (4.5)$$

The full density matrix  $\rho$  can be reconstructed as:

$$\rho = \sum_{ij} m_j A_{ij}^{-1} |\Psi_i\rangle \langle \Psi_i| \quad (4.6)$$

where  $m_j$  is the result of measurement  $M_j$ .

### 4.3.5 Construction of the process matrix

A quantum gate including all error sources can be represented by the operation  $\mathcal{E}(\rho)$ , which can be written in the operator sum representation as

$$\mathcal{E}(\rho) = \sum_{mn} E_m \rho E_n^\dagger \chi_{mn} \quad (4.7)$$

where  $\rho$  is the input state and  $E_i$  is a basis of the set of operators on the state space. The process matrix  $\chi_{mn}$  contains the full gate information. For two qubits, the state space is spanned by 16 basis states, and  $16^2$  elements define the  $\chi$ -matrix, although it only has  $16 \times 15$  independent degrees of freedom because of normalization. This is reflected in the fact that only 15 measurements are needed. The  $\chi$  matrix can be obtained by inverting the above relation. To avoid unphysical results (namely, a non-positive-semidefinite  $\rho$ ,  $\text{Tr}(\rho^2) > 1$ ) caused by statistical quantum error in the experiment, a maximum-likelihood estimation algorithm [JKMW01] is employed to determine the physical operation  $\mathcal{E}$  that most likely generated the measured data. An alternate, iterative algorithm is presented in Ref. [JcvFcvH03].

## 4.4 Results

### 4.4.1 Gate performance

Quantum process tomography was carried out to evaluate the performance of various gates on the two qubits of a single ion. The ion in its motional and atomic ground state is initialized to one of the 16 input states in Table 4.1. Then the gate is applied, and the output state is determined by making all of the measurements listed in Table 4.2. The longest duration of the full measurement sequence (excluding the gate) is 610  $\mu\text{s}$ , and a single CNOT gate takes 230  $\mu\text{s}$ . These durations are determined by the Rabi frequency on the carrier  $\Omega = 2\pi \times 125$  kHz and on the sideband  $\Omega_{\text{BSB}} = 2\pi \times 7.7$  kHz, and the secular frequency  $\omega_{\text{sec}} = 2\pi \times 1.32$  MHz. The resulting  $\chi$  matrix for the CNOT gate is shown in Figure 4-3.

We evaluate the performance of the identity gate (all preparation and measurement sequences performed with no gate in between), the single CNOT gate, and two concatenated CNOT gates (CNOTx2). The results are shown in Table 4.3. The process fidelity is defined as  $F_p = \text{Tr}(\chi_{\text{id}} \chi_{\text{expt}})$ , where  $\chi_{\text{id}}$  is the ideal  $\chi$  matrix calculated with the ideal unitary operation  $U$ , and  $\chi_{\text{expt}}$  is experimentally obtained using maximum-likelihood estimation. We also calculate the mean fidelity  $F_{\text{mean}}$ , based on the overlap between the expected and measured density matrices,  $\text{Tr}(\rho_{\text{id}} \rho_{\text{expt}})$ , averaged over all prepared and measured basis



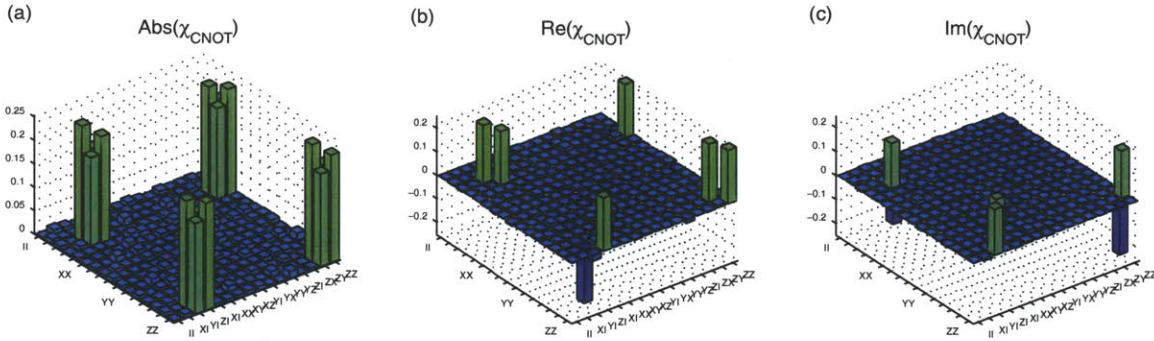


Figure 4-3: Process tomography on the CNOT gate. (a), (b), and (c) show the absolute, real, and imaginary parts of the  $\chi$  matrix, respectively.

Gate	$F_p$ (%)	$F_{\text{mean}}$ (%)
Identity	90(1)	94(3)
CNOT	85(1)	91(5)
CNOTx2	81(1)	89(6)

Table 4.3: Measured gate fidelities for the identity gate, the single CNOT gate, and two concatenated CNOT gates.

states, as in Ref. [OPG<sup>+</sup>04].  $F_p$  characterizes the process matrix whereas  $F_{\text{mean}}$  is a more direct measure of the gate performance. There exists a simple relationship between the two measures,  $F_{\text{mean}} = (dF_p + 1)/(d + 1)$  [GLN05], which is consistent with the independently calculated values for our data. Error bars on  $F_p$  are calculated from quantum projection noise using Monte Carlo methods [RLR<sup>+</sup>04]. The large error bars on  $F_{\text{mean}}$  occur because certain measured basis states consistently have a higher or lower overlap with the ideal states. In general, states that involve multiple pulses to create entanglement are more susceptible to error and therefore have a lower fidelity than states that are closer to pure states. The pulse sequence for some states essentially performs a CNOT gate to create and remove entanglement; thus imperfect state preparation and measurement contributes significantly to the overall infidelity. Using the data for 0, 1, and 2 gates, we can estimate the fidelity of a single CNOT gate normalized with respect to the overall fidelity of the state preparation and measurement steps. Assuming that the fidelity of the  $n$ th gate is  $F_p^n = F_i(F_g)^n$ , where  $F_i$  is the preparation and measurement fidelity, the fitted fidelity per gate,  $F_g$ , is 95%.

#### 4.4.2 Simulation and error sources

A number of possible error sources and their contributions to the process fidelity of the single CNOT gate are listed in Table 4.4. To estimate and understand error sources, we simulated the full system evolution in the (2 atomic state)  $\times$  (3 motional state) manifold using the exact Hamiltonian, including Stark shift and tomographic measurements. The magnitude of each source is measured independently and then added to the simulated pulse sequence. Laser frequency fluctuation is assumed to be the primary cause of decoherence



Error source	Magnitude	Approx. contribution
Off-resonant excitations	1%	10%
Laser frequency fluctuations	300 Hz	5%
Laser intensity fluctuations	1%	1%
Total		15%

Table 4.4: Error budget listing the major sources of errors on the process fidelity of the single CNOT gate, obtained by simulation. Each error source is assumed to be independent. The total error is calculated as the product of individual errors.

and is measured by observing the decay of Ramsey fringes on the carrier transition. The frequency fluctuation is simulated as a random variable on the laser frequency which grows in amplitude over time, and accounted for via Monte Carlo techniques. Laser intensity fluctuations are measured directly with a photodiode. On short time scales comparable to the length of the gate, the fluctuations are  $\sim 0.1\%$  peak to peak; on longer time scales, up to 1% drifts are observed. Both of these effects are accounted for in the simulation. Off-resonant excitations are automatically included in the model of the full Hamiltonian. The effect can be removed from the simulation if decoherence is not included and the simulated pulses are of arbitrarily long lengths, which is equivalent to reducing the laser intensity. The resulting  $\chi$  matrix and fidelity agree well with the measured results, indicating that the observed fidelity is well understood in terms of technical limitations.

Off-resonant excitations, caused by the square pulse shape used to address all transitions, is expected to be the largest source of error, as previous work has found [RKS<sup>+</sup>06]. Square pulses on the blue-sideband transition contain many higher harmonics, which causes residual excitation of the carrier transition. The carrier transition oscillations caused by this can be measured directly, averaged over many scans because of their small amplitude. Although the measured amount of off-resonant excitation is small ( $\sim 1\%$ ) for the laser intensity and secular frequencies used for our gates, both our simulations and previous work [RKS<sup>+</sup>06] have found that up to 10% improvement in gate fidelity can be gained by implementing amplitude pulse shaping.

## 4.5 Outlook

In summary, we have developed a cryogenic, microfabricated ion trap system and demonstrated coherent control of a single ion. The cryogenic environment suppresses anomalous heating of the motional state, as well as enables the use of a compact form of magnetic field stabilization using superconducting rings. A complete set of pulse sequences for performing quantum process tomography on a single ion's atomic and motional state is implemented. These components are sufficient to perform a CNOT gate on the atomic and motional state of a single ion. It is expected that amplitude pulse shaping would further improve the gate fidelity. These techniques, realized in a relatively simple experimental system, make the single ion a possible tool for studying other interesting quantum-mechanical systems.

The control techniques and the CNOT gate demonstrated in this work focus on a single ion and do not constitute a universal gate set for scalable quantum computation. However,

the additional requirements for such a universal two-ion gate, including individual addressing [WLG<sup>+</sup>09] and readout of two ions, have been realized in traditional 3D Paul traps as well as other surface trap experiments, and are not expected to pose significant challenges. The microfabricated surface-electrode ion trap operated in a cryogenic environment thus offers a viable option for realizing a large-scale quantum processor.

#### 4.5.1 Disadvantages and alternatives

While the successful demonstration of the Cirac-Zoller CNOT gate here appear to be limited by only technical sources of noise and error, the gate fidelity obtained in this work as well as in other successful implementations (mean fidelity of 93% in Ref. [RKS<sup>+</sup>06]) are still far from the fault-tolerant limit of 99.99% [Kni10]. There are several practical reasons. The contrast of Rabi flops on the blue sideband is maximized for an ion in the ground state. Therefore the CZ CNOT gate requires ground state cooling, which is a stringent requirement, especially for surface-electrode traps where heating due to anomalous electric field noise or due to moving ions between different zones [BOV<sup>+</sup>09] is of concern. Although not investigated in our work, the CZ CNOT gate also requires individual addressing of ions, which is a significant technical challenge.

Consequently the community has been moving toward alternative multi-ion gates which are less sensitive to such technical noise sources. One example is the bichromatic Mølmer-Sørensen gate [SM99]. This entangling gate is realized with a bichromatic light source: two lasers at near the red and blue sideband frequencies are applied simultaneously. This realizes a  $\sigma_x \otimes \sigma_x$  Hamiltonian, and the resulting operation is equivalent to a controlled-NOT up to single-qubit rotations [KBZ<sup>+</sup>09]. By generating the bichromatic gate from a single light source, intensity fluctuations are canceled out. Furthermore the gate is a global operation and thus does not require individual addressing. Typical gate times, for similar experimental parameters, can also be much shorter than the CZ CNOT gate, thus making it less susceptible to laser frequency fluctuations. The MS gate has been demonstrated with a fidelity of 99.3% in a 3D ion trap [BKRB08]. Another gate known as the conditional phase gate implements the  $\sigma_z \otimes \sigma_z$  operator, and has also been demonstrated in 3D ion traps with 97% fidelity [LDM<sup>+</sup>03]. No clear impediments to implementing these gates in a surface-electrode trap are known.

**Part II**

**Sensors**



## Chapter 5

# Heating rates

In most proposed implementations of trapped ion quantum computing, the shared motional state of ions trapped in a harmonic potential serves as quantum bus [CZ95, SM99]. Motional heating during a computation thus can cause decoherence and information loss in multi-qubit gates. Unlike other trap parameters such as secular frequencies, the heating rate of a given trap is not a direct result of trap design and cannot be predicted ahead of time with any degree of accuracy.

Presently, the cause of motional heating of trapped ions is still an unsolved problem despite a decade of effort in both theory and experiment. For microfabricated traps, where the ion-to-surface distance is on the order of hundreds of  $\mu\text{m}$ , the observed heating rates are orders of magnitude higher than larger traps at room temperatures. While the functional relationship between ion-surface distance and heating rate is still under debate –  $1/d^2$  to  $1/d^4$  scaling are currently considered – the general trend is that heating rate rapidly increases as the distance decreases, such that with microfabricated traps, it can become the dominant source of gate infidelity [ESL<sup>+</sup>07].

The acceptable amount of heating in a trap is determined by the speed of quantum gates. For the Mølmer-Sørensen entangling gate, the gate time is determined by the secular frequency and Rabi frequency; for a typical trap, the gate time is on the order of 10-100  $\mu\text{s}$  [BKR08]. Assuming that a single quantum gain during the computation will cause a gate error, then for a desired gate fidelity of 99%, a heating rate of  $< 100\text{-}1000$  quanta/sec is needed.

In surface-electrode ion traps with characteristic distance  $d < 100 \mu\text{m}$  operating at room temperature, typical heating rates exceed this value [ESL<sup>+</sup>07]. However, by cooling the trap to cryogenic temperatures, the heating rate can be reduced to be comparable to large traps at room temperature and such that it no longer is the limiting factor in gate fidelity. Previous work has been done to characterize the frequency and temperature dependence of heating rate at cryogenic temperatures for silver and gold traps [LGA<sup>+</sup>08, LGL<sup>+</sup>08].

This chapter presents measurements of heating rate of microfabricated traps made of aluminum and niobium. Aluminum is of interest due to its compatibility with conventional CMOS fabrication techniques, but the presence of a native oxide layer may lead to additional complications. Niobium, as a superconductor, is qualitatively different than the normal metals discussed so far, and measuring its heating rate above and below the superconductive

transition can provide insight on whether the source of noise is on the surface or in the bulk of the material. Section 5.1 gives a review of the relevant literature and currently accepted models of anomalous heating, as well as describe the measurement method used in this work. Section 5.2 describes the fabrication, testing, and heating rate characterization of aluminum traps, including two vastly different designs. Section 5.3 describes heating rate experiments with superconducting traps. Section 5.3 has been published in [WGL<sup>+</sup>10].

## 5.1 Electric field noise

Heating of the ion is believed to be caused by electric field noise coming from electrodes located near the ion position. The physical model for the origin of the electric field noise determines the scaling relationship between heating and other physical parameters such as frequency, temperature, and distance.

### 5.1.1 Models of noise sources

Electric field noise at the motional frequency of the ion couples to the ion's charge, and leads to a fluctuating force that drives the ion out of its motional ground state. This relation is given by [TKK<sup>+</sup>00]

$$\dot{n} = \frac{Q^2}{4m\hbar\omega_m} S_E(\omega_m) \quad (5.1)$$

where  $S_E(\omega_m)$  is the spectral density of electric field fluctuations in units of  $V^2/m^2/Hz$ ,  $Q$  is the ion charge,  $m$  is the ion mass,  $\omega_m$  is the secular frequency of motional mode  $m$ , and  $\dot{n}$  is the heating rate in units of quanta/sec.

Several models of the sources of electric field noise have been proposed. Thermal electronic noise (Johnson noise) can be related to the resistance in trap electrodes and connecting circuits. However, for typical values in microfabricated traps, the heating rate predicted by this model is orders of magnitude lower than observed in experiments [TKK<sup>+</sup>00, Lei09]. It also predicts a frequency, distance, and temperature dependence with heating which does not agree with experiment. A more promising theory is that of patch potentials: electric field fluctuations on randomly oriented domains of dipoles which may be uncorrelated [TKK<sup>+</sup>00]. Further refinements of this model include the introduction of correlation lengths [DCG09b], a microscopic model of dipoles formed by surface adsorbed atoms [SNRWS11], and considerations based on electrode geometry [LHC11]. These models predict a distance ( $1/d^4$ ) and frequency ( $1/f$ ) dependence of heating rate which agrees better with those observed in various experiments. However, these models have yet to provide insight on how to reduce heating rates in practice.

### 5.1.2 Measurement method

Measurement of heating rate is generally done by cooling the ion to near the motional ground state, impose a fixed amount of delay, then probing the final motional state. The heating rate is determined by varying the delay before readout and comparing the number of quanta versus delay time. The number of motional quanta is measured by probing the blue

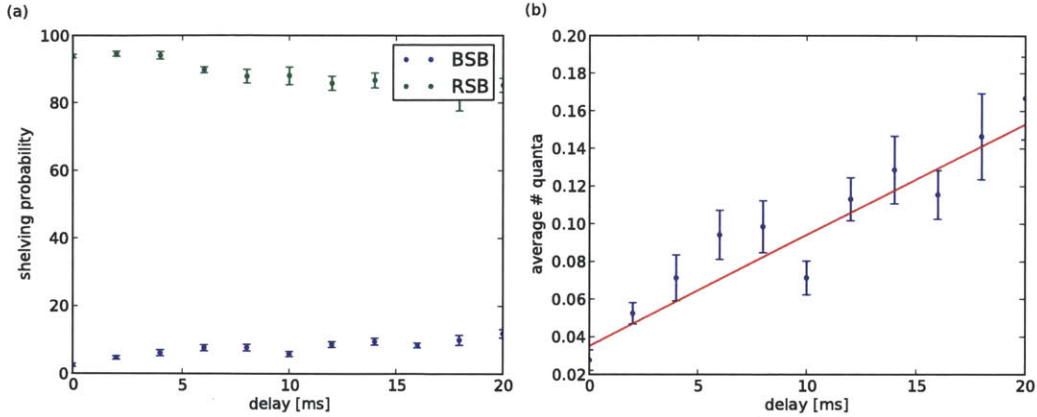


Figure 5-1: Example of heating rate data. (a) Plot of red sideband and blue sideband population as a function of delay time. (b) Average quanta vs delay time.

and red sidebands of the  $S \leftrightarrow D$  transition using the shelving technique, and comparing the ratio of shelving probability on each sideband, as described in Section 2.3.2. The average  $n$  as a function of the transition probabilities are then given by

$$\langle n \rangle = \frac{P_{red}}{P_{blue} - P_{red}}. \quad (5.2)$$

In the experiment, the transition probabilities are determined by making 100 measurements at each value of delay time (see Figure 5-1). Thus there is an error on each measured probability given by a binomial distribution. For  $\langle n \rangle = 0$ , the expected transition probabilities are  $P_{blue} = 1$  and  $P_{red} = 0$ . Propagating the errors on  $P_{blue}$  and  $P_{red}$  as  $\langle n \rangle$  increases into equation 5.2 gives the error on  $\langle n \rangle$ . In practice, this means that this measurement method is only valid for small values of  $\langle n \rangle$ ; a desired error of 50% bounds  $\langle n \rangle < 2$ . The maximum delay time is chosen such that the maximum  $\langle n \rangle$  at the end of the measurement time is below this value. Each measurement of  $\langle n \rangle$  is repeated 10 times to obtain the mean and the standard error. The heating rate is given by the slope of a linear fit to the data of  $\langle n \rangle$  vs delay time. The error on the fitting parameter is quoted as the error on the heating rate. This constitutes a single measurement of the heating rate for a particular instance of a trap.

A number of factors have been observed to contribute to larger measured heating rates. Noise on the RF source and DC electrodes at the trap's secular frequency can couple to the ion and drive the motional state. In theory, these effects should be minimized by adjusting the DC compensation voltages such that the ion is trapped at the overlap of the RF null and the DC null. However in some cases, direct observation of micromotion compensation in some axis is not possible and thus complicates the measurement. In our experiment, the projection of the vertical axis on the cooling laser is small and prevents a simple measurement of the micromotion along that axis. By measuring the heating rate at various RF and DC settings, we have observed that some traps have heating rates that strongly depend on finding an optimal RF and DC setting, whereas other traps exhibit

weak dependence.

In practice, there are other factors which were not fully characterized in our study, but which may lead to a higher measured heating rate than what is intrinsic to a particular material used for a trap. These include surface contamination from exposure to air or ion loading from a thermal oven; electrical noise on the DC/RF voltage supplies; magnetic field noise from the environment; etc. However, none of these sources should cancel out local electric field noise on the trap. Therefore any single measurement of heating rate should be considered as an upper bound to the “intrinsic” heating rate of a particular trap material. In the following section, we quote the “lowest heating rate” to mean such a bound.

## 5.2 Aluminum traps

Aluminum is an appealing material for microfabricated ion traps, due to its compatibility with sophisticated CMOS fabrication techniques. However, aluminum rapidly forms a native oxide layer,  $\text{Al}_2\text{O}_3$ , which may cause undesired effects such as increased heating rates and light-induced charging (see Chapter 6). The contribution to electric field noise from dielectrics is supported by experiments using cantilevers which found increases in non-contact friction with dielectric layers [KLM06, SMS<sup>+</sup>01].

Our goal is to characterize heating rates of aluminum traps at cryogenic temperatures. We fabricated 9 aluminum traps in total; out of those, 7 were able to trap ions with long enough lifetimes to measure heating rates. Three of the traps tested have additional layers of oxide deposited to test the hypothesis that heating rate increases with oxide thickness. We also tested and characterized a multi-zone trap with very different geometry and trapping parameters.

### 5.2.1 Fabrication

Two main fabrication methods were developed, one with wet etch and one with lift-off, and will be briefly summarized here. The detailed process can be found in [Lac10]. All traps were fabricated on single-crystal quartz wafers, chosen for its high thermal conductivity at cryogenic temperatures and low RF loss.

#### Wet Etch

- diesaw the wafer into 1 cm<sup>2</sup> pieces; clean with a 3:1  $\text{H}_2\text{SO}_4:\text{H}_2\text{O}_2$  solution
- E-beam evaporate 400 nm of aluminum onto the substrate
- spin NR9-3000P photoresist at 3000rpm for 60 sec
- expose the pattern with a chrome mask (soft contact)
- develop, then remove the unwanted photoresist
- etch the aluminum with a phosphoric acid-based etchant

#### Liftoff

- diesaw and clean the wafer as above
- spin NR9-3000P photoresist
- expose the pattern with a chrome mask



Trap	ion height ( $\mu\text{m}$ )	fabrication	Al thickness ( $\mu\text{m}$ )	oxide thickness	comments
I	75	wet etch	1.5	native	
IIa	75	wet etch	1.2	native	no ions trapped
IIb	75	wet etch	1.2	native	no ions trapped
IIIa	100	wet etch	1.0	native	
IIIb	100	wet etch	1.0	native	
IIIc	100	wet etch	1.0	native	
O-5	100	lift off	0.4	5 nm	
O-10	100	lift off	0.4	10 nm	
O-20	100	lift off	0.4	20 nm	

Table 5.1: Summary of ion height and fabrication parameters for all aluminum traps tested.

- develop, then remove the unwanted photoresist
- E-beam evaporate 400 nm of aluminum onto the patterned substrate
- soak in acetone to remove the photoresist and any metal deposited on top

For storage, all traps were coated with S1813 resist to keep out dust and stored in normal room environments. About a day prior to installation, traps were packaged in a standard 101-pin Ceramic Pin Grid Array (CPGA) by gluing with epoxy (Varian TorrSeal) along with a 0.5 mm thick niobium spacer. Each trap electrode was wirebonded to two pins using an ultrasonic aluminum wirebonder. Filter capacitors (9 nF, Novacap 1206B103K501P) were glued to the CPGA and connected to each DC electrode to reduce RF pickup. No additional cleaning is done, apart from removing individual dust particles with a wirebonder’s wire, between packaging and installation in the cryostat. In cases where the same trap is tested over multiple cooldown cycles and removed from the vacuum in between, the traps are stored in a normal room environment with no additional cleaning steps done between installations.

A list of traps and fabrication-related characteristics are listed in Table 5.1.

## 5.2.2 Heating rate

As observed in the past with silver and gold traps, the aluminum traps showed large variation in heating rates between different traps fabricated in the same batch, between different cooling cycles on the same trap, even within the same trap with different DC compensation voltages and RF amplitudes. Figure 5-2 shows all the heating rates that were measured in the aluminum traps, on various days over several months. Each data point represents a single measurement along with its error as defined in Section 5.1.2. The secular frequencies used for all the measurements here are in the range of 0.74-0.90 MHz. Clearly, the error from a single measurement is insufficient to account for variations occurring in practice, such as temperature cycling, fluctuating magnetic field noise, etc.

No strong dependence on RF amplitude is observed in any of the traps. To determine dependence on DC electrodes, for each measurement, we changed the voltage on one of the radial compensation electrode, and then adjusted the voltage on the corresponding electrode to minimize micromotion in the radial direction. The predominant effect is on the vertical

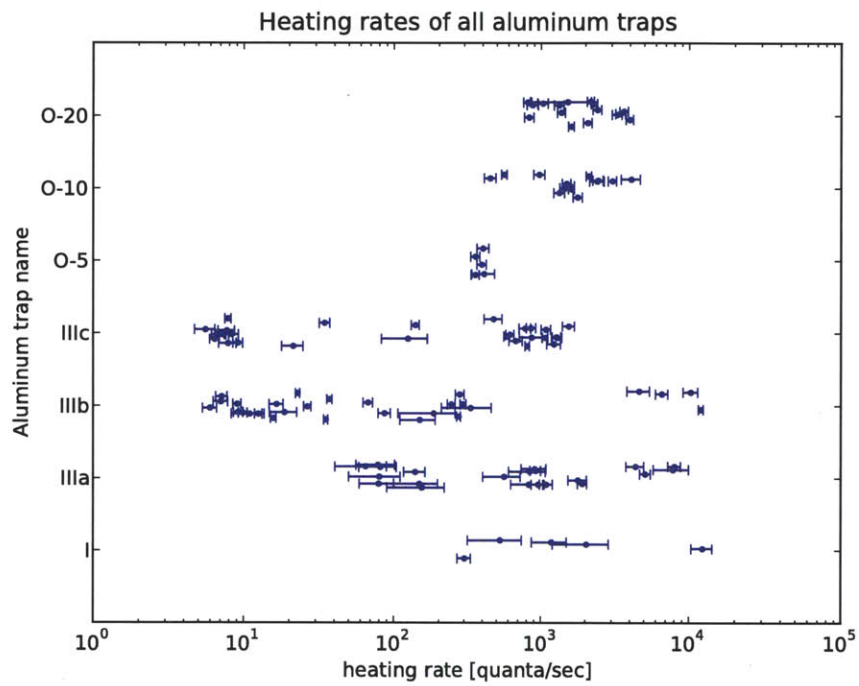


Figure 5-2: Summary of all heating rate data obtained in 7 aluminum traps. Each point represents a single measurement. Y-axis placement of each data point for each trap is arbitrary.

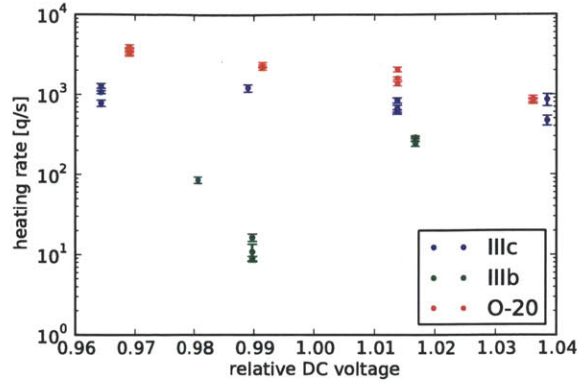


Figure 5-3: Heating rate as a function of radial DC compensation voltage for three aluminum traps. Voltages are relative to their mean value.

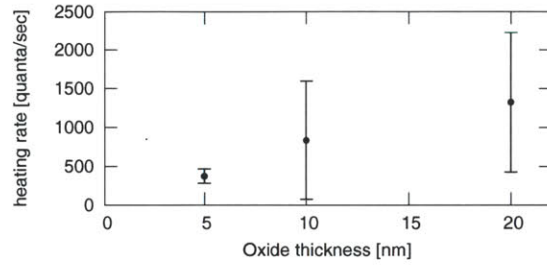


Figure 5-4: Heating rate as a function of oxide thickness for aluminum traps with extra layer of deposited oxide. Data is a weighted average of all data points shown in Figure 5-2.

position of the ion possibly relative to the RF null. In principle, if the ion is in the RF null, it becomes insensitive to noise on the RF supply. Two traps, O-20 and IIIc, showed some dependence on the voltage applied to one of the radial DC electrodes. In trap IIIb, a local minimum is observed (see Figure 5-3).

The lowest heating rate measured in aluminum traps is  $7.0(8)$  q/s, corresponding to an electric field noise of  $S_E = 84(10) \times 10^{-15}$  V<sup>2</sup>/m<sup>2</sup>/Hz. This value is comparable to the lowest heating rate measured in gold [LGL<sup>+</sup>08] and silver [LGA<sup>+</sup>08] traps in the same experiment. Thus, as an upper bound measurement, heating rates in aluminum traps appear not to be practically limited by the presence of an oxide layer.

The measured heating rates in aluminum traps with additional oxide layers appear to increase with respect to oxide layer thickness (Figure 5-4). However, given the orders of magnitude spread in measured heating rates of the native oxide traps, the possibility that not enough data was taken under sufficiently varying conditions for the oxide traps cannot be excluded.

### 5.2.3 SMIT-II trap

The SMIT-II trap (SMIT stands for Scalable Multiplexed Ion Trap) is an aluminum trap fabricated by the team of Jason Amini and Richard Slusher at Georgia Tech Research



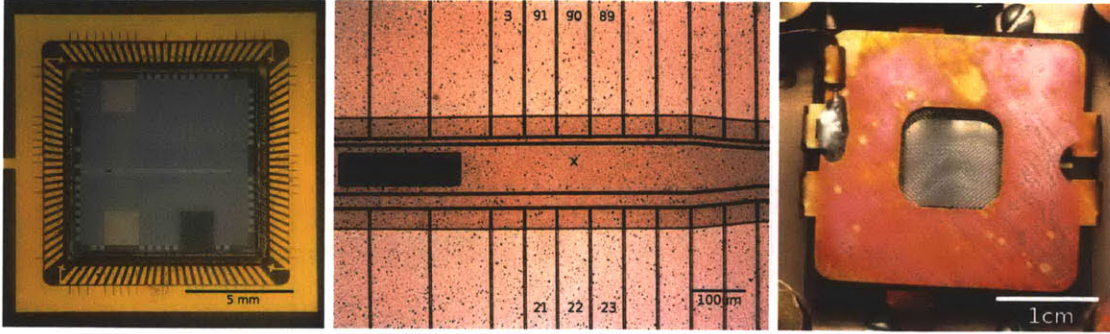


Figure 5-5: (a) Photographs of SMIT trap. (b) Microscope image of the trap showing the loading (not used) and trapping regions. The expected ion location is marked x. Numbered are DC electrodes 89-91 and 21-23; electrode 3 is used for RF tickle measurements. All other electrodes are grounded. (c) Trap installed in the cryostat with a stainless steel mesh to shield the trap from stray charges.

Institute. They consist of a top layer of aluminum, separated by  $10 \mu\text{m}$  of aluminum oxide from a ground plane. The expected ion height is  $67.7 \mu\text{m}$  and the expected secular frequency is 0.8 MHz in the axial direction and 3.40, 3.85 MHz in the radial directions.

The trap was tested in the same apparatus as all of the aluminum traps described in the previous section. A stainless steel mesh ( $\sim 80\%$  optical transmission) was installed 2 mm above the trap to reduce the effect of laser-induced charging on nearby dielectrics such as the detection optics.

We measured the axial secular frequency by taking a spectrum of the S-D transition and found the axial frequency to be 0.81 MHz at the designed DC voltages. The radial secular frequencies were measured by the RF tickle method: a small oscillating voltage is applied to one of the electrodes (in this case a nearby DC electrode) and the ion fluorescence is monitored while we sweep the oscillation frequency. Across a secular frequency, the oscillation drives the ion and induces an additional Doppler shift which can be detected as a change in the fluorescence signal. The additional DC electrode has an RC filter inside the cryostat with a cutoff frequency at 4 kHz, so up to 5 Vpp signal on the frequency generator was necessary to excite the radial modes. We measured radial mode frequencies of 3.408 and 3.887 MHz, in good agreement with the simulated values.

Heating rates in this trap were measured by taking a full scan of the blue and red sidebands, and extracting the maximum shelving probabilities on the blue and red sidebands for each delay. The full scan was necessary in this case because of drifts in the measured sideband frequency that is likely unrelated to the trap. For three different set of DC voltages corresponding to different degrees of vertical compensation, we obtained heating rates below 100 q/s as summarized in Figure 5-6. The three measurements of heating rates are consistent with each other within error bounds. These values are also an improvement over an earlier generation with a similar design and fabrication process described in Ref. [LLC<sup>+</sup>09].

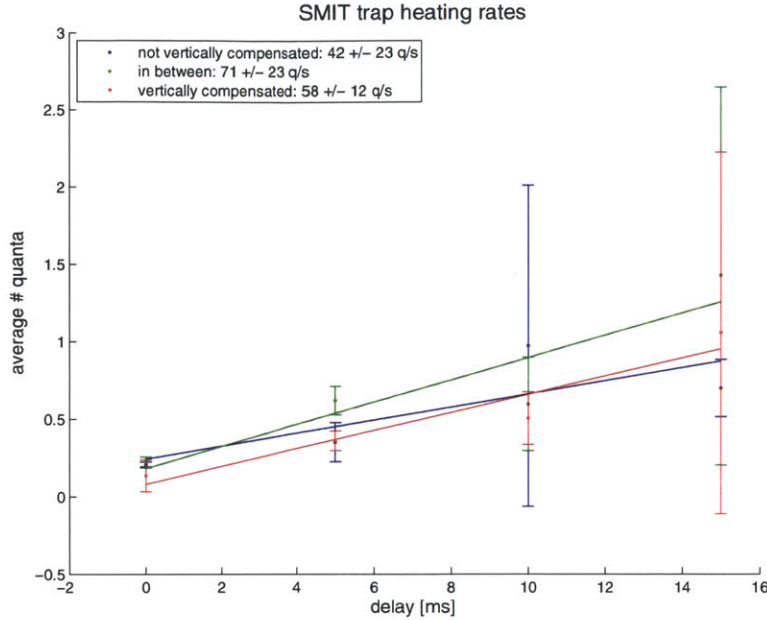


Figure 5-6: Heating rates measured in the SMIT-II trap.

### 5.2.4 Summary

We fabricated and measured heating rate in 7 aluminum traps. The lowest heating rate measured in an aluminum trap is comparable to those measured in gold and silver traps at the same cryogenic temperatures, indicating that the presence of an oxide layer is not a fundamental limit for anomalous heating in aluminum traps. However, all measured heating rates span up to several orders of magnitude for any single trap, making it difficult to summarize trends, such as with respect to oxide thickness. The heating rates in the SMIT-II trap, which has a very different design and fabrication process, are consistent with those measured in our aluminum traps, suggesting that trap design and fabrication may also not be the dominant variable in determining variations in the measured heating rates.

The additional data gathered in this section, in combination with the studies in silver and gold traps from previous work, suggests that trap material is not the dominant contribution to anomalous heating; heating rate data for more exotic materials such as molybdenum and boron-doped silicon [TKK<sup>+</sup>00, BLB<sup>+</sup>06] are also similar. Presently, the discovery of a “magic bullet” material which will make heating rate significantly lower seems unlikely.

## 5.3 Superconductor

In this section, we describe the fabrication and testing of superconducting ion traps. The anomalous heating affecting the gold, silver, and aluminum traps studied thus far provides significant motivation to explore qualitatively different materials for microfabricated ion traps, such as superconductors. In particular, the fact that a superconductor expels electric fields provides an opportunity to test the theoretical understanding that anomalous noise results from surface patch potentials [TKK<sup>+</sup>00, DCG09b], rather than sources in the



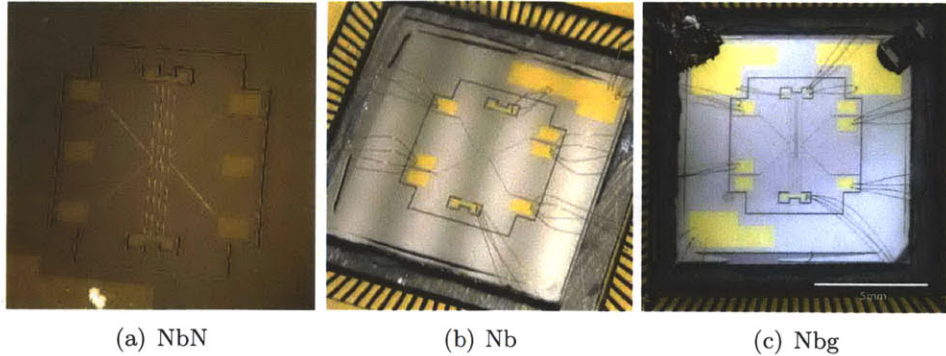


Figure 5-7: Photographs of (a) NbN and (b,c) Nb traps. The Nb trap in (c) has a  $1\text{ k}\Omega$  resistor for heating the trap (top right corner) and a copper braid to thermally connect the trap to the helium baseplate (top left corner).

bulk, since the bulk noise sources would be screened by the superconductor. A similar approach was taken for neutral atoms, where in superconducting traps it was found that magnetic near-field noise is suppressed resulting in lower heating rate and longer spin-flip lifetimes [FMZ<sup>+</sup>10, KHC<sup>+</sup>10]. For a thin film superconducting ion trap, blue lasers are typically employed for Doppler cooling and state detection of trapped ions, and the short (279–422 nm [Jam98]) wavelengths may create quasiparticles in the superconductor, driving it into a normal state. Therefore, verifying that the superconductor employed is actually superconducting during an experiment is required.

We demonstrate the operation of several superconducting microfabricated ion traps made with niobium and niobium nitride, describe how superconductivity is verified during trap operation, and measure heating rates above and below  $T_c$  to test the physical source of anomalous heating. The demonstration of superconducting ion traps opens up possibilities for integrating trapped ions and molecular ions with superconducting devices, such as photon counting detectors, microwave resonators [SBC<sup>+</sup>11], and circuit-QED systems [TRBZ04].

### 5.3.1 Fabrication

The superconducting ion traps consist of Nb or NbN on a sapphire substrate. Two Nb trap (Nbg-Ia and Nbg-Ib) include a thin wire structure (Figure 5-8) on the center ground electrode that is electrically connected in a 4-wire configuration to measure the resistivity of the electrode. The thinnest part of the wire is  $10\text{ }\mu\text{m}$  wide.

The fabrication procedure is as follows. A 400 nm layer of Nb is grown by DC magnetron sputtering of a niobium target in Ar gas; NbN is grown by adding  $\text{N}_2$  gas during sputtering. Electrodes are defined by optical lithography using NR9-3000P photoresist, exposed through a chrome mask and developed in RD6 developer. Reactive ion etch with  $\text{CF}_4$  and  $\text{O}_2$  is used to etch exposed metal. Gold contact pads for wirebonding are then defined by optical lithography using S1813 or NR9-3000PY photoresist, deposited using evaporation and created with a lift-off process. After the initial Nb sputtering, the trap is maintained below a temperature of  $90^\circ\text{C}$  during all steps of the fabrication and packaging process to

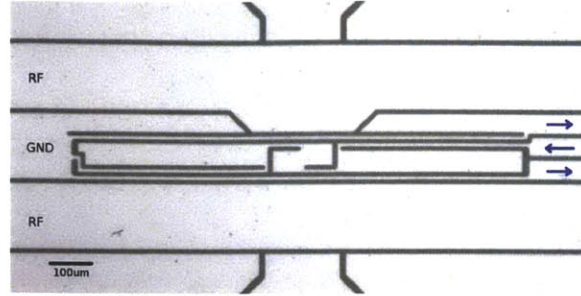


Figure 5-8: Microscope image of Nb trap with wire structure for measuring electrode resistivity.

minimize oxide formation on the surface. For trap Nbg, a surface-mount resistor (0603, 1 k $\Omega$ ) is glued to one trap corner and used as a heater for controlling the trap temperature. We estimate the trap surface temperature to be  $\sim 6$  K [ASA<sup>+</sup>09].

### 5.3.2 Superconductivity

We verify that the traps are superconducting by observing three variables: resistance, critical current, and reflected RF power. Resistance is measured on the wire structure in Nbg as the trap cools or warms up. Figure 5-9(a) shows the resistance as a function of measured baseplate temperature (with a Lakeshore RX103 calibrated diode) during a slow warm-up of the cryostat. The trap is heated to above  $T_c$  during ion loading, but cools to below  $T_c$  within 5-10 minutes. Superconductivity is maintained on the trap when 150 V (amplitude) of RF drive is applied to the trap RF electrodes to create the trapping potential. This corresponds to  $\sim 250$  mA of current on the RF electrodes, given a capacitance to ground of  $\sim 8$  pF. The critical current of the wire structure, both with and without the trapping lasers, is 180(1) mA, as shown in Figure 5-9(b). This corresponds to a critical current density of  $4 \times 10^6$  A/cm<sup>2</sup>, typical in order-of-magnitude for thin-film Nb. Based on this measured critical current density and electrode dimensions ( $400$  nm  $\times$   $150$   $\mu$ m), the calculated current limit on the RF electrodes is 2.7 A, well above what is needed for typical trap operations. The superconducting transition is also observed by looking at the reflected RF power in the NbN trap, which is more resistive immediately above  $T_c$ . RF reflected power is measured with a directional coupler mounted before the helical resonator, which is inside the cryostat. As shown in Figure 5-9(c), almost all power is reflected back on resonance below  $T_c$ . These methods confirm that the traps are superconducting while ions are trapped in the presence of lasers and RF current drive.

When the wire structure is current biased with 1-10 mA less than the critical current, the 405 nm, 422 nm, and 460 nm lasers grazing incident on the trap cause it to transition to the normal state. However, under normal trapping conditions, the lasers have no effect on the measured critical current.

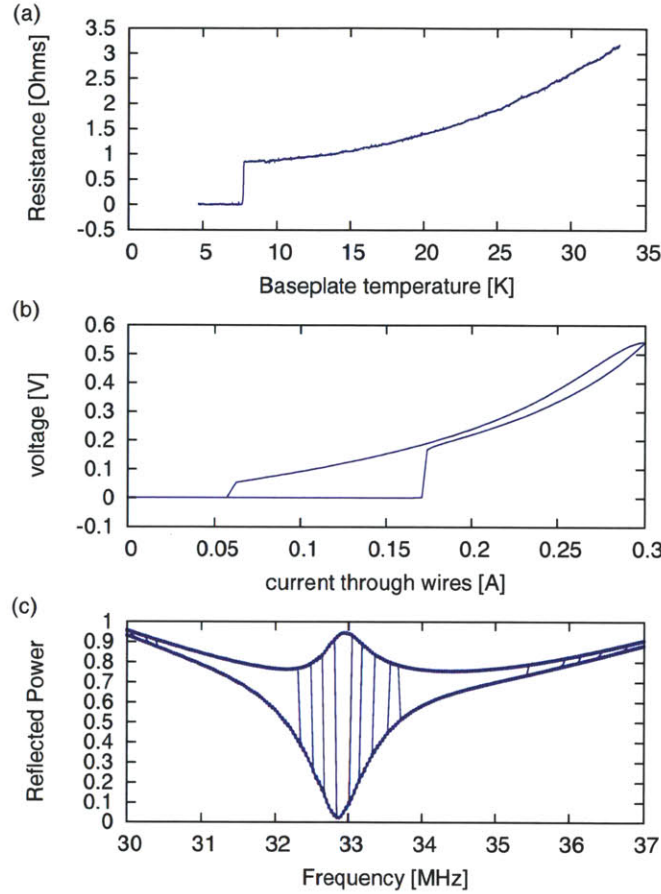


Figure 5-9: (a) Resistance in trap Nbg during warm-up of the cryostat. The baseplate temperature is lower than the trap temperature. (b) Measurement of critical current through the wire structure on trap Nbg. Transition to the normal state occurs at 180 mA. (c) Fraction of reflected RF power vs frequency in the NbN trap during warmup of the cryostat; top curve: below  $T_c$ , bottom curve: above  $T_c$ . The observed value switches multiple times between the two curves due to warmup and cooldown of the trap as the RF is moved on and off the resonant frequency.



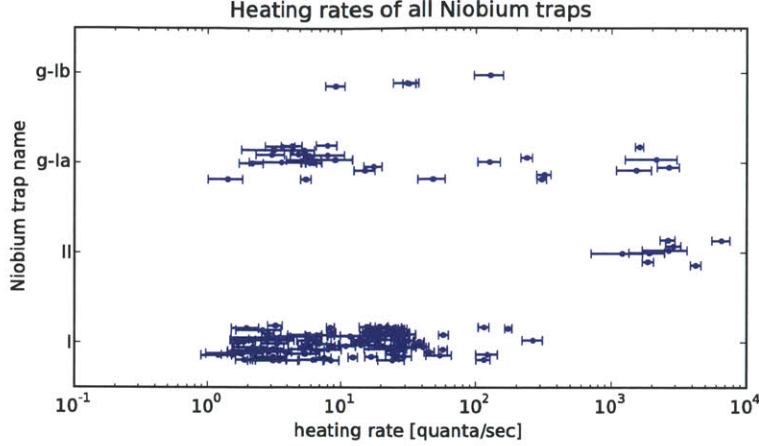


Figure 5-10: Summary of all heating rate data obtained in 4 niobium traps. Each point represents a single measurement. Y-axis placement of each data point for each trap is arbitrary.

Trap	heating rate (q/s)	$S_E$ ( $10^{-15}$ V <sup>2</sup> /m <sup>2</sup> /Hz)
NbN	16(1)	192(12)
Nb	2.1(3)	25(4)
Nbg	4.2(8)	48(12)
Au <sup>a</sup>	2.1(4)	25(5)
Ag <sup>b</sup>	2.1(2)	25(3)
Al	7.0(8)	84(10)

Table 5.2: Lowest heating rate in quanta/second of traps made of superconducting and normal metals measured at cryogenic temperatures. Conversion to electric field noise  $S_E$  is scaled to 1 MHz assuming  $1/f$  scaling [LGA<sup>+</sup>08]. References: <sup>a</sup> [LGL<sup>+</sup>08], <sup>b</sup> [LGA<sup>+</sup>08].

### 5.3.3 Heating rate

The heating rates of Nb and NbN traps are comparable to the lowest heating rates of traps of the same design and tested in the same cryogenic experiment but made with normal metals including Au, Ag, and Al, as listed in Table 5.2. The span of all heating rates measured similarly spans orders of magnitude for any single trap (Figure 5-10).

We measured the heating rate above and below the superconducting transition in the Nbg trap. For the data above  $T_c$ , 3 mA of current is driven to the 1 k $\Omega$  resistor so as to heat the trap just past  $T_c$  as observed by monitoring resistance of the wire structure, corresponding to 9 mW of power dissipated on the trap. In a subsequent cooldown, we mounted RuO<sub>2</sub> temperature sensors on the trap [LGL<sup>+</sup>08] and estimate that the operating temperature in the normal state is  $\sim 2$  K above  $T_c$ . The trap heating rate is measured immediately before and after this change as shown in Figure 5-11. Measurements above and below  $T_c$  are interleaved and taken in quick succession, and they are found to be comparable. All data were taken within one cooldown over two days.

The negligible change in heating rate across  $T_c$  suggests that buried defects have little

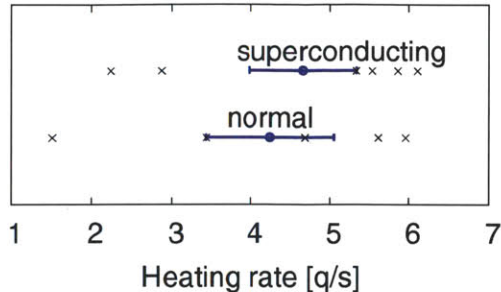


Figure 5-11: Heating rates in trap Nbg-Ia in the normal and superconducting state, with mean and errorbars at 1 standard deviation. Individual data points are shown. The difference between normal and superconducting data is not significant.

effect on anomalous heating. First it is useful to note that at cryogenic temperatures, the expected level of Johnson noise on the trap electrodes is on the order of  $10^{-20}$   $\text{V}^2/\text{m}^2/\text{Hz}$ , while the field noise as measured by the ion is on the order of  $10^{-14}$   $\text{V}^2/\text{m}^2/\text{Hz}$ . Thus it is not surprising that removing the Johnson noise may not have much effect on anomalous heating. The remaining explanation is that anomalous heating is predominantly a surface effect and is unrelated to resistivity. The distinction between surface and bulk is given by the London penetration depth, which in Nb is about an order of magnitude less than the 400 nm film thickness. The results here are still consistent with the current theory of patch potentials on metal surfaces. For superconductors, a recent theory proposed that surface plasmons can be an additional source of electromagnetic noise [HH08].

In traps Nb-I and Nbg-Ia, the heating rate was measured multiple times over the period of over one year. During this time, the trap was installed and removed from the cryostat multiple times and exposed to air in between with no processing or cleaning. The lowest heating rate obtained in Nb-I during any data run shows little variation. However, in both Nb-I and Nbg-Ia, the aggregate of all heating rates shows a slow logarithmic increase over 16 months (Figure 5-12). The cause for this is unclear, though several sources have indicated that Sr deposition during ion loading may be to blame [DNM<sup>+</sup>11, ESL<sup>+</sup>07].

### 5.3.4 Summary

We have demonstrated superconducting ion traps that show good trapping stability and low heating rates. The heating rate does not change appreciably across  $T_c$ , indicating that anomalous heating is primarily a surface effect unrelated to bulk resistivity.

The feasibility of superconducting ion traps invite the possibility of integrating trapped ions with superconducting devices such as Josephson junctions and SQUIDs, though the compatibility of such devices is open to investigation. Recent progress in using the ion as an extremely sensitive detector of forces and charges [MLB<sup>+</sup>09, BUB<sup>+</sup>10, HBHB10] also suggest the possibility of detecting superconducting vortices with trapped ions. Magnetic flux trapped in vortices would modify the magnetic field above the superconductor. The vortex density is determined by the applied external field during cooling across the superconducting transition. The resulting change in local magnetic field can be detected by

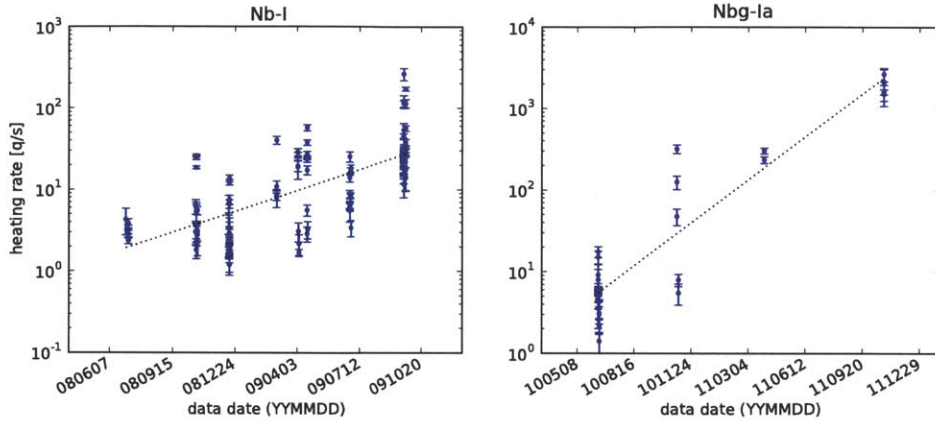


Figure 5-12: Heating rates measured in traps Nb-I and Nbg-Ia over a period of 1.5 years. Dashed lines are fits to  $10^{ax+b}$  where  $x$  is the date in days. The fitted slopes are  $2.6(3) \times 10^{-3}$  q/s/day for Nb-I and  $5.3(4) \times 10^{-3}$  q/s/day for Nbg-Ia.

the ion via the Ramsey method on a narrow transition. An estimate of the ion's sensitivity to magnetic field [MLB<sup>+</sup>09] of  $1.1 \times 10^{-11}$  T/ $\sqrt{\tau/\text{Hz}}$  and typical ion height of 100  $\mu\text{m}$  are comparable to parameters in early experiments that demonstrated vortex detection using SQUIDs [MGC<sup>+</sup>92, MSGW92]. Such coupling to superconducting vortices have been demonstrated with trapped neutral atoms in a recent experiment [MZF<sup>+</sup>10].

## 5.4 Discussion

Much work remains to be done in understanding the source of and reducing anomalous heating in microfabricated ion traps. The variation in heating rates measured in different samples of the same trap material, fabricated under similar conditions and tested in the same experiment, points out the necessity of studying heating rates in a more controlled manner. For example, measurement of distance dependence in the same trap, or modifying the surface composition of a trap in-situ, would remove uncertainties associated with surface contamination and unknown variations in trap fabrication. The data we have obtained for a superconducting trap above and below  $T_c$  is an example of this kind of measurements.



## Chapter 6

# Laser-induced charging

While anomalous heating of ions trapped in microfabricated traps has been studied extensively both by theory and experiment [TKK<sup>+</sup>00, DCG09b, LGL<sup>+</sup>08], laser-induced charging has only seen a few systematic experiments recently. So far, laser-induced charging has been studied on the glass substrate of planar gold traps [DKM<sup>+</sup>08], on copper traps including insulators brought close to the trap surface [HBHB10], and aluminum traps [AHJ<sup>+</sup>11]. Several unknown issues, including material dependence and the role of oxide layers on the metal, remain.

The lasers used for any typical ion trap experiment span a wide range of wavelengths. In a microfabricated trap, they are much closer to the trap surface, and as traps become smaller in size, it is increasingly difficult to avoid scatter caused by lasers illuminating the trap. In some experiments, lasers are deliberately shone onto the trap for the purpose of micromotion compensation [ASS<sup>+</sup>10]. This could be expected to cause buildup of electrical charges on the trap surface due to the photoelectric effect. The typical shortest wavelengths needed for ion traps range from 194 nm for Hg<sup>+</sup> to 493 nm for Ba<sup>+</sup>, corresponding to 6.4–2.5 eV. Typical work functions for metals used for ion traps such as Au, Ag, Al, Cu, etc. are  $\sim 4$  eV or higher, but may change due to surface effects such as the presence of an oxide layer.

The choice of material for ion traps is an important consideration. Gold has been a popular choice due to its chemical inertness, and it has a high work function of greater than 5 eV, but is incompatible with traditional CMOS fabrication. Consequently there has been some interest in using aluminum [LLC<sup>+</sup>09] or copper for microfabricated ion traps, which can take advantage of sophisticated CMOS fabrication techniques. Pure aluminum has a high work function at 4.2 eV and is expected not to release electrons when illuminated with light at 405 nm for Sr<sup>+</sup>. However, aluminum is also known to quickly form a native oxide layer, Al<sub>2</sub>O<sub>3</sub>, which may lower the work function and thus make it susceptible to blue light. Such effects have been observed in previous studies of the photoelectrochemical effects of blue light on aluminum and other materials [Sem69, DCPS91]. Local charges formed on the Al<sub>2</sub>O<sub>3</sub> may not dissipate, changing the trapping potential and leading to excess micromotion [BMB<sup>+</sup>98], which can affect the stability of the trap.

In this chapter, we study the charging behavior of aluminum, copper, and gold microfabricated traps when illuminated with lasers at 674, 460, 405, and 370 nm. All traps are



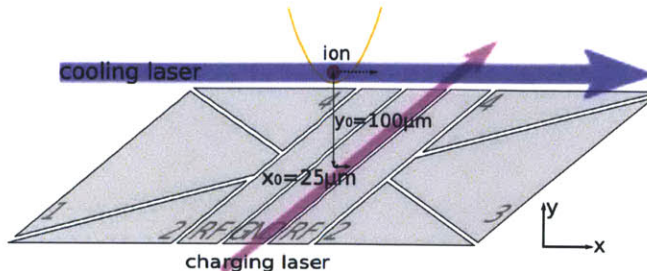


Figure 6-1: Diagram of charging experiment setup (not to scale). The 4 DC compensation electrodes are labeled 1, 2, 3, 4. The charging laser is displaced along the  $x$  axis as shown by  $x_0 = 25 \mu\text{m}$ , such that the ion's displacement has a non-zero projection (dotted line) along the direction of the cooling laser. The axis origin is taken to be the point along the charging laser's waist nearest to the ion (endpoint of the left arrow of  $x_0$ ).

operated in a cryogenic system at 6 K. Charging is measured by observing the micromotion amplitude of a single trapped  $^{88}\text{Sr}^+$  ion and relating it to ion displacement. In the aluminum traps, we find a wavelength dependence of the charging behavior: the laser at 405 nm charges the trap noticeably on timescales of minutes, whereas minimal charging is observed with 460 nm and 674 nm lasers over the same timescales. Copper traps exhibit charging at all wavelengths. No charging is observed at any of these wavelengths for gold traps, but some is observed at 370 nm. A schematic of our charging laser and trap geometry is shown in Figure 6-1.

We describe these experiments beginning in Section 6.1, which covers the physical model, the charging dynamics using rate equations, and the measurement method. Section 6.2 covers trap fabrication and the experimental setup and measurement. Section 6.3 describes the results and presents numerical estimates for the relative charging of different materials and wavelengths. Material in this chapter was published in [WLL<sup>+</sup>11].

## 6.1 Model

We postulate a basic model of the charging process as the photoelectric effect on a metal modified by a thin-film oxide layer, similar to the approach taken in Ref. [HBHB10]. Electron-hole pairs are created near the metal-oxide interface with an initial rate that is proportional to the power of the incident light. As electrons accumulate in the oxide layer, the charging rate decreases due to screening. At the same time, the electrons diffuse at a rate set by the material properties of the oxide layer. We assume that the dissipation of holes in the metal is much faster than the rate of electron diffusion and screening, due to the higher conductivity of the metal.

Based on the simple picture of the photoelectric effect modified by oxides, one would expect that light of lower wavelength and materials with oxide layers or lower conductivity would charge more.

Here we describe the relations between the measured quantities and physical parameters in the model (see Figure 6-2) and the rate-equation model used to fit the time evolution of the charging behavior. Section 6.1.1 defines the micromotion amplitude, relating it to

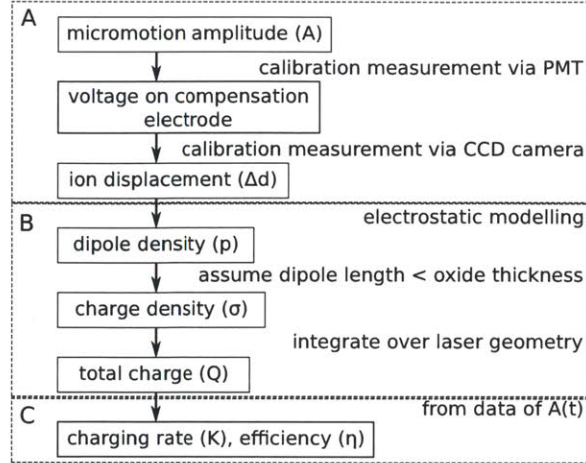


Figure 6-2: Block diagram illustrating the conversion between the measured quantity, micromotion amplitude  $A(t)$ , and the desired quantities, (A) ion displacement  $\Delta d$  and (B) total charge  $Q$ .

the ion displacement and electric field. Section 6.1.2 describes the conversion from ion displacement and electric field to a quantitative estimate of charges on the trap. Section 6.1.3 describes the charging dynamic using a charge accumulation rate, a dissipation rate and screening rate as parameters which leads to a rate equation for fitting to the measured micromotion amplitude vs time.

### 6.1.1 Micromotion and ion displacement

The dynamics of a trapped ion is described in Ref. [BMB<sup>+</sup>98] and the relevant parts are summarized here. The motion of a single trapped ion in a Paul trap with a quadratic pseudopotential is characterized by a low-frequency “secular” oscillation and an oscillation called “micromotion” at the frequency of the applied RF field. For a surface-electrode trap, the trapping potential is slightly modified [Hou08, Wes08], but the nature of the motion (with two characteristic frequencies) is the same. The intrinsic micromotion which occurs when the secular motion carries the ion through the nodal point of the RF field (RF null) is small and will not be of concern in this work. We focus on the “excess” micromotion discussed next.

Assuming the ion is initially located in the RF null such that no micromotion is present, any additional charges generate an electric field which displaces the potential minimum point such that the ion is no longer located in the RF null. With an ion displacement of  $\Delta d$ , the micromotion amplitude is  $\frac{q_i}{2}|\Delta d|$  along the direction of displacement, where  $q_i$  is the Mathieu  $q$  parameter along the same axis. This excess micromotion cannot be significantly reduced by Doppler cooling because it is driven by the RF field [BMB<sup>+</sup>98]. Experiments generally seek to minimize micromotion due to its effect on spectral properties of the ion [BMB<sup>+</sup>98], but here we take advantage of the well-defined temporal behavior of micromotion to discern small displacements in the ion position. This technique is closely related to the Doppler velocimetry technique that has recently been used for ultra-sensitive



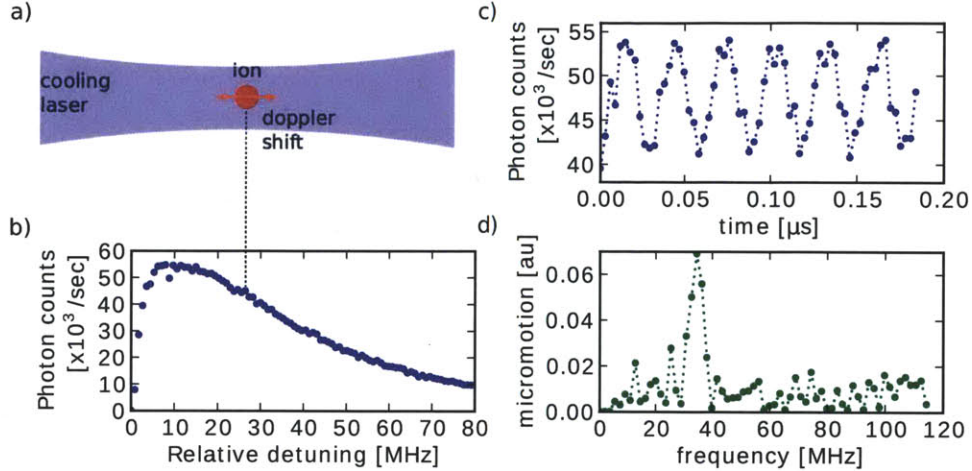


Figure 6-3: Measurement of micromotion signal. (a) Diagram of an ion executing micromotion along the direction of a Doppler cooling laser beam. (b) Typical scan of the fluorescence curve. Dotted line between (a) and (b) indicates the ion’s scattering rate in the absence of micromotion. (c) The oscillating fluorescence signal due to Doppler shift from the micromotion. Photon counts are normalized. (d) Fast Fourier transform of the fluorescence signal normalized to total fluorescence. The maximum value gives the micromotion amplitude.

force detection in Penning traps [BUB<sup>+</sup>10].

Micromotion of the ion is measured using the fluorescence detection method [BMB<sup>+</sup>98] (see Figure 6-3). A photomultiplier tube (PMT) detects fluorescence of the ion, and single photon arrival times are binned to 3 ns bins. This is fast enough to capture the modulation of the fluorescence due to the Doppler shift at near the RF drive frequency of the trap, which is typically between 34-37 MHz. The amplitude of these oscillations,  $A(t)$ , gives a measure of the amplitude of the micromotion along the propagation direction of the cooling laser, and is obtained by performing a fast Fourier transform of the PMT signal. This observed amplitude is proportional to the ion displacement  $\Delta d$ , as verified by two calibration measurements described in detail in Section 6.2.3. The micromotion amplitude is observed to vary linearly with the applied voltage on one of the compensation electrodes (electrode 1, see Figure 6-1), and the ion displacement also varies linearly with this voltage.

### 6.1.2 Ion displacement and charge distribution

In the approximation of a harmonic potential, the ion displacement  $\Delta d$  can be related to the electric field  $\mathbf{E}$  at the ion location generated by the laser-induced charges as  $\Delta d = e\mathbf{E} \cdot \hat{x} / m\omega^2$ , where  $e$  is the ion’s charge,  $m$  is the ion’s mass, and  $\omega$  is the secular frequency along the direction of the ion’s displacement [BMB<sup>+</sup>98]. For simplicity we only consider the ion displacement along  $\hat{x}$ , the radial axis parallel to the trap surface (axis  $x$  in Figure 6-1), so that all the analysis can be done in the one-dimensional model.

The laser-induced charges are located above a conducting surface and thus should be considered to be dipoles, due to the image charge induced in the conductor [HBHB10]. The size of the dipole  $r_d$  in the expression for the dipole moment,  $qr_d$ , is unknown and thus



the number of dipoles created by the laser cannot be easily determined with the techniques described here; however, a rough estimate of the order-of-magnitude of the charge generation rate can be obtained by bounding the dipole size by twice the thickness of the oxide layer. For aluminum, the thickness is taken to be 3-5 nm from the literature [CKN<sup>+</sup>99]. The growth of oxide on copper is not self-limiting as in aluminum and thus its thickness is difficult to estimate; it is assumed that gold has no native oxide layer.

The spatial distribution of the laser-induced charge is taken to be an area of dipoles as follows. The laser intensity distribution on the trap at grazing incidence can be approximated as a line with constant intensity along the trap axis and gaussian distributed intensity profile along the axis perpendicular to it, with waist (radius)  $\omega_0$ . The Rayleigh range of all charging lasers is longer than the length of the trap, so intensity variations along the axial direction can be ignored. We approximate the distribution of charges created on the trap as a Gaussian along the radial direction and constant along the axial direction, directly proportional to the laser intensity. Let  $p$  be the dipole moment density, which is related to the charge density  $\sigma$  as  $p = \sigma r_d$  where  $r_d$  is the size of the dipole. The potential due to such an infinite gaussian line of dipoles is given by:

$$V_{\text{line}}(x_0, y_0) = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} e^{-\frac{2x^2}{\omega_0^2}} \frac{py_0}{2\pi((x-x_0)^2 + y_0^2)\epsilon_0} dx \quad (6.1)$$

where  $x_0$  and  $y_0$  are the horizontal and vertical displacement of the charges from the ion respectively (see Figure 6-1), and  $\epsilon_0$  is the vacuum permittivity constant.

To summarize, from the ion displacement  $\Delta d$  we obtain the electric field and thus the potential created by the laser-induced charges at the ion's location. By assuming a spatial distribution, the potential can be converted to dipole and charge density.

### 6.1.3 Charge accumulation & dissipation

Let  $Q(t)$  be the amount of charge present in the oxide layer generated by a laser incident on the trap as a function of time, with  $Q(0) = 0$ . The charging rate is modeled by two processes. Let  $K$  be the (constant) rate of charge accumulation due to the incident laser. The presence of existing electrons modifies the charging rate over time due to screening, represented by a rate  $-\delta Q$ . Discharging through the oxide can be modeled by  $-\gamma Q$  where  $\gamma$  is a constant set by material properties [DKM<sup>+</sup>08]. Solving the rate equation  $\dot{Q} = K - \delta Q - \gamma Q$  with  $Q(0) = 0$  gives  $Q(t) = \frac{K}{\delta + \gamma}(1 - e^{-\delta t - \gamma t})$ . The time constant for this charging/discharging process is then  $\tau = 1/(\delta + \gamma)$ . At 6 K the conductivity of insulators is expected to be lower than at room temperature, leading to a longer time constant of discharging. Figure 6-4 illustrates the model and rates.

Let  $A(t)$  be the measured micromotion amplitude as a function of time and  $A_\infty$  be the limiting value of  $A$  as  $t \rightarrow \infty$ . The relation between the measured micromotion amplitude  $A(t)$  and charge  $Q(t)$ , therefore  $K$  and  $A_\infty$ , is obtained as follows: micromotion amplitude  $\rightarrow$  voltage on compensation electrode  $\rightarrow$  ion displacement and electric field  $\rightarrow$  dipole density  $\rightarrow$  charge density  $\rightarrow$  total charge (see Figure 6-2). The measured micromotion amplitude

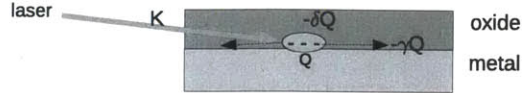


Figure 6-4: Illustration of rate constants in the model of charge accumulation and dissipation.  $K$  is the rate of electron creation,  $-\delta Q$  is the modification to the charging rate due to screening, and  $-\gamma Q$  is the rate of discharging through the oxide.

vs time can thus be written as:

$$A(t) = A_{\infty}(1 - e^{-\frac{t}{\tau}}) \quad (6.2)$$

where  $A_{\infty}$  is the saturated micromotion amplitude (as  $t \rightarrow \infty$ ) and is proportional to the term  $\frac{K}{\delta + \gamma}$ . We use this phenomenological model to fit the experimental data of micromotion vs time and extract the values of  $\Delta d$  at saturation (corresponding to  $A_{\infty}$ ) and the time constant  $\tau$ . Finally we estimate the initial charging rate  $K$  and the charging efficiency  $\eta$ , the latter defined as the number of charges created per photon at  $t = 0$ .

## 6.2 Experiment

The fabrication of the surface-electrode traps used in this work follows the standard optical lithography procedures described in Section 6.2.1. The experimental setup and measurement method are described in Section 6.2.2. Section 6.2.3 describes the calibration measurements to convert the observed micromotion amplitude to ion displacement and electric field.

### 6.2.1 Trap fabrication

We fabricate 1 aluminum, 1 copper, and 2 gold traps for the charging tests described here. We also fabricate 3 aluminum traps with additional deposited layers of oxide in thicknesses of 5, 10, and 20 nm. The traps are of a 5-rod surface-electrode design as described in Section 3.3. All traps are fabricated with optical lithography on 0.5 mm-thick quartz substrates. The aluminum trap with no additional oxide layer is made by first evaporating 1  $\mu\text{m}$  of aluminum on the substrate at a rate of 0.45 nm/s. After lithography using NR9-3000P photoresist, the trap is patterned with wet chemical etch using Transene aluminum etchant type A. No attempt is made to modify the native aluminum oxide formed via contact with air between fabrication and testing. Aluminum traps with extra layers of deposited oxide are fabricated using the lift-off process. After lithography on photoresist, 400 nm of aluminum is evaporated at a rate of 0.33 nm/s, followed by 5, 10, or 20 nm of  $\text{Al}_2\text{O}_3$  at a rate of 0.11 nm/s. After fabrication, the traps are coated with a protective layer of photoresist. Copper and gold traps with electrode thicknesses of 400  $\mu\text{m}$  are fabricated using a very similar lift-off process, except that a 10 nm initial layer of Ti is needed for adhesion during evaporation. Photos of the traps are shown in Figure 6-5.

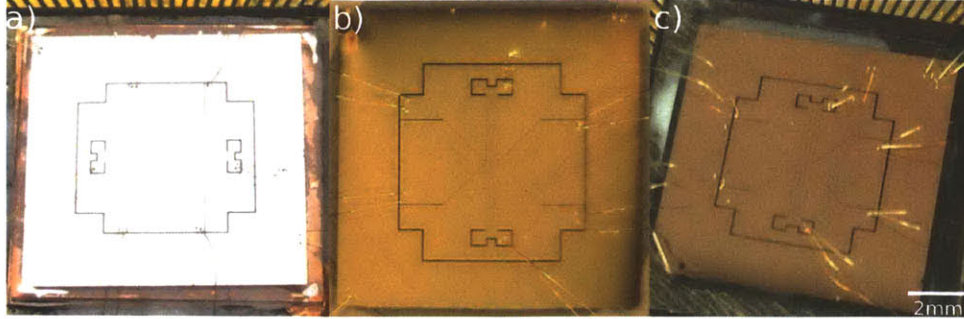


Figure 6-5: Photographs of (a) aluminum, (b) gold, and (c) copper traps.

It is well-known that surfaces exposed to ordinary laboratory environments absorb a few monolayers of hydrocarbon contaminants within a few hours [DPM<sup>+</sup>09, KLM06]. In our experiment, no attempt was made to clean the trap surfaces in situ, so it may be argued that surface contaminants will play a role in the charging effects that we observe. In addition, the exposed dielectrics between electrodes have also been suspected to contribute significantly to charging. To minimize the effects of varying surface preparation, the protective layer of photoresist on all traps is removed only immediately before packaging and installing. The process for packaging and installing in vacuum takes between 12-16 hours.

### 6.2.2 Experimental setup & methods

Lasers at 674, 460, 405, and 370 nm are used for the charging measurements. They propagate along the axial direction of the trap. For the measurements, they are brought to grazing incidence on the trap as confirmed by observing their scatter on the trap surface using the CCD camera. The 370, 405, and 460 nm lasers have a beam diameter of 100  $\mu\text{m}$  and 100  $\mu\text{W}$  of power. The 674 nm laser has a beam diameter of 34  $\mu\text{m}$  and power of 200  $\mu\text{W}$ . Based on the geometry of the experiment we estimate the grazing incidence angle to be no more than 1 degree. From this we can calculate the peak photon flux to be  $\sim 10^{14} \text{ cm}^{-2} \text{ s}^{-1}$ . The lasers are incident on the trap with a horizontal offset of  $x_0 = 25(5) \mu\text{m}$  from directly below the ion, such that there is a discernible displacement of the ion along the radial axis,  $x$ . A schematic of the trap and laser beams is shown in Figure 6-1.

In the absence of deliberate charging by aligning the laser to graze the trap surface, the ion's micromotion signal and compensation voltages are observed to be stable for a long time, on the order of a day. Before each charging measurement, the micromotion of the ion is minimized by applying voltages on the four DC compensation electrodes. Typically the observable micromotion along the radial or axial direction of the trap is sensitive to a 0.01 V change in the compensation voltages. For the measurements described here, we focus only on the micromotion caused by the radial displacement of the ion, parallel to the trap surface.



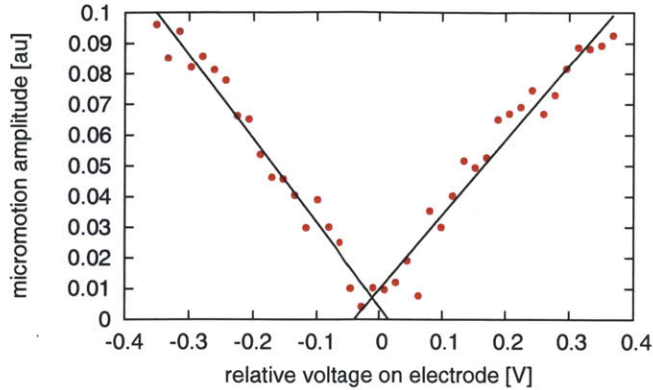


Figure 6-6: Micromotion amplitude vs voltage (relative to optimal compensation) on a compensation electrode, used to convert the measured micromotion amplitude to ion displacement. Lines are linear fits.

### 6.2.3 Calibration

The observed micromotion amplitude is converted to the displacement of the ion and electric-field changes at the ion location via calibration measurements and modeling of the trap potential. We calibrate the micromotion amplitude to the voltage applied to one of the compensation electrodes, 1. The voltage on electrode 1 is scanned and the resulting changes in micromotion is measured as shown in Figure 6-6. Linear fits to this data give the conversion between micromotion amplitude and voltage on the compensation electrode 1,  $c_1 = 0.24(1)$  [au]/V. The ion displacement as a function of voltage is also measured by applying a voltage on the electrode and measuring the ion displacement with the CCD camera. The resolution of the imaging optics is insufficient for measuring the ion displacement directly during charging measurements, so a larger voltage must be applied to obtain this calibration. The ion is displaced  $c_2 = 0.75(3)$   $\mu\text{m}/\text{V}$ . From electrostatic modeling of the trap geometry, for an ion height of 100  $\mu\text{m}$ , the electric-field sensitivity is 50(2) V/m for every 1 V applied to the electrode. From these calibration, the fitting parameter  $A_\infty$  can be converted to ion displacement:  $\Delta d = (1/c_1)c_2A_\infty$ .

## 6.3 Results

In Section 6.3.1 we describe the material and wavelength dependence of the observed charging behavior. Section 6.3.2 describes the measurement to determine the sign of charge, as well as estimates for the initial charging rate  $K$  and efficiency  $\eta$ .

### 6.3.1 Material & wavelength dependence

We tested one trap each of copper, gold, and one each of aluminum with different thicknesses of oxide layers: native oxide, 10, and 20 nm. Aluminum traps, both with and without the additional oxide layer, exhibit charging behavior when the 405 nm laser is incident on the trap. Figure 6-7 shows measured evolutions of micromotion amplitude over time as

the laser is turned on at  $t = 0$  for each of the trap materials. The two parameters that describe the micromotion amplitude over time,  $A_\infty$  (or  $\Delta d$ ) and  $\tau$ , are obtained by fitting Eq. 6.2 to the data. For the aluminum traps and 405 nm, typical time constants are 400-800 s and the saturated micromotion amplitude corresponds to an ion displacement of  $\Delta d = 0.34(3) \mu\text{m}$  at the end of the measurement time, or an electric field at the ion location of  $\sim 20 \text{ V/m}$ . No significant variation of the charging rate or time constant as a function of oxide thickness is observed. The ion displacements measured here are about an order of magnitude smaller than those reported in Ref. [HBHB10], likely due to the much smaller trap size and differences in the laser/trap geometry. The electric field is slightly smaller than that observed in Ref. [AHJ<sup>+</sup>11], where again a different trap geometry and laser wavelengths are used.

In the copper trap, the most pronounced charging effect is observed with the 460 nm laser. Less charging is observed with the 405 and 674 nm lasers. The reversed wavelength dependence of charging in copper, which was also observed in Ref. [HBHB10], is inconsistent with the photoelectric effect hypothesis, suggesting other mechanisms in effect. The charging time constant is typically shorter in copper traps, 100-200 s. The saturated micromotion amplitudes for aluminum is  $\sim 15$  times higher than copper at 405 nm and 5 times higher than copper at 460 nm. In the gold trap, some charging is observed at 370 nm, and not observed for any other wavelengths tested. Fitting to the cases where the micromotion signal appears to stay constant over the measurement time of 1000 s indicates a measurement sensitivity of  $0.01 \mu\text{m}$ .

Comparisons of  $\Delta d$  and  $\tau$  for these materials and wavelengths are listed in Table 6.1. Errors are estimated from repeating the same measurements on different days with the same trap. In some cases such as aluminum at 460 nm and gold at 370 nm, the micromotion signal vs time appears closer to linear, suggesting that the time constant of charging is very long,  $>1000 \text{ s}$ . These data are marked with (\*) in Table 6.1. After blocking the beam, the micromotion amplitude stays constant for at least 20 minutes, suggesting that discharging occurs on a much longer time scale than charging. By comparison, in previous work the discharging time constants were measured to be 654 s for aluminum [AHJ<sup>+</sup>11] and 120 s for copper [HBHB10]. This is consistent with the expectation that the conductivity of the oxide material becomes negligibly small at cryogenic temperatures. Because the rate of discharge is slow, it's possible that the charges created in one experiment continue to contribute a screening effect to the subsequent measurement and thus the measurements taken on the same day are not independent of each other. To minimize such effects, we measured the wavelength dependence starting with the longest wavelength, and in the cases where only the shortest wavelengths exhibited significant charging, the screening effect should be minimal between successive measurements.

### 6.3.2 Quantifying charge

The charging measurements are performed by displacing the laser from directly below the ion by  $25(5) \mu\text{m}$  in order to enhance the detected micromotion signal due to ion movement. For one aluminum trap, the measurement was repeated with the laser displaced on either side of the ion. From the sign of the change in compensation voltages needed to minimize

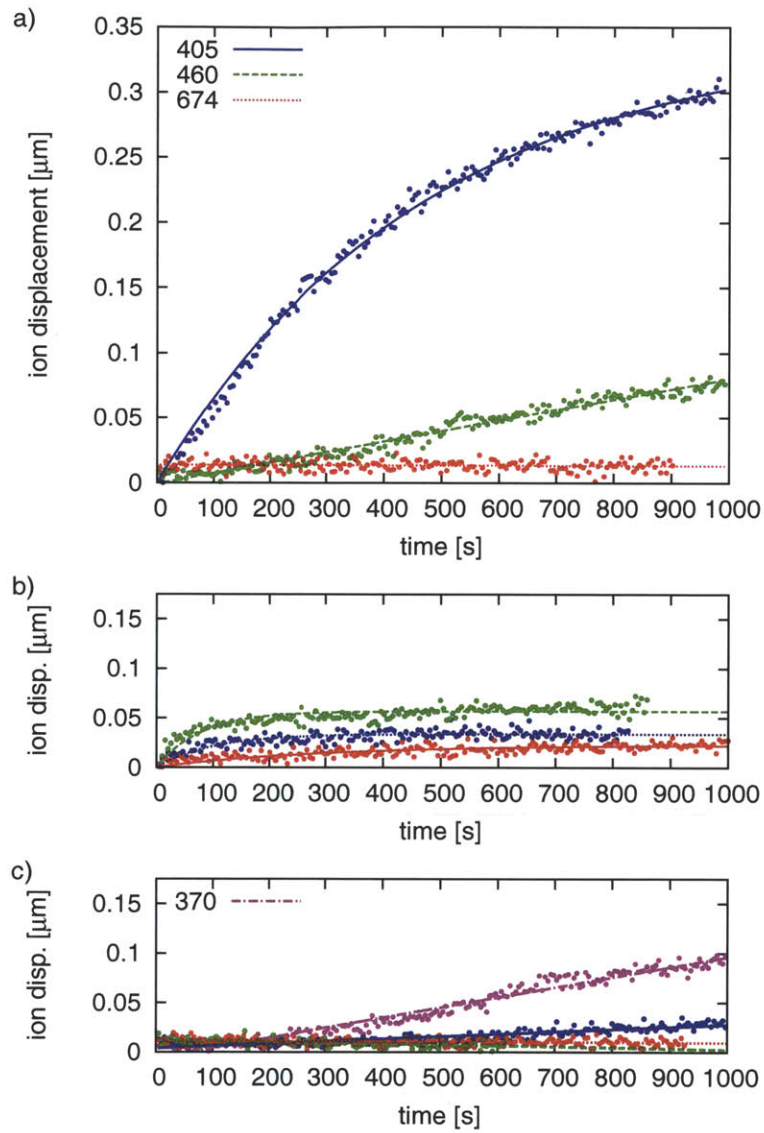


Figure 6-7: Typical plot of ion displacement over time in (a) aluminum, (b) copper, and (c) gold traps showing charging dynamic for all wavelengths: 405 nm (blue, top, solid), 460 nm (green, middle, dashed), 674 nm (red, bottom, dotted), and 370 nm (magenta, dash-dot). Data is smoothed over 5 second intervals. Plots for 405 nm for the aluminum trap and all wavelengths for copper are fit to Eq. 6.2.

Trap	370 nm		405 nm		460 nm		674 nm	
	$\Delta d$	$\tau$	$\Delta d(\mu m)$	$\tau(s)$	$\Delta d$	$\tau$	$\Delta d$	$\tau$
Al	o	o	0.25(5)	500(70)	*	*	-	-
Al-10	o	o	0.28(11)	770(140)	*	*	-	-
Al-20	o	o	0.34(2)	420	*	*	-	-
Cu	o	o	0.02	300(30)	0.05	80(5)	0.034	100(30)
Au	*	*	*	*	-	-	-	-

Table 6.1: Summary of fit parameters from charging data in aluminum, copper, and gold traps. Number after “Al” indicate the thickness of additionally evaporated aluminum oxide. Errors are estimates based on different measurements performed on the same trap. Values without errors indicate that only one measurement was done and errors from the fit are very small, unlikely to represent actual uncertainties. Dashes (-) indicate that the fitted charging rate is consistent with zero, given that the ion displacement at the end of the measurement time is within the measurement resolution of  $0.01 \mu m$ . Asterisks (\*) indicate that fitting to an exponential function resulted in poor constraint on the fitting parameters, and a linear fit is used with different parameters. Circles (o) indicate that the data was not obtained.

micromotion of the ion after charging, one can determine the direction of the movement of the ion and the sign of the charge. We find that the voltage on the electrode closest to the charging laser needs to be increased to re-compensate the ion, indicating that the ion moves toward the electrode and the laser beam due to charging. These observations agree with the hypothesis that the sign of the light-induced charge is negative.

The number of charges created by the laser can be estimated by considering the trap and experiment geometry (Eqs. 1 and 2), assuming a linear relationship between the initial charging rate and the laser power. Such a relationship was observed previously [HBHB10]. Figure 6-8 shows the result of such measurements, but note that the slow rate of discharging in the cryogenic environment means that measurements of charging rate vs power may not be independent. The data shown in Figure 6-8 cannot conclusively rule out either a single-photon or two-photon processes for the charging effect in aluminum. Nevertheless we give an order-of-magnitude estimate for the number of charges produced from the geometry as follows. The center of the gaussian profile of charges is located  $\sim 100 \mu m$  below and  $25(5) \mu m$  to the side of the ion. The size of the dipole  $r_d$  is unknown, but physical estimates of  $r_d \simeq 1$  to  $10$  nm (corresponding to twice the thickness of the oxide layer) results in a dipole density of  $\sim 4 \times 10^6$  dipoles/cm<sup>2</sup> at saturation for the data of the aluminum trap at 405 nm. The initial charging rate as calculated from data fitting is  $K \simeq 1.5 \times 10^4$  charges/sec. The charging efficiency is then estimated to be  $\eta \simeq 10^{-10}$  charges/photon. The charging efficiency for Cu and Au is not calculated since we don’t have a good estimate of the dipole length.

## 6.4 Conclusion

We have observed and characterized effects of laser-induced charging on microfabricated aluminum, gold, and copper ion traps. In aluminum, charging is only clearly observed

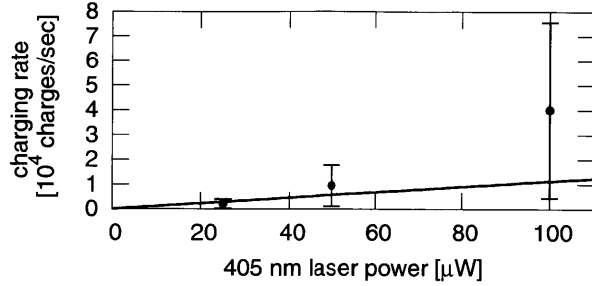


Figure 6-8: Charging rate  $K$  (at  $t = 0$ ) measured in an aluminum trap with 20 nm of oxide on the surface, as a function of incident 405 nm laser power.

for the shortest tested wavelength of 405 nm, suggesting a mechanism dominated by the photoelectric effect. No significant variation is observed for aluminum traps with varying amounts of deposited aluminum oxide. Copper traps exhibit less charging at 405 nm, but some charging is observed at all wavelengths. No charging is observed in gold traps except at 370 nm, consistent with both its higher work function compared to aluminum and copper, and the absence of a native oxide. These measurements suggest that gold may be a preferable material for small-scale ion trap quantum computing.

In surface-electrode traps it is difficult to avoid hitting the trap surface during routine laser alignment when loading ions, but with long ion lifetimes ( $\sim$  few hours in our system) and otherwise stable trapping voltages, the problems with charging may be mostly avoided. However, with lasers at shorter wavelengths (such as those needed for most species other than  $\text{Sr}^+$  currently considered for trapped-ion quantum computing) or smaller ion heights, the charging issue may have greater impact. The timescales of charging observed in our experiments ( $\sim$  100s of seconds) are long compared to most gate operations ( $\sim$   $\mu\text{s}$ -ms), but become relevant in experiments that require many repeated measurements over long periods of time (minutes to hours), such as precision measurements [RHS<sup>+</sup>08] or process tomography [MKH<sup>+</sup>09b, HHR<sup>+</sup>05]. In such cases, care should be taken to detect and correct for changes in micromotion and ion position due to charging.



**Part III**  
**Systems**



# Chapter 7

## TIQC-SPICE

Trapped ions have long satisfied the basic five DiVincenzo criteria for the physical realization of quantum systems. For example, state initialization was demonstrated with fidelities exceeding 99.9% [RCK<sup>+</sup>06]; qubit-specific readout was demonstrated with detection efficiencies of 99.99% [MSW<sup>+</sup>08]; single-qubit gates were demonstrated with error rates of  $2.0 \times 10^{-5}$  [BWC<sup>+</sup>11]; and two-qubit entangling gates were realized with fidelities of 99.3% [BKRB08].

At this point it is worthwhile to identify and examine the key impediments to realizing more complex algorithms with larger systems. Previous chapters have briefly mentioned a number of technical and intrinsic noise sources that limits the fidelities of 1- and 2-qubit gates. For more complex algorithms, the effects of such error sources as a function of algorithm complexity (as measured by, for example, the number of pulses needed to implement it) need to be characterized. Ideally, such information would be used to design pulse sequences that are less sensitive to actual noise sources in the experiment.

In this chapter, we describe a first step toward such a system. TIQC-SPICE (Trapped Ion Quantum Computing – SPICE) is a modeling system for trapped ion quantum computing experiments. A key component of it simulates practical pulse sequences for realizing quantum algorithms by numerically evolving the system Hamiltonian in the presence of various noise sources corresponding to physical errors and technical imperfections in the experiment. Section 7.1 gives an overview of the project, including its relationship with previous simulation programs, and the operator set used in our system. Section 7.2 describes the TIQC simulator in detail, specifically the mapping of experimental noise sources to simulation parameters, and the Monte-Carlo methods used to simulate state evolution. Section 7.3 list some specific examples demonstrating various features and uses of the TIQC simulator. Section 7.4 describes an application of the TIQC simulator for in-circuit gate fidelity evaluation.

### 7.1 System overview

The inspiration for TIQC-SPICE came from SPICE (Simulation Program with Integrated Circuits Emphasis), a software system developed in the 1970s for assisting in the design of integrated circuits [NP73]. SPICE was conceived at a time when increasing size and

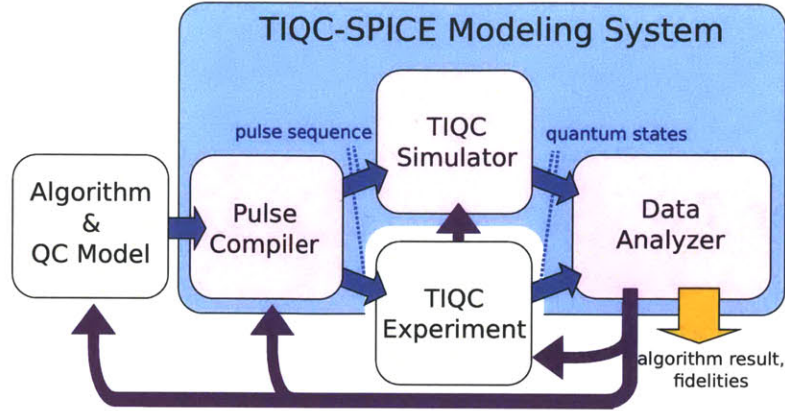


Figure 7-1: Trapped ion quantum computing simulator system overview.

complexity of analog circuits, particularly in integrated circuits, made manual analysis of even simple circuits impractical. Quantum information processing experiments are easily approaching the same level of complexity today.

### 7.1.1 System components

TIQC-SPIICE is a modeling system for assisting in designing experiments and evaluating system performance. A schematic of the parts of the system is shown in Figure 7-1. The input consists of an algorithm, expressed with some assumption about the model of quantum computing to be used (in this case that of trapped ions). This can take the form of a single unitary operator, or a number of such operators with descriptions of how they are interconnected. Next, a pulse compiler converts a quantum algorithm to a specific pulse sequence, given information about a basis of pulses (operators) available in the experiment. This pulse sequence is then sent to the experiment, and simultaneously evaluated in the TIQC simulator. The outputs of both the experiment and the simulator are quantum states during and after the pulse sequence. A data analyzer then processes this data and compares their results with the expected results to obtain desired metrics, including the outcome of the quantum algorithm and their fidelities.

The pulse compiler used in this work was described in Ref. [NHR09]. Its main function is to decompose an  $N$ -qubit unitary operation into a sequence of gates, corresponding to actual pulses in the operational gate set in an experiment. It uses a modified gradient descent algorithm to find a pulse sequence of minimal length which most closely matches the desired  $N$ -qubit unitary operation. A typical pulse sequence obtained this way can be more efficient (shorter) than that obtained analytically by decomposing into 1- and 2-qubit gates [NHR09]. The pulse sequences described in the next chapter are the outputs of this pulse compiler.

The TIQC simulator is useful tool for experimentalists for several reasons. First, it can be used to obtain information that is otherwise difficult or impossible to get directly from an experiment; for example, the effect of an individual error source on the entire sequence. Secondly, it can be used as a predictive tool, where information from a simulation can feed back to design of an experiment such as to minimize the effect of a particular noise source.

Finally, a simulation that matches experimental results is a validation of our understanding of the physical system.

To these ends, our primary goal for a first implementation of the TIQC simulator is to evaluate practical pulse sequences for specific quantum algorithms, predict experimental outcomes with a good degree of accuracy, and identify necessary changes and improvements for an experimental setup. There are two major challenges to this implementation. First, the size of the unitary operator matrix grows exponentially with the size of the quantum system, quickly exceeding the memory capabilities of even today's best classical computers. This poses a practical limit on the size of quantum systems that we can simulate, but in practice even the ability to simulate a small system would be useful and perhaps yield insight on the behavior of larger systems. Secondly, physically there are many sources of noise, and it may be impractical to categorize and simulate all of them. Thus we depend on some experimental insight to identify the most important sources of noise, and validate the choice by comparing experiment and simulation.

### 7.1.2 Prior art

A large number of simulation programs for quantum computing exists. These generally fall under one of two categories. One category of programs defines a gate set based on Pauli matrices, which typically include the single-qubit rotations, CNOTs, and Toffolis, and their operations on quantum state objects. State evolution is computed applying the gate operators on the quantum states, either symbolically or numerically. These programs include, for example, `libquantum` [lib], `eqcs`, `Sympy.quantum`, and many others. Another category of programs provide the tools to define gates, such as Pauli matrices, but leaves the task of specifying a system Hamiltonian representing actual operations to the user. State evolution is computed by numerically evolving the Hamiltonian. These programs include, for example, `qotoobox` [Tan99] and its successor `QuTiP` [JNN11].

The requirements for the TIQC simulator deviates from the goal of most of these programs. The operator set used in our trapped ion experiment is different from the standard gate set consisting of 1- and 2-qubit operations. The programs which defines a complete gate set tend to focus on simulating ideal quantum circuits; in the cases where decoherence mechanisms are defined, they don't include technical sources of error in the experiment. On the other hand, the programs that allow arbitrary definitions of Hamiltonians don't always include the concept of a sequence of gates (equivalent to the system Hamiltonian changing during a single evolution).

The ideal framework for our purpose is a combination of some of these features: ability to specify an operator set, define a sequence of gates, and include a broad range of possible error sources including technical and systematic errors in the actual experiment. The TIQC simulator aims to provide a more directly useful tool for experiments by using physically relevant parameters. The TIQC simulator is so far custom-designed for trapped ion experiments, which limits its broad application, but the principles used in its design should be easily adaptable to other physical systems.

The TIQC simulator is a direct descendant of `quantumcomputer`, a Matlab toolbox to simulate ion trap quantum computers [HRB08]. It defines the Hilbert space and a list of

parameters, and numerically evolves the system Hamiltonian including modifications to it due to a variety of noise sources including laser intensity fluctuations, dephasing, addressing error, etc. We wish to update this framework for several reasons. First, many of the noise sources implemented in `quantumcomputer` are made constant for the duration of a pulse. This is a good approximation, but is not entirely representative of the experiment. In particular, it does not account for correlation functions of the noise, which has recently been characterized [MSB<sup>+</sup>11]. Secondly, it relies on the proprietary software package Matlab, which limits its wide adoption, especially for parallelizing over a cluster (every computer needs a license). We choose to implement the TIQC simulator in Python, which is a full programming language with a large numerical library (Scipy/Numpy) and easy to use syntax.

### 7.1.3 Operator set

Although the framework we have designed should be general enough to be applicable for any ion trap experiment, for the set of test experiments and example algorithms described here, we focus on the <sup>40</sup>Ca linear trap experiment in operation at the University of Innsbruck (see Section 8.2.3).

It was an important result in quantum computing that a two-qubit entangling gate (such as the CNOT gate) plus the set of single-qubit rotations in  $SU(2)$  comprises a universal gate set: any arbitrary  $N$ -qubit operation can be decomposed into such a gate set [BBC<sup>+</sup>95]. A commonly used gate set consists of the set of single qubit rotations generated from the Pauli matrices, plus a two-qubit CNOT gate. However, as discussed in Chapter 4, the 2-qubit CZ CNOT gate for ion traps is more sensitive to experimental errors than the global Mølmer-Sørensen (MS) gate. Similarly, individual qubit rotations in a multi-ion chain also requires precise individual addressing and is sensitive to addressing error.

Therefore we choose an operator set consisting of individual phase rotations  $Z(\theta, \text{ion})$ , global qubit rotations  $R(\theta, \phi)$ , and the global MS gate  $MS(\theta, \phi)$ , as described in Table 7.1. Here,  $\theta$  are the rotation angles and  $\phi$  is the phase, both ranging from 0 to  $2\pi$ . The global qubit rotation is simply the sum of individual qubit rotations, realized by a global laser beam applied to all ions on resonance. The operator is

$$R(\theta, \phi) = \exp\left(-i\frac{\theta}{2}(\cos(\phi)S_x + \sin(\phi)S_y)\right) \quad (7.1)$$

where

$$\begin{aligned} S_x &= \sum_i \sigma_x^i & (7.2) \\ &= \sigma_x \otimes \sigma_0 \otimes \dots \otimes \sigma_0 + \\ &\quad \sigma_0 \otimes \sigma_x \otimes \dots \otimes \sigma_0 + \dots + \\ &\quad \sigma_0 \otimes \sigma_0 \otimes \dots \otimes \sigma_x. \end{aligned}$$

The single-qubit rotation is a phase shift on one ion. It is realized by an addressed,

Gate type	Schematic	Laser path	Operator
Global rotation			$R(\theta, \phi) = \exp\left(-i\frac{\theta}{2}(\cos(\phi)S_x + \sin(\phi)S_y)\right)$
Addressed phase shift			$Z(\theta, j) = \exp\left(-i\frac{\theta}{2}\sigma_z^j\right)$
MS gate			$MS(\theta, \phi) = \exp\left(-i\frac{\theta}{4}(\cos(\phi)S_x + \sin(\phi)S_y)^2\right)$

Table 7.1: Operator set for implementing the quantum algorithms in this chapter.

far-detuned laser, which induces a Stark shift on the ion. The operator is

$$Z(\theta, j) = \exp\left(-i\frac{\theta}{2}\sigma_z^j\right) \quad (7.3)$$

where

$$\sigma_z^j = \sigma_0 \otimes \dots \sigma_z \dots \otimes \sigma_0 \quad (\text{at position } j). \quad (7.4)$$

The global MS gate is realized by a bichromatic laser pulse, detuned by  $\omega_z + \delta$  where  $\omega_z$  is the secular frequency of the trap. Starting from an initial state  $|00\dots 0\rangle$ , the gate generates a global GHZ state in time  $t_{\text{gate}} = 2\pi/\delta$  [BKRB08]. It realizes the operator

$$MS(\theta, \phi) = \exp\left(-i\frac{\theta}{4}(\cos(\phi)S_x + \sin(\phi)S_y)^2\right). \quad (7.5)$$

Two-qubit versions of the global gates can be constructed from the global gates plus single-qubit phase rotations via refocusing, as described in Ref. [NHR09]. However, the process is not necessarily efficient. In particular, the total number of pulses required to generate a 2-qubit MS gate from a global MS gate on  $N$  qubits scale exponentially with  $N$ . Therefore in this framework of single-qubit phase gates and global gates, it is not always constructive to attempt to derive a pulse sequence for an arbitrary operation by reducing it to 1- and 2-qubit gates.

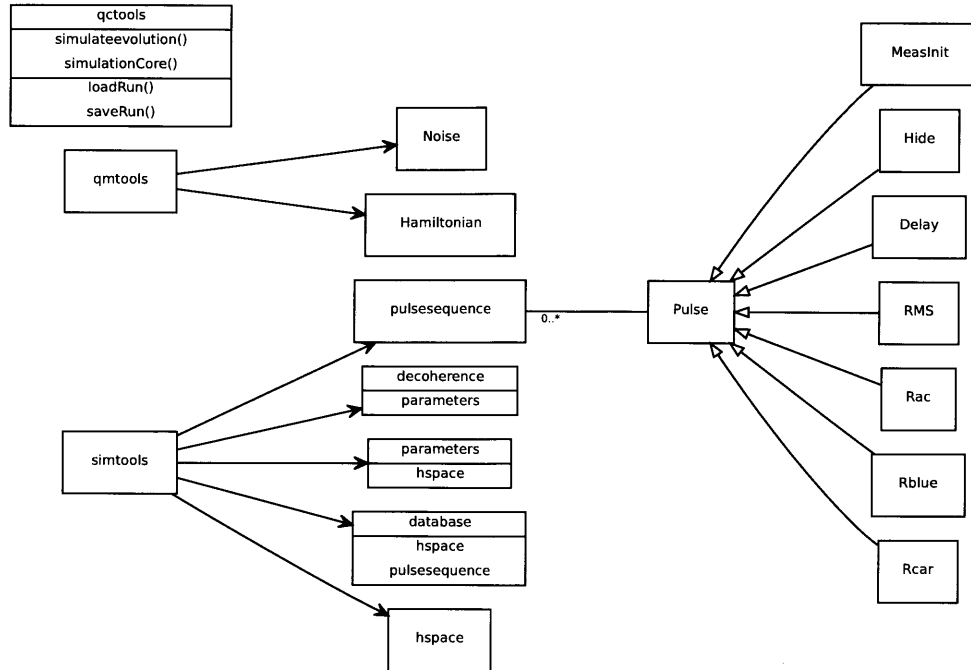


Figure 7-2: Module dependency diagram for the TIQC simulator. The module `qctools` contains the core functions for performing the simulation. `simtools` contain a number of classes used to specific the Hilbert space, parameters, and pulse sequence, and defines the `database` object for storing simulation results. `qmtools` defines the Hamiltonian and Noise classes.

## 7.2 TIQC simulator

This section describes the object-oriented framework with which the TIQC simulator is implemented. We give an overview of the main classes, and describe the Monte-Carlo method used to compute the system state evolution. We describe the implementation of the various decoherence sources and how they relate to physical parameters in the experiment. A module dependency diagram listing the main modules and classes is shown in Figure 7-2.

### 7.2.1 State evolution

The starting point of any simulation of a quantum system is to define the set of variables and operators that completely specifies the system. This is accomplished with a class object `hspace`, with members representing the Hilbert space of the experimental system. A list of its members are shown in Table 7.2. The class also includes a dictionary of operators used to construct the Hamiltonian, as listed in Table 7.3.

The evolution of a state  $|\psi\rangle$  in a quantum system with a Hamiltonian  $H$  is given by the Schrödinger equation (SE):

$$i\hbar \frac{d}{dt} |\psi\rangle = H |\psi\rangle. \quad (7.6)$$

The state is expressed as a state vector with dimension  $D = (n + 1)L^N$ , where  $n$  is the maximum number of phonons in the COM mode,  $L$  is the number of levels (usually 2 or



Member	Description
<b>nuions</b>	number of ions, $N$
<b>levels</b>	levels per ion (2 or 3), $L$
<b>maxphonons</b>	maximum number of phonons, $n$
<b>dimensions</b>	$(n + 1)L^N$

Table 7.2: List of members for the `hspace` class.

Operator	Description
<b>zero</b>	matrix of zero's
<b>proj</b>	projection to the S state (spontaneous decay)
$a$	lowering operator for phonons
$a^\dagger$	raising operator for phonons
$\sigma_m$	lowering operator for qubits
$\sigma_p$	raising operator for qubits
$\sigma_z$	Pauli $z$ operator

Table 7.3: List of operators in the `hspace` operator dictionary.

3) for each ion, and  $N$  is the number of ions.  $H$  is then a  $D \times D$  matrix representing the Hamiltonian.

If  $H$  is time-independent, one standard way to solve this equation is to diagonalize  $H$  and compute the unitary operator  $U = \exp(-iHt/\hbar)$ . In this case we can choose a timestep  $t$  arbitrarily and compute the state with any desired precision in time. In the absence of decoherence, this provides the fastest solution to the final state. Diagonalizing the Hamiltonian, however, is the most time-consuming step in computing the state evolution for a single pulse.

Alternatively, if  $H$  is time-dependent (for example in the case of the bichromatic gate, where one cannot move to a reference frame in which  $H$  is time-independent), the most direct way to solve the SE is to employ an ODE solver and numerically integrate the SE to find the state at each desired timestep. We use the ZVODE solver, a standard ODE solver which uses a method based on backward differentiation formulas with adaptive step size, and is a part of the `scipy.integrate` package derived from the Fortran Lapack library [zvo].

To simulate decoherence mechanisms such as dephasing and spontaneous decay, we use a quantum Monte-Carlo approach [MCD93, DCM92] in which a random trajectory of a particular parameter is generated, and a single simulation realizes an individual evolution with a single trajectory. The ensemble average is obtained by repeating the simulation with different random trajectories and taking the average of the populations and density matrices. More detail of this approach pertaining to specific decoherence mechanisms are described in the next section.

For simulation of open quantum systems with decoherence, another standard method to compute the system evolution is to use the master equation with density matrix formalism. A density matrix  $\rho$  completely specifies an ensemble average of a state. The equivalent of

Member	Description
T	vector of times in the state evolution
Y	state vector vs time
YP	population vector vs time
YtrN	all state populations traced over phonons
TP, YP, YRP	time, states, and populations at the end of each pulse in a pulse sequence
YRPN	YRP traced over motional states
RhoPN	density matrices traced over motional states at the end of each pulse
RhoPNA11	all density matrices for all MC instances. The shape of this matrix is (number of MC instances, number of pulses, <code>hspace.dimensions</code> , <code>hspace.dimensions</code> ).

Table 7.4: List of member variables and functions of the `database` object.

the SE in the density matrix formalism is:

$$\dot{\rho}(t) = -\frac{i}{\hbar}[H, \rho(t)]. \quad (7.7)$$

This is the default system equation used in `qotoolbox` and `QuTiP` [Tan99, JNN11]. However, it is not implemented in the current version of the TIQC simulator due to its poor scaling with system size. A density matrix requires  $D^2$  elements to specify whereas a state only requires  $D$ . For small  $D$ , the density matrix approach is faster since it does not require averaging over many trajectories. But for large  $D$ , the Monte-Carlo approach becomes much more practical.

The module `qctools` provide the functions `simulationCore` and `simulateevolution` which performs the core of the state evolution computations. `simulationCore` represents a single Monte-Carlo (MC) instance. It takes a pulse sequence, a parameters object, and a decoherence object as an input, and outputs a `database` object. `simulateevolution` takes the same inputs, but sets up and calls the multiple MC instances, and processes the multiple `database` objects to return a single data object to the user.

## 7.2.2 Output and analysis

A quantum state is represented by a numpy array of complex values with dimension `hspace.dimensions`. The default initial state of any simulation is the  $|S\dots S, 0\rangle$  state, though this can be user-defined to be any arbitrary state. The result of a simulation is stored in a `database` object. The key components of an instance of this object is a vector `T` of timesteps for which states were computed, and a vector of states `Y` at each timestep. From `Y`, the state populations, populations traced over motional states, density matrices, and other desired views of a simulated evolution are computed. `database` objects are equipped with arithmetic functions `add` and `mean` for summing and averaging multiple MC instances. Key member variables and functions are listed in Table 7.4.

The output of a `simulateevolution` function call is a single `database` object that includes the combined results of all MC instances. For each instance, the populations

Pulse	Description
<code>Rcar(params, <math>\theta</math>, <math>\phi</math>, ion)</code>	carrier pulse
<code>Rblue(params, <math>\theta</math>, <math>\phi</math>, ion)</code>	blue sideband pulse
<code>Rac(params, <math>\theta</math>, ion)</code>	AC stark pulse, approximates a Z-rotation
<code>RMS(params, <math>\theta</math>, <math>\phi</math>)</code>	MS pulse; constructs a bichromatic Hamiltonian
<code>Delay(params, duration)</code>	represents the absence of a laser pulse
<code>Hide(params, hide=True)</code>	instructs certain ion to be hidden/un-hidden
<code>MeasInit(params, ion)</code>	perform measurement and reinitialize one ion

Table 7.5: Pulses defined in the simulation and their mapping to operations.

and density matrices are computed, and these are averaged at the conclusion of all MC instances. Density matrices are only computed at the end of each pulse and not at all timesteps. Density matrices from every instance are stored in the variable `RhoPNAll` while the averaged density matrix is stored in `RhoPN`.

The `database` class also includes a number of functions to plot the evolution of populations over time using the library `matplotlib`, as well as functions to save the data into Python `shelve` file objects. Functions to perform evaluation of fidelities are defined in a separate module, `EvaluateData`. This module also includes functions for performing data analysis on raw data from the lab.

### 7.2.3 Pulse and pulse sequence

Central to the design of our simulator is the representation of a pulse, which realizes individual operators used to construct a quantum algorithm. A `pulse` object represents an operator  $U_{gate}$  (one of  $R$ ,  $Z$ ,  $MS$ ), corresponding to a laser pulse of a fixed duration, frequency, and phase in the experiment. The `pulse` object is an input parameter to the Hamiltonian constructor, and contains all the information necessary to construct a Hamiltonian, along with stored variables for the starting time and ending time of the pulse in a specific pulse sequence. A `pulse` object is specified by a pulse type, rotation angle, phase, and ion. Rotation angle and phase are specified in radians. All other parameters for constructing a Hamiltonian is stored in a global `parameter` object.

A list of `pulses` mapping to the operators in the chosen operator set is listed in Table 7.5. `Rcar`, `Rblue`, `Rac`, and `RMS` represent the physical pulses: carrier, blue sideband pulse, AC stark shift, and MS, respectively. `Delay` represents the absence of a pulse: the resulting Hamiltonian (without noise) is the zero matrix. `Hide` is a special pulse type that represents the combination of pulses, including two carrier pulses on the other Zeeman transitions which are used to hide ions in states that do not couple to the main qubit laser (see Section 8.2.3). In the ideal case, the `Hide` pulse modifies the addressing matrix in the middle of a pulse sequence to achieve this. `MeasInit` is another special pulse type that represents a partial measurement on one ion, storing the result in a classical register, then recooling and re-initializing the ion to the  $|S\rangle$  state. In the simulation, this is implemented by tracing over all other qubits, storing the population of the desired qubit in a classical register, and projecting the qubit to  $|S\rangle$ .

`PulseSequence` is an object that contains a list of pulses along with a set of functions to operate on the list. The member function `makepulsesequence` sets the start and end times of each pulse in the sequence to be consistent with their ordering and durations. This avoids having to declare time in a pulse sequence as a global variable. Another function, `addDelays`, is used to insert delay pulses between each operation pulse to represent the time that it takes in the experiment to switch between a global beam and an addressed beam.

## 7.2.4 Decoherence modeling

Several sources of decoherence are modeled, including dynamic errors (effects that occur during a pulse sequence, such as laser intensity fluctuations, dephasing, etc) and static errors (from technical sources such as addressing error). The static errors are implemented by randomizing the input parameters in each instance of a simulation. The dynamic errors are implemented by either modifying the Hamiltonian, or inserting a (possibly non-unitary) operator at randomly chosen timesteps. A list of error sources currently included are listed in Table 7.6 and are described below.

The dynamic sources of error, including dephasing, intensity fluctuations, spontaneous decay, and heating, are implemented as modifications to the system Hamiltonian by additive and multiplicative terms,  $N_+$  and  $N_*$ , and projection matrix  $N_p$ :

$$\begin{aligned}
 H' &= N_*(t)H + N_+(t) & (7.8) \\
 &\dots \\
 \text{if } N_p(t) \neq 0 : & \\
 &\psi \leftarrow N_p\psi \\
 &\psi \leftarrow \text{normalize}(\psi).
 \end{aligned}$$

The modifications are defined in the `Noise` object, which contains functions to construct these terms.

The static sources of error, including addressing error, initialization error, and spectator mode coupling, are implemented by modifying the runtime parameter (in these cases the addressing matrix) before each simulation instance. The addressing matrix is a member variable of the class `parameters` with  $N$  columns and  $N+1$  rows; each column corresponds to an ion and each row corresponds to the coupling of an addressed laser to that ion. The last row is the coupling of the global beam to each ion. An example of an addressing matrix for 3 ions is:

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}. \quad (7.9)$$

**Dephasing** is an important source of noise for ion traps and correspond directly to the  $T_2$  decoherence commonly discussed in the context of quantum computing. For trapped ions, the physical origin of dephasing are laser phase fluctuations and magnetic field fluctuations. For our current purposes we will not make a distinction between these two physical

Error source	parameters	typical values
addressing error	<code>addressingerr</code>	0.035
	<code>addressingerr_global</code>	0.02
heating rate	<code>heatingrate</code>	141000 $\mu\text{s}$ per quanta
dephasing	<code>coherenceTime</code>	5000 $\mu\text{s}$
	<code>correlationTime</code>	333 $\mu\text{s}$
intensity fluctuations	<code>intensityfluct</code>	0.02
initialization error	<code>stateiniterr</code>	0.003
spectator mode coupling	<code>specmodecoupling</code>	0.02
spontaneous decay	<code>lifetime</code>	1168000 $\mu\text{s}$

Table 7.6: List of error sources, variable names within the class `parameters`, and typical values used for the Innsbruck experiment.

sources, and instead model the dephasing as a phase error which is an extra term on the Hamiltonian:  $N_+ = \phi(t) \sum_{k=1}^N \sigma_z^k$  where  $\sigma_z^k$  is a phase flip on the  $k$ th ion [MSB<sup>+</sup>11].  $\phi(t)$  is a time-sequence of random variables generated from a gaussian distribution with width  $dE$  denoting the strength of the fluctuations. A conversion between the gaussian width  $\sigma$  and the experimentally determined coherence time `coherenceTime` from a Ramsey experiment is necessary. This conversion takes the form  $dE = K/T$  [Lab08, Chw09] where  $K$  is a constant obtained by simulating the Ramsey experiment for various delay times and widths.

A recent experiment on GHZ states [MSB<sup>+</sup>11] postulated a noise model for dephasing that has an exponential correlation in the gaussian width  $\phi(t)$ . In simulation, this is represented by a parameter `correlationTime` and generated from a purely gaussian random sequence by the method described in Ref. [Des02].

**Intensity fluctuations** refer to that of the qubit laser during a pulse sequence. The intensity is proportional to the Rabi frequency. The fluctuation is modeled as a multiplicative term on the Hamiltonian:  $N_* = A(t)$  where  $A(t)$  is a random variable, normally distributed with width `intensityfluct`. The amount of intensity fluctuation is determined by measuring the intensity of the laser beam with a photodiode.

**Spontaneous decay** refers to the lifetime of the  $D$  state of the ion, corresponding to the  $T_1$  time in quantum computing. For example, the lifetime of the  $D_{5/2}$  state of  $^{40}\text{Ca}^+$  is 1.168 s [BDL<sup>+</sup>00]. In the simulation, spontaneous decay is modeled using a non-Hermitian Hamiltonian,  $N_+$ :

$$N_+ = -\frac{i}{2}C^\dagger C \quad (7.10)$$

where  $C$  is a collapse operator, here  $C = \sigma_m$ . After every timestep of the evolution, the state is normalized. In addition, spontaneous decay is simulated by a Poisson process where at every time step, the probability of a spontaneous decay event occurring is equal to the ratio of the time step size to the lifetime. At the beginning of an MC instance, a list of times  $\tau$  where a decay event occurs is generated according to a Poisson process. We evolve the SE as usual until a time  $t$  in  $\tau$ , at which point a projection matrix is applied to the

state followed by a normalization. The projection matrix is

$$N_p = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \otimes I_a \quad (7.11)$$

where  $I_a$  is an identity matrix with the same dimensions as the phonon operators. Then the SE evolution continues until the next  $t$  in  $\tau$ , until the end of the pulse sequence is reached.

**Heating** is the heating rate as described in Chapter 5. This is modeled similarly as the spontaneous decay, but using the raising operator for phonons instead of the lowering operator for qubits:

$$N_p = I \otimes a^\dagger \quad (7.12)$$

where  $I$  is the identity matrix on all ions. The heating rate parameter, expressed in units of  $\mu\text{s}$  per quanta, is obtained by measurements of the type described in Section 5.1.2.

**Addressing error** takes two forms: one is the excitation on a neighbouring ion due to the Gaussian beam of an addressed laser pulse; the other is the non-uniformity of the global beam. The individual addressing error can be measured directly by driving addressed Rabi oscillations on one ion and observing the Rabi frequency on all neighbouring ions. The global addressing error can be measured similarly by driving Rabi oscillations on all ions and measuring the discrepancy on the Rabi frequency between different ions. An example of the modified addressing matrix for 3 ions with global addressing error of 2% and individual addressed error of 3.5% is:

$$\begin{bmatrix} 0 & 0.035 & 1 \\ 0.035 & 1 & 0.035 \\ 1 & 0.035 & 0 \\ 0.98 & 1 & 0.98 \end{bmatrix}. \quad (7.13)$$

**Initialization error** occurs when, during the state preparation phase before a pulse sequence begins, an ion fails to be initialized to the  $|S\rangle$  state. This could be due to imperfect optical pumping. The parameter `stateiniterr` is the probability of this occurring for a single simulation; when such an error occurs and one ion is not initialized to the correct state, the corresponding column in the addressing matrix is set to 0, signifying no coupling to the qubit laser.

**Spectator mode coupling** refers to coupling to motional modes which are not explicitly sideband cooled. The amount of coupling depends on the frequency difference between the desired motional mode and all other modes, and on the thermal occupation of the other modes. This is represented as a randomized intensity shift which is initialized to a different value for each simulation.

## 7.3 Examples

The typical work flow for creating and executing a simulation of a given pulse sequence is as follows:

1. Import the python modules `simtools` as `sim` and `qctools` as `qc`.



2. Create instances of the hilbert space object `sim.hspace` specifying the number of ions/levels/phonons, the parameters object `sim.parameters`, and the decoherence object `sim.decoherence`. All parameters have a default value, which may be directly modified if necessary.
3. Define the pulse sequence object, `sim.PulseSequence`, initialized with a list of pulse objects.
4. Run the simulation by calling `qc.simulateevolution` with the pulse sequence, parameters, and decoherence objects as parameters. Results are returned as a `database` object.
5. Optionally plot the data or save to a file.

Figure 7-3 shows a few sample outputs for various routine experiments. Each of these are described below. The complete code for generating the plots can be found in Appendix C. All frequencies are in units of  $2\pi$  MHz and all times are in units of  $\mu\text{s}$ .

### Rabi oscillations

The simplest example is the simulation of Rabi oscillations on the carrier transition for one ion. After importing the required modules, the `hspace`, `parameters`, and `decoherence` objects are defined:

```
hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)
params = sim.parameters(hspace)
dec = sim.decoherence(params)
```

A pulse sequence consisting of a single pulse is defined:

```
pulseseq = sim.PulseSequence( [
    sim.Rcar(params, 8*pi, 0),
  ] )
```

We run the simulation with the following line:

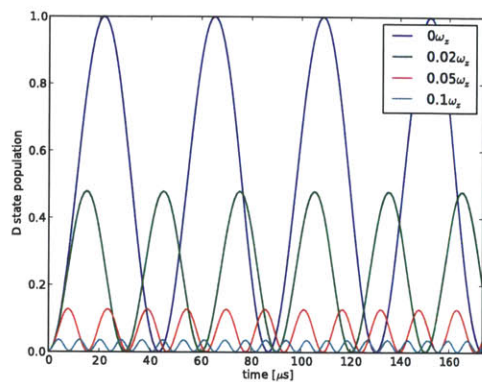
```
data = qc.simulateevolution(pulseseq, params, dec)
```

The D state population can be plotted by first tracing over the phonons, then plotting the time variable (`data.T`) vs the traced population (`data.YtrN[:,0]`).

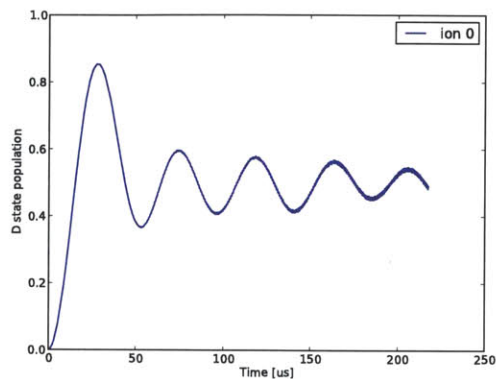
To simulate a carrier pulse that is detuned from the carrier transition, set the detuning parameter for the pulse. Here the pulse can be accessed by indexing the `pulseseq` object. In the following lines, for example, the pulse detuning is set to be  $0.05\times$  the secular frequency `params.omz`.

```
pulseseq[0].detuning = 0.05*params.omz
```

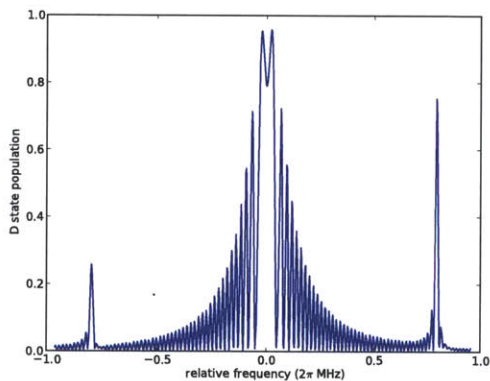
The default initial state is the ground state  $|S,0\rangle$ . Alternatively, the function `initial_state` provides methods to initialize thermal or other quantum states. For example, to initialize a thermal state with  $\bar{n} = 0.5$ :



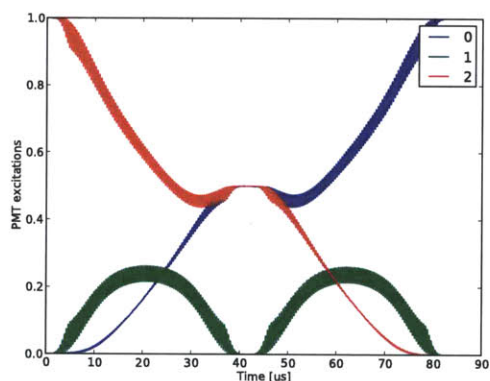
(a) Rabi with detunings



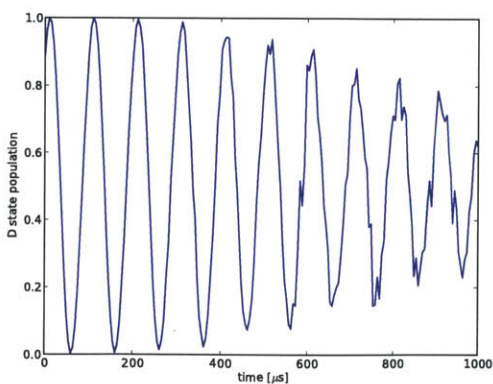
(b) Rabi on thermal state



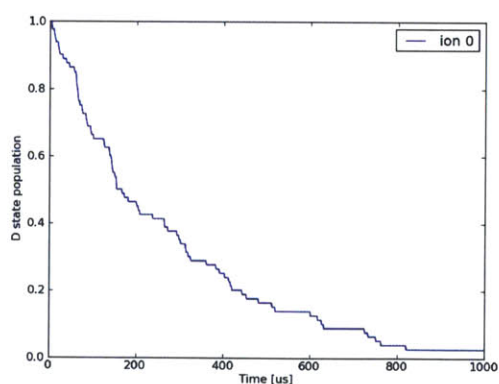
(c) Spectrum



(d) MS gate



(e) Ramsey oscillations



(f) Spontaneous decay

Figure 7-3: Example simulations of standard experiments. The code to generate these plots are in Appendix C.

```
hspace.initial_state('thermal', nBar=0.5)
```

By default, all pulse types except the MS gate uses the method of diagonalizing  $H$  and computing the unitary operator to get the state evolution, while the MS gate uses the ODE solver to evolve the state. An option is provided for all pulses to switch to doing time evolution via the ODE solver, by setting

```
pulseseq[0].dotimedepPulse = True
```

This is mostly useful for verifying that the two methods of solving the SE give the same results. Also, the default timestep size is set to be  $0.1 \mu\text{s}$  for the ODE solver and  $1 \mu\text{s}$  for the method of diagonalizing  $H$ , so greater precision might be gained by switching to the ODE solver at the expense of the runtime of the simulation.

## Spectrum

A common experiment is scanning across the main carrier transition over a specific range to locate the transition frequencies. In the experiment, this typically involves fixing the pulse duration of a single pulse, and measuring the final D state population as a function of the pulse detuning.

To simulate this, we need to evaluate the pulse sequence for each detuning. For each of these evaluations, we obtain a `database` object, but we are only interested in the final D state populations. In this experiment, the operation is a carrier pulse which has a time-independent Hamiltonian that can be diagonalized. Therefore, a larger stepsize (here in units of  $\mu\text{s}$ ) can be selected if no intermediate state information is desired.

```
params.stepsize = 600
```

In addition to being returned by `simulateevolution`, `database` objects can also be directly created and accessed by the user. This allows the user to take advantage of the plotting functions defined in `database` objects. Extra arguments to the constructor indicate that the arrays used to instantiate the object should not be interpreted as states (they are populations in this case) and no particular pulse sequence is defined, so that default functions for computing density matrices and extracting states at each pulse are not invoked:

```
data1 = sim.database(detuning, YR, hspace,  
                    pulseseq=None, statesvalid = False)
```

## MS gate

The simplest simulation of an MS gate follows the same steps as a single carrier pulse, except the pulse sequence consists of the MS pulse:

```
pulseseq = sim.PulseSequence( [  
    sim.RMS(params, pi, 0),  
    ] )
```

In an experiment involving a pulse sequence consisting of several MS gates with possibly different durations, the parameters of the MS gate including the detuning from the sideband, power of the bichromatic beams, and Stark shift correction are optimized for the smallest unit of the MS gate. For example, setting up a pulse sequence consisting of  $MS(\pi/8)$  gates in the experiment involves first running a pulse sequence that consists of 4 such gates concatenated together to make  $MS(\pi/2)$ , and adjusting the parameters to optimize the fidelity of the resulting GHZ state. This is represented in the simulation by the parameter `params.shortestMS`, which is the denominator of the angle for the smallest MS gate, and is set to 2 by default corresponding to an  $MS(\pi/2)$  gate. For example, for a sequence that contains  $MS(\pi/4)$  gates, the parameter is set to 4 before defining the pulse sequence, and the function `calcPulseParams` is called to adjust other parameters accordingly:

```
params.shortestMS = 4
params.calcPulseParams()
```

Populations during an MS pulse is better visualized by the function

```
data.displayPMTpopulations(1)
```

A PMT population measurement gives the total number of bright ions in a chain of ions, but does not identify which of the ions are bright. The plot shows  $n$  lines, each line representing the measurement outcome that  $n = 0..N$  ions in the  $|S\rangle$  state where  $N$  is the total number of ions.

Pulse shaping is crucial in the experimental implementation of the MS gate, due to the small detuning of the bichromatic laser from the sideband frequencies. In the absence of pulse shaping, high frequency components in the sharp edges of a pulse cause residual excitations of the sidebands and is the source of the rapid oscillations seen in Figure 7-3(d). Without pulse shaping, the gate fidelity at the gate time of  $\pi/2$  becomes very sensitive to the exact duration of the pulse [KBZ<sup>+</sup>09]. In the experiment and the simulation, pulse shaping with a duration of 5  $\mu\text{s}$  is enabled by default. Figure 7-3(d) shows the evolution of the state populations as measured with a PMT for two  $MS(\pi/2)$  pulses. Due to pulse shaping, the widths disappear at the beginning and end of each pulse.

## Ramsey

The Ramsey experiment is used to measure the phase coherence time of the system. The pulse sequence consists of two  $R(\pi/2)$  pulses separated by a variable delay time. Similar to the spectrum scan, this requires simulating the pulse sequence multiple times while varying the delay time. In the case of simulating a dynamic source of error such as dephasing, each simulation should be repeated multiple times, corresponding to multiple MC instances, to obtain sufficient statistics to observe the decay of the Ramsey oscillation. The number of MC instances is set by the variable `dec.doRandNtimes`:

```
dec.doRandNtimes = 40
```

Dephasing is enabled by setting the entry in `dec.dict` to True:

```
dec.dict['dephase'] = True
```

In this case it would be advantageous to use parallel processing to speed up the total evaluation time for all the MC instances. Parallel processing, using the external module Parallel Python (`pp`), can be enabled by setting:

```
dec.doPP = True
```

By default, `pp` uses all the available processors on the local computer. Additional servers can be specified by:

```
params.use_servers( ['local', 'server0', 'server1'] ) # etc
```

Each element is a hostname; a text file `server_dict` specifies their translation to IP addresses. On each server besides the local computer, a `pp server python` module must be running to receive and process the job requests. Further information on the functionality of `pp`, including the server module, can be found on its website [pp].

Parallel processing log files and printing of the job statistics, enabled by default, can be suppressed by:

```
params.pplog = False  
dec.doPPprintstats = False
```

## Spontaneous decay

The simulation of spontaneous decay involves initializing the ion in the D state and observing the population vs time for a single Delay pulse. Similar to the Ramsey simulation, we also enable parallel processing and turn on the relevant error source (`dec.dict['spontdecay']`).

### 7.3.1 Performance

Table 7.7 lists some typical runtimes for various experiments. The QFT, order-finding, and Shor pulse sequences are described in more detail in Chapter 8. The Hilbert space dimensions and pulse sequence lengths are given. The runtimes refer to a single MC instance. Runtimes are quoted for simulations run on a Xeon E3-1245 3.30 GHz computer with 16 GB of RAM, except for the Shor simulations which were run on a Xeon X3480 3.07 GHz computer with 16 GB of RAM.

In most cases the `maxphonons` is chosen to be 7, because for the simulation of a single  $MS(\pi/2)$  gate, 7 phonons are required to get 3 digits of accuracy on the resulting GHZ state.

Runtimes shown in the table indicate that turning on the error sources add a significant factor to the total run time. The extra terms on the Hamiltonian (the `Noise` object) are instantiated when any error source is turned on regardless of the error parameter. Because of this, the addition of systematic errors, such as addressing error where only the global parameters are modified, also significantly increase the runtimes due to the extra zero matrices that are still computed and stored. Aspects of this implementation of noise can be optimized for efficiency in the future.

Example	pulse sequence length	hspace. dimensions	runtime (no errors)	runtime (all errors)
MS gate, 2 ions	1	40	2 sec	9 sec
MS gate, 4 ions	1	192	45 sec	4 min
MS gate, 6 ions	1	896	38 min	-
MS gate, 8 ions	1	4096	~36 hrs	-
QFT	18	64	30 sec	3 min
Order-finding, $\pi_2(y)$	24	64	2 min	13 min
Order-finding, $\pi_3(y)$	75	64	3 min	20 min
Shor $a = 11$	52	256	12 min	45 min
Shor $a = 7$	95	256	1 hr	8 hrs

Table 7.7: Run times for simulations with various dimensions and pulse sequence length. Hilbert space dimensions are given by  $(\text{maxphonons}+1)2^{\text{ions}}$  where  $\text{maxphonons} = \text{ions} + 7$  for MS gate simulations and  $\text{maxphonons} = 7$  for all other cases.

## 7.4 In-circuit fidelity

In this section, we apply the TIQC simulator on evaluating gate fidelities in experiments of specific pulse sequences. Often in practical experiments, a full characterization of the implemented gates is desired, and quantum process tomography provide one method of such characterization. However, for  $N$  qubits, this requires making  $12^N$  measurements<sup>1</sup>; this limits the technique to only small systems such that the experiment and data processing can be carried out in a practical amount of time. For 3 qubits, this takes several hours in a typical ion trap experiment. For a 5-qubit gate, the measurement time is for all practical purposes unfeasibly long, and alternative evaluation methods must be used.

Various other methods for obtaining bounds on gate fidelity have been proposed, including parity measurements for GHZ states [MSB<sup>+</sup>11], Hofmann bounds [Hof05], and compressed sensing via measurements on a subset of Pauli operators [FL11]. Most of these methods still involve measurements of isolated gates. Here, we discuss a proposal for a new method of fidelity evaluation based on simulation of noise for gates embedded in an algorithm: in-circuit fidelity (ICF).

### 7.4.1 Definition

Consider a quantum algorithm consisting of multiple quantum operations, represented by gate sequences  $V_i$ , and one particular gate type labelled  $\mathcal{E}$  as shown in Figure 7-4. For example, in an algorithm consisting of a sequence of  $R$ ,  $Z$ , and  $MS$  gates,  $\mathcal{E}$  could be all the  $R(\pi/2)$  gates in the pulse sequence. In an experiment, we measure observables (such as populations or density matrices) between the gates in the algorithm. Let  $E$  be the error between the measured observables and the ideal observables between each gate. Here  $E$  can be a fidelity, for which several possible definitions are possible. Let  $\Delta(F_{\text{sim}}, F_{\text{exp}})$  be the difference between the outcomes of the simulation and the experimental data. Suppose

<sup>1</sup>4 settings to specify the elements of the  $2 \times 2$  matrix for a single qubit, and 3 measurements to project along the  $x, y, z$  axes corresponding to Pauli matrices for the output.

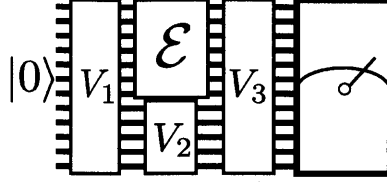


Figure 7-4: A schematic of a quantum algorithm illustrating the general idea of in-circuit fidelity. A gate  $\mathcal{E}$  is embedded in an algorithm operating on a qubit register.  $V_1$ ,  $V_2$ ,  $V_3$  represent sequences of gates before, during, and after the gate  $\mathcal{E}$ .

all gates  $V_i$  have no error except for one gate,  $\mathcal{E}$ . We adjust the amount of error on  $\mathcal{E}$  such that it minimizes  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ , and denote this gate by  $\mathcal{E}_{\text{max}}$ . Then the in-circuit quantum gate fidelity of  $\mathcal{E}$  is the fidelity of  $\mathcal{E}_{\text{max}}$ .

Given an ideal gate (the desired quantum operation)  $U$  and an experimentally realized gate  $\mathcal{E}$ , the gate fidelity is defined as [Nie02]:

$$F = \int d\psi \langle \psi | U^\dagger \mathcal{E}(\psi) U | \psi \rangle \quad (7.14)$$

where the integral is over the uniform (Haar) measure on the state space.  $F = 1$  if and only if  $\mathcal{E}$  perfectly implements the desired gate, and will be less than 1 in the presence of noise. We parametrize the noise on the gate by  $\lambda$ . The physical interpretation of  $\lambda$  depends on the noise source; for example, for dephasing, we can choose  $\lambda = dE$ , the strengths of the phase fluctuations (see Section 7.2.4).

If the algorithm in which the gate  $\mathcal{E}$  is embedded is complex enough, and the gate appears multiple times, then it is likely that each instance of the gate will have different input states. Over the entire algorithm, the gate will efficiently sample a linearly independent combination of all possible input states. Similarly for state tomography on the output state: subsequent gates after  $\mathcal{E}$  will serve as measurements in different bases. This is similar to quantum process tomography in which the input and output states cover the complete input Hilbert space. The advantage is that if the gate does not appear enough times to sample the input or output states, ICF should give a worse bound on the gate fidelity; but in the case of process tomography, all  $12^N$  measurements are required to obtain any gate fidelity at all.

Note that ICF assumes that the algorithm and pulse sequence is in the Markovian regime: errors on each gate is independent of each other. For a sufficiently long pulse sequence (longer than the measured correlation time of the noise), this should be a good assumption. In this case, the gate fidelities estimated in this manner allow us to predict the performance of longer pulse sequences. Consider for example the 3-qubit QFT pulse sequence. Performing an experiment with state tomography at each step takes approximately 30 minutes. Using ICF, we obtain not only the state fidelity at each step but an upper bound on the fidelity of each gate. This allows us to obtain an upper bound on the fidelities of a longer 3-qubit pulse sequence by taking the products of the individual fidelities, assuming that the same gates appear in both. By contrast, performing state tomography on a single 3-qubit gate in the experiment takes >2 hours.

To implement ICF evaluation, in the experiment we require state tomography measure-



ments after each pulse in the pulse sequence. The number of measurements required for state tomography scales as  $3^N$ , which is much preferred to  $12^N$  for process tomography but still scales exponentially with the number of qubits. However, the ICF method can potentially be used given only population information. To generalize the population measurements to density matrices for comparison to the simulation, we make the assumption that the phases in the measured state are correct as compared to the ideal state: the coherence element is then given by  $\rho_{ij} = \sqrt{p_i p_j} \exp(i\phi)$  where  $p_i, p_j$  are the measured experimental populations and  $\phi$  is the ideal phase.

## 7.4.2 Simulation

We evaluate the ICF on various gates for two pulse sequences: the 18-pulse sequence used to implement the 3-qubit Quantum Fourier Transform as described in Section 8.3, and a random permutation of 10 basic operations repeated to form a 40-pulse sequence on 2 qubits.

In the experiment, state tomography is performed after each pulse, as shown in the data in Section 8.3.3. Measurements of population are a subset of the measurements required for state tomography.

The obtained data is compared to simulated data derived from the TIQC simulator. For this ICF evaluation, we focus on errors caused by correlated dephasing, which (apart from addressing error) is the dominant error source for the QFT circuit. The dephasing model is as described in Section 7.2.4, and the error is parametrized by  $\lambda = dE$ .

There are several methods to convert the error parameter  $\lambda$  to the gate fidelity. In the case of global carrier rotations, the actual unitaries (including dephasing error) used in each simulation instance are saved and converted to the process matrix  $\chi$  (see Section 4.3.4) from which the gate fidelity is computed as  $\text{Tr}(\chi_{\text{id}}\chi_{\text{sim}})$ . The conversion can also be obtained by numerically generating the total operator matrix and computing the process fidelity as defined by Eqn. 7.14. For global rotations, these two conversion methods give identical results. For the MS gate, the actual unitaries are not available from the simulation since the Hamiltonian used is time-dependent; therefore only the second method is used.

Performing an ICF evaluation for a particular gate  $\mathcal{E}$  involves only turning on the noise for  $\mathcal{E}$  in the simulation, while all other gates are simulated with no errors. The experimentally measured density matrix before  $\mathcal{E}$  is the input state to the simulation, which begins at  $\mathcal{E}$  and continues to the end of the pulse sequence. The noise in the simulation is set such as to minimize the difference between the outcomes of the simulation and the experimental data,  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ .

There are several possibilities for defining  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ . The goal is to choose a definition which makes the most physical sense and gives the most useful bounds. Table 7.8 lists a number of possibilities. Here  $i$  is the position of  $\mathcal{E}$ ,  $F_g^i$  where  $g = \{\text{sim}, \text{exp}\}$  is the state fidelity of the output of gate  $\mathcal{E}$  in the simulation and experiment,  $n_{\text{end}}$  is the total pulse sequence length, and  $i \leq n \leq n_{\text{end}}$ . For example,  $F_{\text{exp}}^i$  is the experimentally measured state fidelity immediately before  $\mathcal{E}$ , and  $F_{\text{exp}}^{i+1}$  is the measured state fidelity immediately after.

These definitions have varying influences on the resulting gate fidelity bounds. `mean` and `last` are expected to give relatively poor bounds, because a single gate has to account

name	definition	comment
<b>mean</b>	$\sum_{n=i}^{n_{\text{end}}} F_g^i - F_g^n$	sum of errors from $\mathcal{E}$ to the end.
<b>last</b>	$1 - F_g^{n_{\text{end}}}$	final fidelity of the whole sequence.
<b>one</b>	$1 - F_g^{i+1}$	fidelity immediately after $\mathcal{E}$ .
<b>wtail</b>	$\sum_{n=i}^{n_{\text{end}}} (1 - F_g^n) \exp(-i - n_{\text{end}})$	sum of errors starting from $\mathcal{E}$ to the end, weighted by an exponential factor.
<b>wtail_rel</b>	$\sum_{n=i}^{n_{\text{end}}} (F_g^i - F_g^n) \exp(-i - n_{\text{end}})$	same as <b>wtail</b> , but relative to initial state fidelity.

Table 7.8: Possible definitions for the difference between simulated and experimental results,  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ , to be minimized during ICF evaluation.

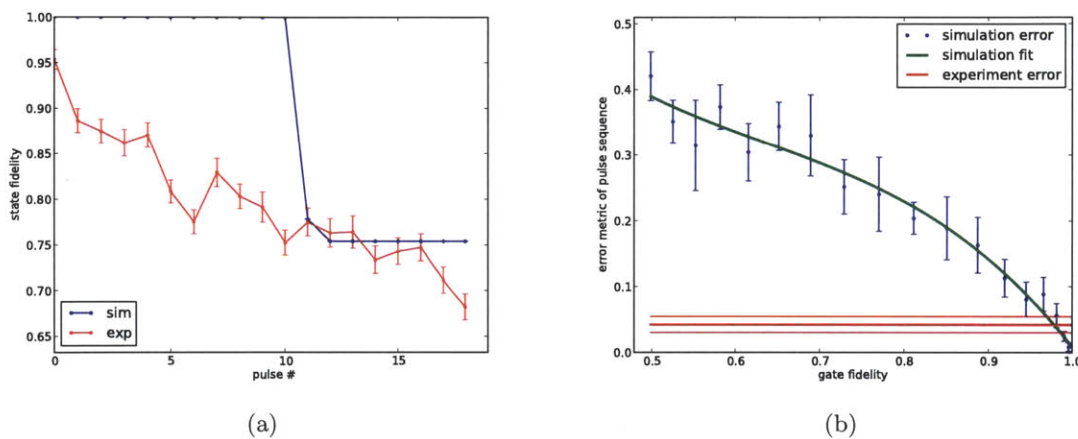


Figure 7-5: (a) Plot of fidelity (compared to ideal state) in the experiment and simulation, where a single gate  $\mathcal{E}$  is subject to noise, while the remaining gates are assumed to be ideal. (b) When the difference between simulated and experimental results is minimized, an upper bound on the fidelity of gate  $\mathcal{E}$  has been determined.

for the infidelity in the remainder of the pulse sequence. **one** should give very good bounds since it's the extension of state tomography on the gate, but there are large variation in the fidelity bounds resulting from experimental randomness. **wtail** is expected to give better bounds than **mean**, and be less sensitive to experimental randomness than **one**. **wtail\_rel** has the same advantages but is normalized with respect to the starting state, and thus makes a good definition to use for evaluating the ICF.

Figure 7-5(a) shows a typical example of the state fidelity vs pulse number as dephasing is turned on for only one gate. Figure 7-5(b) shows how matching the experiment to the simulation (intersection of the red and green lines) gives an estimate for the gate fidelity.

For the typical case, where a gate appears in the middle of a sequence in which the state fidelity falls monotonically vs the pulse number, the above method should provide a fidelity estimate. However, there are two ways that the method may fail, indicated by the lack of an intersection between the simulated and experimental error lines in Figure 7-5(b): (1) the minimum error in the simulation is larger than the experimental error. This happens when the infidelity is dominated by the low fidelity of the experimental input state of gate

sequence	gate	mean	last	one	wtail	wtail_rel
ICF2q	MS( $\pi/2$ )	95(0)	89(1)	95(2)	98(2)	98(2)
ICF2q	R( $\pi/2$ )	91(7)	74(12)	96(7)	97(8)	97(4)
ICF2q	R( $\pi/4$ )	90(5)	n/a	99(5)	95(4)	96(6)
QFT	R( $\pi/2$ )	86(9)	97(2)	95(2)	95(4)	96(1)
QFT	R( $\pi/4$ )	97(3)	96(2)	97(2)	95(5)	99(1)

Table 7.9: Gate fidelity bounds obtained with ICF for 3 types of gates in the 40-pulse 2-qubit random pulse sequence (“ICF2q”) and the 18-pulse QFT sequence for various definitions of  $\Delta(F_{\text{sim}}, F_{\text{exp}})$  as listed in Table 7.8. Errorbars are averaged over the upper and lower deviations.

$\mathcal{E}$  compared to the ideal input state, instead of by the error during  $\mathcal{E}$ . (2) The maximum error in the simulation is smaller than the experimental error. This means the experimental error is too large, or equivalently, that error during gate  $\mathcal{E}$  is not enough to account for the measured error and the bound is very low.

### 7.4.3 Results

Following the fidelity estimation simulation discussed in the previous section, the fidelity bounds are evaluated for global carrier rotation and MS gates. Table 7.9 shows the ICF bounds obtained for gates in the 40-pulse 2-qubit sequence, for each of the definitions of  $\Delta(F_{\text{sim}}, F_{\text{exp}})$  as listed in Table 7.8. Each of the displayed gates appear 8 times in the pulse sequence (4 times each with phases  $x$  and  $y$ ). The values are averaged over the 8 instances, ignoring cases where the ICF fails. These values are generally consistent with the expectations of fidelity bounds for each  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ . The table also shows the ICF bounds obtained for two global rotation gates in the QFT sequence. The  $R(\pi/2)$  gate appears twice in the sequence and the  $R(\pi/4)$  gate appears once.

While it would be desirable to compare the outputs of ICF to outputs of full process tomography, in this particular experiment the process tomography infidelities are dominated by errors during the state initialization, state measurement, and projection noise. The estimated combined error of state initialization and measurement is 1%, and projection noise contributes 2-3% to the error [Mon11]. Thus a direct measurement of gate fidelity via full process tomography in the experiment can never exceed 99% with 2-3% errorbars. Instead, based on measurements of states resulting from repeated applications of these gates, it is estimated that the fidelity of the 2- and 3-qubit gates being evaluated here should be around 99%.

Table 7.10 shows a comparison between the results of the ICF evaluation based on full density matrices and based on populations only, using `wtail_rel` for  $\Delta(F_{\text{sim}}, F_{\text{exp}})$ . The ICF bounds resulting from measuring population only are comparable to those obtained by full state tomography. This speeds up the evaluation significantly, since much fewer measurements are required to obtain direct population information in the experiments.

sequence/gate	ICF ( $\rho$ )	ICF (p)
ICF2q: $R(\pi/2)$	97(4)	95(6)
ICF2q: $MS(\pi/2)$	98(2)	96(6)
QFT: $R(\pi/2)$	96(1)	98(1)
QFT: $R(\pi/4)$	99(1)	99(1)

Table 7.10: Fidelity bounds on individual gates obtained via ICF from density matrix ( $\rho$ ) and population (p) measurements.

## 7.5 Summary and future work

In this chapter, we presented a modeling system for designing practical pulse sequences for realizing quantum algorithms in a trapped ion quantum computer. The TIQC simulator computes the state evolution given a pulse sequence and an initial state, taking into account physical and technical noise sources present in an actual experiment. Some simple examples were provided to demonstrate the application of the simulator for typical ion trap experiments. As a further example, we showed how to use the simulator to provide estimates on gate fidelities given only partial state information from the experiment.

One major objection to classically simulating quantum computing systems is obvious and is one of the motivations for building quantum information processors: the size of the Hilbert space scales exponentially, so in practice the size of the system amenable to classical simulations is severely restricted. The TIQC simulator does not overcome this challenge: already the simulation of a 5-qubit,  $< 100$ -pulse sequence takes on the order of several hours. But there is still a lot to be gained in the optimizations of the classical simulation itself. A few ideas for improving the speed is listed below.

- errors can be approximated as a perturbation on the Hamiltonian:  $\exp(i(H + H_\epsilon)t)$  may be computed faster when given  $\exp(iHt)$ .
- choose a faster ODE solver for time-dependent Hamiltonians. Currently the numerical ODE solving for MS gates takes the majority of the time in simulating a typical pulse sequence.
- use sparse matrices, or save the Hamiltonian matrix between each timestep.
- dynamically reshape the Hilbert space: remove the phonon states when they are not needed.
- parallelize the computation at the matrix multiplication level rather than the Monte Carlo instance level.

The recent demonstration of the classical simulation of 42 qubits on a supercomputer cluster [RMR<sup>+</sup>07] shows that much more classical computing resources can be exploited than what was done in this work.

Also in this chapter we have mostly discussed the TIQC simulator in isolation, and taken the pulse compiler to be a separate and external system component. In fact, the issue of compiling a quantum algorithm into a pulse sequence (essentially, “machine code” specific to a particular physical quantum computing implementation) is an interesting one and there is much ongoing work to investigate the various possibilities. From a practical point of

view, a logical step is to combine the pulse compiler and the simulator, by feeding back the outputs of the simulator to the pulse compiler. This would enable the pulse compiler to generate pulse sequences that are optimized based on the noise sources present in a specific experiment, rather than relying only on a generic optimization strategy such as shortening the total pulse sequence length.

In the next chapter, we use the TIQC simulator in conjunction with experimental work to demonstrate a number of key quantum algorithms using trapped ions. These remain small systems which are amenable to classical simulation. Of course, eventually the quantum devices we can build will exceed our ability to classically simulate them. But just as the simulations of the next generation classical processors often demand the best available computational power of the current generation, the development of future quantum processors will likely require the best of their predecessors. Today, one cannot build a classical computer without a classical computer, and perhaps this will be true for quantum computers as well.

## Chapter 8

# Quantum Algorithms

The basic operations required by the DiVincenzo criteria have all been achieved with trapped ions with above 99% fidelity in recent years. With these fidelities, the realization of actual quantum algorithms seem to be within reach; indeed, the Deutsch-Jozsa algorithm was demonstrated on a single ion [GRL<sup>+</sup>03], and the Grover search algorithm was demonstrated on 2 ions [BHL<sup>+</sup>05]. More recently, universal two-qubit gates were realized with 2 ions [HHJ<sup>+</sup>09a, HHJ<sup>+</sup>09b], quantum error correction with 3 ions [CLS<sup>+</sup>04, SBM<sup>+</sup>11], and universal quantum simulation with 5 ions [LHN<sup>+</sup>11]. Other significant recent achievements include quantum teleportation [OMM<sup>+</sup>09] and a 14-ion entangled state [MSB<sup>+</sup>11].

In this chapter, we describe the experimental realizations of Shor’s algorithm [Sho94], one of the most well-known quantum algorithms with an exponential speedup over its classical counterpart. This requires good classical and quantum control, and we work toward this goal by first realizing the quantum Fourier transform and the order-finding algorithm with trapped ions. We also use TIQC-SPICE, described in Chapter 7, for simulating these experiments including noises sources to compare the validity of our noise model with the experimental data.

This chapter is organized as follows. Section 8.1 gives a general overview of Shor’s algorithm and its relationship with order-finding and quantum Fourier transform. Section 8.2 discusses prior implementations of quantum algorithms in other physical systems, highlights some unique properties of the trapped ion system, and describes the experimental system used to realize the quantum algorithms described in this chapter. Section 8.3 describes the Quantum Fourier Transform (QFT), including the experimental implementation and comparison with simulation results. Both the fully quantum version using 3 ions and a semiclassical version using only 1 ion are implemented. Section 8.4 covers the same for the order-finding algorithm on 3 ions. Section 8.5 describes the implementation Shor’s algorithm for factoring  $N = 15$  using 5 ions. The demonstration of Shor’s algorithm represents one of the most complex quantum algorithms realized in trapped ion quantum computing to date.

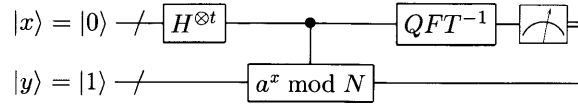


Figure 8-1: General quantum circuit for Shor’s algorithm.

## 8.1 Motivation

Shor’s algorithm was one of the earliest quantum algorithms discovered that gives an exponential speedup over all currently known classical algorithms. The algorithm finds prime factors  $p, q$  of an integer  $N = p \times q$ . This is a hard problem on classical computers, as the number of computational steps scale exponentially with the number of bits in  $N$ . Shor’s algorithm has implications for practical applications: the difficulty of factoring on a classical computer underlies the security of the RSA public-key encryption system, and the quantum algorithm can be used to break it.

Shor’s algorithm is a good candidate for experimental demonstrations for several reasons. First, it solves a problem of practical interest for which classical algorithms are also available, and has been extensively studied. Second, its demonstration requires good classical and quantum control; entanglement and coherence are crucial to algorithm success. Finally, it includes the quantum Fourier transform, a key component that is responsible for the speedup of many other quantum algorithms over their classical counterparts.

Given a number,  $N$ , the steps to factor it using Shor’s algorithm proceeds as follows.

1. classical pre-processing
  - (a) choose  $a < N$  where  $a$  is not already a factor
  - (b) construct unitary operators for controlled-  $a^x \bmod N$  for  $x = 2^0, 2^1, 2^2 \dots 2^{\log_2(N)}$
2. quantum part
  - (a) prepare the input state  $|x\rangle |y\rangle = |00\dots 1\rangle$
  - (b) perform Hadamard gate on register  $|x\rangle$  to put it in a superposition state
  - (c) perform controlled modular exponentiations  $a^x y \bmod N$
  - (d) do an inverse quantum Fourier transform on register  $|x\rangle$  and find the period  $r$  of  $a \bmod N$  from continued fraction expansion
3. classical post-processing:
  - (a) compute  $\gcd(a^{r/2} \pm 1, N)$ .

The classical pre- and post-processing steps can all be carried out in polynomial time on classical computers<sup>1</sup>. The quantum computer is used to speed up the sub-task that otherwise takes time exponential in the number of bits of  $N$ . The general quantum circuit is shown in Figure 8-1

Shor’s algorithm is a special case of a more general problem called order-finding. The order-finding problem is defined as follows. Let  $M$  represent a set of  $2^m$  discrete elements, and  $\pi$  represent a permutation on an element  $y$  of  $M$ . The order,  $r$ , of a permutation  $\pi$

<sup>1</sup>Constructing the operators for controlled-  $a^x \bmod N$  can be done in various ways; decomposing it into multiplications is described in Ref. [NC00]. Ref. [MS12] gives an explicit construction into quantum gates. The gcd can be efficiently computed using Euler’s algorithm.



of  $y$  is defined as the minimum number of applications of  $\pi$  such that  $y = \pi^r(y)$ . Shor’s algorithm reduces to finding the order of the permutation  $\pi(y) = ay \bmod N$  where  $y = 1$ . The quantum algorithm for order-finding has a provable exponential speedup over classical algorithms [Cle00].

## 8.2 Implementations

Many of the current leading physical implementations of quantum computing has successfully demonstrated small instances of the algorithms described in this chapter. In this section, we summarize some of these prior art and highlight some important and unique advantages of trapped ions.

### 8.2.1 Prior art

NMR was one of the first physical systems to demonstrate the QFT [WPF<sup>+</sup>01], order-finding [VSB<sup>+</sup>00], and Shor’s algorithms [VSB<sup>+</sup>01]. The QFT was done on 3 qubits, the order-finding on 5 qubits, and Shor’s on  $N = 15$  with 7 qubits. In particular, the implementation of Shor’s algorithm used  $\sim 300$  pulses, many of which are refocusing pulses to cancel unwanted couplings between qubits. While NMR was instrumental in showing that practical realizations of such algorithms are possible, NMR has several disadvantages which pose difficulties in scaling. For example, each of these algorithms required the synthesis of a new molecule in order to obtain the desired number of qubits with favourable frequencies. Refocusing pulses are required because the qubits are always coupled in a molecule, and as the number of qubits in the molecule grow, the requirement for refocusing also grows rapidly. Finally, high-fidelity measurements on qubits require them to decohere, and therefore one cannot use intermediate measurement results to control further operations, which is a key element of quantum error correction.

Photons have followed NMR’s lead in demonstrating Shor’s algorithm in a number of experiments [LWL<sup>+</sup>07, PMO09, LLL<sup>+</sup>11]. Ref. [LWL<sup>+</sup>07] performed factoring using 3 qubits on  $N = 15$  with 13 gates, while Ref. [LLL<sup>+</sup>11] factored  $N = 21$  using 4 qubits encoded in the higher order states of 2 photons, using 26 gates. Photonic qubits overcome some of the challenges of NMR qubits, but the main drawback remains that conditional branching in an algorithm is not straightforward to implement with photonic circuits. Instead, Ref. [LLL<sup>+</sup>11] used post-selection, in which all possible subsequent operations resulting from all possible outcomes of intermediate are performed, and the correct results are selected from all the measurements via post-processing. Post-selection is disadvantageous for scaling because the number of experiments needed scales exponentially with the number of branching points in an algorithm.

Superconducting qubits have only recently resolved the problem of short coherence times, and are making rapid progress in demonstrating gates and other key quantum computing elements with high fidelity [CGT<sup>+</sup>09]. A recent demonstration of Shor’s algorithm for factoring  $N = 15$ , albeit “compiled” to 3 qubits and using only 2 CNOT gates, was performed [LBY<sup>+</sup>11]. Superconducting qubits benefit from well-developed fabrication techniques and existing capabilities in microwave electronics. However, many implementation

details required for demonstrating complete algorithms remain to be resolved. For example, feedforwarding requires measurement times to be comparable to gate times, which for superconducting qubits are very short (10-100 ns). Characterization of process fidelities for longer pulse sequences also needs to be done to determine the limitations for realizing complex algorithms.

### 8.2.2 TIQC implementation

Trapped ions are well-equipped for realizing complex quantum algorithms at the present time, not only benefitting from a longer history of well-characterized gates and small algorithms, but also from having a few capabilities not found in some of the other physical systems.

Unlike in NMR, ion qubits in a trap are not coupled by default and therefore do not require refocusing pulses to undo unwanted coupling. However, our choice of an operator set that includes only global and addressed gates necessitates the usage of refocusing pulses to realize 2-qubit operations. Like in NMR, this adds significant overhead to the number of pulses required in implementing an algorithm. However, instead of using refocusing to remove specific qubits from global operations, ion qubits can be “hidden”, or stored in auxiliary atomic states outside of the computational space (see Figure 8-2). This ability greatly simplifies the mapping of complex circuits onto a qubit register.

Trapped ion experiments are also capable of fast measurements, which take only slightly longer than typical gate times (a few 100  $\mu\text{s}$  vs  $\sim 10 \mu\text{s}$ ), but are still well within the coherence time. Combined with hiding, measurements can be performed on a subset of qubits while leaving the hidden qubits intact. Ions can also be initialized quickly: Raman cooling [Pos10] can cool the motional state to  $\sim 1$ -2 quanta in  $\sim 400 \mu\text{s}$ . These two features combined allows a feed-forward circuit design, in which an intermediate measurement can be used to control subsequent gates in the rest of the quantum circuit.

The demonstrations of algorithms described in the rest of this chapter use these capabilities to fully explore the potentials of the trapped ion system for realizing quantum algorithms.

### 8.2.3 Experimental setup

The experiments described in this chapter were all performed on the linear trap experiment operating at the University of Innsbruck. The experiment consists of a 3D Paul trap and uses  $^{40}\text{Ca}$  ions. The  $^{40}\text{Ca}$  ion has a similar atomic structure as the  $^{88}\text{Sr}$  ion, and its qubit transition is at 729 nm with a lifetime of  $\sim 1$  second (see Figure 8-2). Detailed descriptions of the experiment apparatus, the laser system, and the control software can be found in Refs [Gul03, Rie05, Chw09, Sch08].

Table 8.1 lists some capabilities of the Innsbruck experiment, including qubit preparation, measurement, coherence times, and gate times and fidelities.

The algorithms described in this chapter uses the additional abilities of ion trap experiments: hiding ions in auxiliary states, performing a measurement in the middle of the sequence and use the result to control the presence of subsequent pulses (feedforwarding),

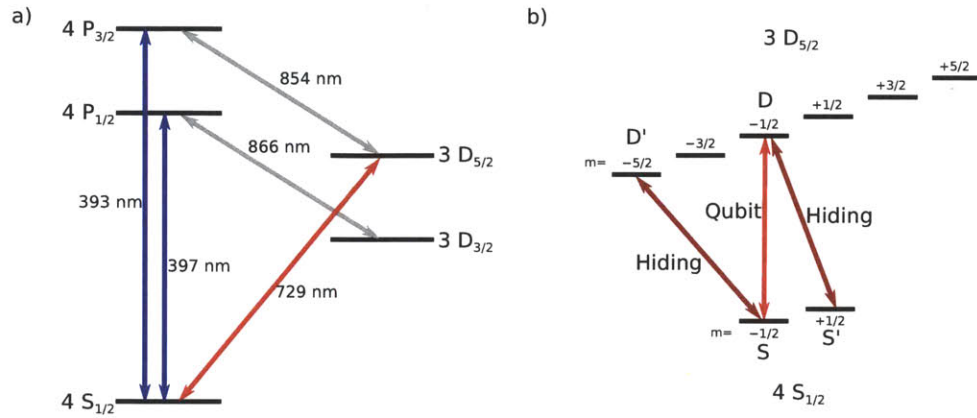


Figure 8-2: (a) Atomic structure for the  $^{40}\text{Ca}$  ion. (b) Zeeman sublevels of the S-D manifold showing the qubit transition and other transitions used for ion hiding.

ions	$^{40}\text{Ca}$ , up to 14	
lasers	1 qubit, 2 cooling, 2 repump, 2 PI	
qubit preparation	sideband cool	$\sim 2$ ms, 99.5%
	Raman cool	$\sim 400$ $\mu\text{s}$ , to $\bar{n} \sim 1$
coherence time	T1	1.1 s
	T2	5-50 ms
measurement time	PMT	400 $\mu\text{s}$
	CCD	8 ms
gate time	$R(\pi, \phi)$	20 $\mu\text{s}$
	$Z(\pi)$	20 $\mu\text{s}$
	$MS(\pi/2, \phi)$	45 $\mu\text{s}$
	gate fidelity	$R(\pi, \phi)$
$Z(\pi)$		>99%
$MS(\pi/2, \phi)$		99% (2 ions), 94% (5 ions)
hiding	50 $\mu\text{s}$ , 99.5%	

Table 8.1: Capabilities of the ion trap experiment at the University of Innsbruck, including durations and fidelities of various operations.

and reinitializing an ion to the default ground state in the middle of the pulse (recooling).

Hiding of ions is implemented by detuning the global qubit laser to one of the other carrier transitions ( $S' \rightarrow D$  or  $S \rightarrow D'$ ). For performing a partial measurement, only hiding the  $S$  state into the  $D'$  state is necessary. For decoupling ions completely from a set of operations, hiding both the  $S$  and  $D$  states are necessary.

Feedforwarding is implemented in the experiment via electronics hardware: a 4-bit counter counts the number of photons accumulated during a set detection time by the PMT. The resulting number of counts is compared to a threshold and the resulting binary signal is stored in a flip-flop. The current hardware can store the result of two such measurements in a given sequence. The binary signal is used to control the RF output of a subsequent pulse. The latency is determined by the propagation time through a few logic gates and should be no more than a few tens of ns. Details of this feedforwarding hardware is described in the Appendix of Ref. [Rie05].

Recooling can in principle be done by sideband cooling, but this has the disadvantage that it takes a few ms to cool to the ground state, which is significant amount of time compared to the entire pulse sequence length. An alternative cooling technique is Raman cooling: two lasers detuned from the  $S$ - $P$  and  $S'$ - $P$  transitions drive the transition  $|S, n\rangle \leftrightarrow |S', n+1\rangle$ , and a third laser resonant with the  $S$ - $P$  transition preferentially scatters photons without adding a phonon. Experimentally all three beams are turned on for  $\sim 200 \mu\text{s}$  and cools the ion to 1-2 quanta. The Raman cooling apparatus is similar to that described in Section 4.6 of Ref. [Pos10].

## 8.3 Quantum Fourier transform

The quantum Fourier transform (QFT) is an important component of order-finding, which is a generalization of the Shor's algorithm. The QFT is a quantum analog of the DFT which is employed in the order-finding algorithm as a method to determine the period of the amplitudes in a superposition of quantum states. The QFT is responsible for many of the known instances of exponential speedups over classical algorithms [NC00, EJ96]. The QFT has been implemented with NMR qubits [WPF<sup>+</sup>01] and trapped ions [CBL<sup>+</sup>05], in both cases with 3 qubits.

### 8.3.1 Algorithm

The QFT performs the same transformation as the classical discrete Fourier transform (DFT). The DFT takes in a length  $N$  string of complex values,  $x_j$ , and outputs a string of  $N$  complex values,  $y_k$ , such that

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j e^{2\pi i \frac{k}{N} j}. \quad (8.1)$$

The DFT inverts the periodicity of the input string: if the input string has entries that repeat itself every  $r$  entries (it has period  $r$ ), the output string has entries with period  $N/r$ . If  $r$  does not divide  $N$ , the result is approximate. The DFT also converts the phase of the

$r$	input	output	$N/r$
8	1 0 0 0 0 0 0 0	1 1 1 1 1 1 1 1	1
4	1 0 0 0 1 0 0 0	1 0 1 0 1 0 1 0	2
2	1 0 1 0 1 0 1 0	1 0 0 0 1 0 0 0	4
1	1 1 1 1 1 1 1 1	1 0 0 0 0 0 0 0	8

Table 8.2: Examples of input strings with different periodicities and the output strings after applying the DFT operation.

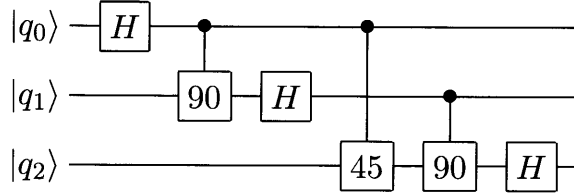


Figure 8-3: Quantum circuit for the 3-qubit inverse quantum Fourier transform algorithm.

input string (an offset in the location of the first element) into phase factors in front of the terms of the output string. Several examples of the DFT input and outputs on a string of 8 values is shown in Table 8.2.

The QFT performs the same transformation, except that the input and output strings are stored in the amplitude and phase of the states of  $n = \log_2 N$  qubits. For example, with  $N = 8$  states, there are 3 qubits, and the possible states are  $|000\rangle, |001\rangle \dots |111\rangle$ .

The quantum circuit for the QFT requires  $O(n^2)$  Hadamard gates and controlled phase gates, whereas the classical DFT requires  $O(n2^n)$  gates, where  $n$  is the number of (qu)bits [NC00]. The QFT appears in many quantum algorithms, including the order-finding algorithm and Shor's algorithm. In those cases it is actually the inverse QFT that is needed. The quantum circuit for the 3-qubit inverse QFT is shown in Figure 8-3.

### 8.3.2 Pulse sequence

The gate set we have chosen in our experiment (Section 7.1.3) does not include single-qubit Hadamard gates or controlled phase gates. As stated earlier, conversion between our chosen gate set including global gates, to single- and two-qubit gates is possible but not necessarily efficient. Therefore for the QFT and all subsequent algorithms described here, we use the algorithm's unitary operator as an input to the pulse compiler to find optimized pulse sequences. The unitary operator for the QFT on  $n = \log_2 N$  qubits is:

$$U_{qft} = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 & 1 & 1 & 1 & \dots & 1 \\ 1 & \omega & \omega^2 & \omega^3 & \dots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \dots & \omega^{2(N-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \dots & \omega^{3(N-1)} \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \omega^{3(N-1)} & \dots & \omega^{(N-1)(N-1)} \end{pmatrix} \quad (8.2)$$

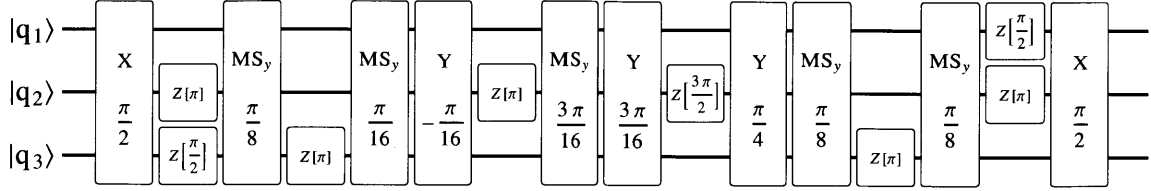


Figure 8-4: Optimized pulse sequence for the 3-qubit QFT. X, Y = global carrier rotations;  $MS_x$ ,  $MS_y$  = global MS gates; Z = addressed AC Stark shifts.

where  $\omega = e^{\frac{2\pi i}{N}}$  is the  $N$ th root of unity. The optimized pulse sequence is shown in Figure 8-4. The pulse sequence written in the TIQC simulator syntax is included in Appendix D.

We experimentally realize the QFT on 3 ions, and prepare a number of input states that are representative of the different periods including those listed in Table 8.2. These states, denoted by  $|\psi_r\rangle$  where  $r$  is the period (normalization factors omitted), are:

$$\begin{aligned}
 |\psi_1\rangle &= |111\rangle + |110\rangle + \dots + |000\rangle \\
 |\psi_2\rangle &= |110\rangle + |100\rangle + |010\rangle + |000\rangle \\
 |\psi_3\rangle &= |110\rangle + |100\rangle + |011\rangle + |001\rangle \\
 |\psi_4\rangle &= |011\rangle + |000\rangle \\
 |\psi_8\rangle &= |000\rangle
 \end{aligned}$$

The state  $|\psi_8\rangle$  corresponds to all ions being in the ground state, and is the default initialized state for any experiment. All other states are generated by additional pulse sequences before the start of the main QFT pulse sequence. For example, the state  $|\psi_1\rangle$  is generated by applying a global  $R(\pi/2, 0)$  gate. The pulse sequences to generate all the starting states are listed in Appendix D. The state for period = 3 is an approximate representation, identical to the one used in Ref. [CBL<sup>+</sup>05]. The expected output of the QFT with these input states are states which invert the periodicity. For example,  $|\psi_1\rangle$  is the output of the QFT on  $|\psi_8\rangle$  and vice versa;  $|\psi_2\rangle$  and  $|\psi_4\rangle$  forms a similar pair. The expected output of the QFT on the state  $|\psi_3\rangle$  has a period equal to the approximate value of  $c\frac{8}{3}$  for integers  $c$ , which in this case gives 3 or 5.

### 8.3.3 Experiment and simulation

We test the QFT pulse sequence on the various initial states representing different periods. Each experiment begins with Doppler cooling and sideband cooling to initialize all 3 ions to the ground state, then optionally performing one of the initialization pulse sequences to generate the desired initial state with a known period. The QFT pulse sequence is applied, followed by full state tomography with a CCD camera which is capable of individual ion resolution. Each experiment is repeated 100 times to estimate the probability of the output state using observed frequency statistics. In addition, we perform measurements in which the QFT pulse sequence is only realized up to  $n$  pulses, where  $0 \leq n \leq \text{length of sequence}$ ,

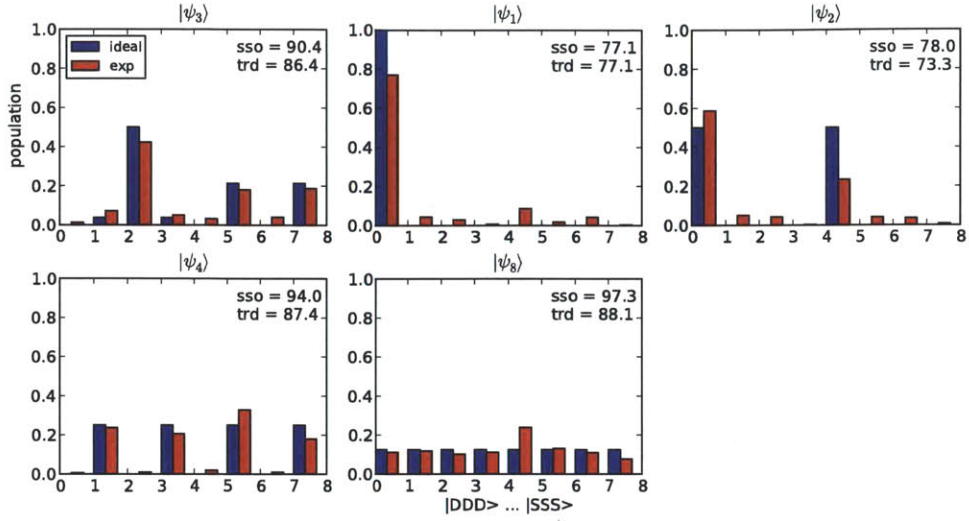


Figure 8-5: Initial and final state populations in the QFT pulse sequence for input states of various periods. The squared statistical overlap (sso) and trace distance (trd) for each output are shown.

and full state tomography is performed after such partial sequences. This allows us to characterize the state fidelity as a function of the number of pulses applied so far in the pulse sequence. Number of pulses = 0 indicates initial state preparation only.

The Rabi frequency of the global carrier pulse is 23 kHz and the Rabi frequency of the individual AC stark pulse is 25 kHz. The detuning of the MS gate is set such that the pulse length of an  $MS(\pi/16)$  gate is  $35 \mu s$ . In total the QFT pulse sequence is 1.3 ms, which includes the time that it takes to switch the addressing beam between different ions ( $46 \mu s$ ) or from addressed beam to the global beam, and to switch between consecutive global beams ( $7 \mu s$ ). This is still well below the measured decoherence time of 5 ms.

The output state populations for the input states  $|\psi_r\rangle$  for  $r = 1, 2, 3, 4, 8$  and the population fidelities (see Table 8.3) are shown in Figure 8-5. The output states agree with the expected outputs indicating periodicities of the input state, confirming the correctness of the QFT algorithm and implemented pulse sequence.

We simulate the QFT pulse sequence with the default set of decoherence parameters. Figure 8-6 shows the state populations for all 8 input states as a function of pulse number for the  $|\psi_8\rangle$  input state, for the ideal case, experiment, and simulation. The ideal case is computed with the ideal unitary operators for each pulse in the sequence. The simulation, while it does not capture all details of the experimentally measured state evolution, nevertheless reproduces some features such as relative ordering of the state populations.

Figure 8-7 shows the experimental and simulated results as a function of the number of pulses, for each of four distance measures, for the input state  $|\psi_8\rangle$ . For a given state  $|\psi\rangle$  or a density matrix  $\rho$ , Table 8.3 lists the definitions of distance measures between them and the ideal states or density matrices. The experiment and simulation shows reasonable agreement for most of the distance measures calculated here. A number of error sources



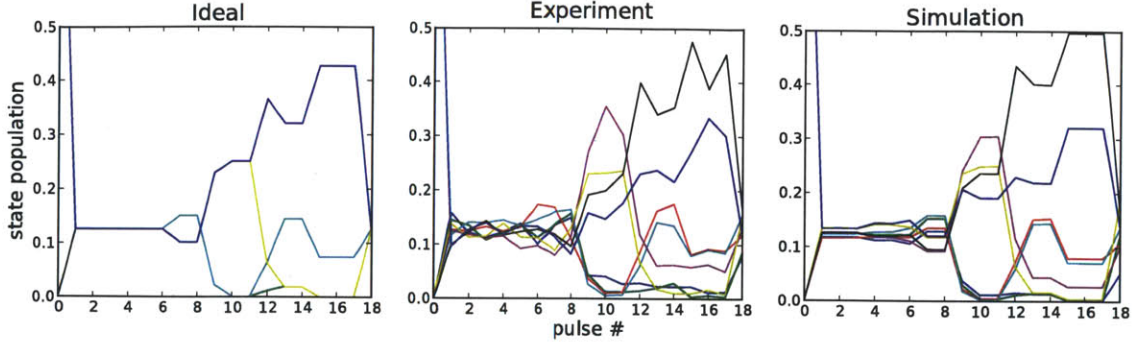


Figure 8-6: Comparison of population after each pulse for ideal, experiment, simulation in the QFT pulse sequence for input state  $\psi = |000\rangle$  (period = 8). Each color corresponds to a different basis state ( $|000\rangle, |001\rangle \dots |111\rangle$ ).

distance measure	definition
Jozsa fidelity	$\text{Tr}(\sqrt{\sqrt{\rho_1} \times \rho_2 \times \sqrt{\rho_1}})^2$
squared statistical overlap (sso)	$\left(\sum_{i=0}^n \sqrt{p_1^i} \sqrt{p_2^i}\right)^2$
trace distance (density matrix)	$1 - \frac{1}{2} \text{tr} \rho_1 - \rho_2 $
trace distance (population)	$1 - \frac{1}{2} \text{tr} p_1^i - p_2^i $

Table 8.3: Definitions of various distance measures used to define distance between density matrices  $\rho_1, \rho_2$ , or states  $\psi_{1,2} = \sum_{i=0}^n p_{1,2}^i |\psi_b\rangle$  where  $|\psi_b\rangle$  are the basis states. For the trace distance for density matrices,  $|X| = \sqrt{X^\dagger X}$ .

in the experiment are not included in the simulation yet. One of them is error during the state tomography pulses. The state tomography pulse sequences are up to 6 pulses long (depending on the measurement basis) for the 3 qubits, which is a significant fraction of the total pulse length in this case. The direct measurement of population requires no additional pulses. This may explain the greater discrepancy in the distance measures based on the full density matrix, compared to the distance measures computed from populations only. Spontaneous decay during measurement is not simulated, though it is expected to have a small effect since the measurement time of 8 ms is about 3 orders of magnitude shorter than the lifetime of the D state. Also, improperly set parameters and slow drifts over the course of several hours are not included in the simulation, but these effects are minimized during the experiment by automatically performing standard measurements (such as Ramsey sequences and spectrum scans) to correct for drifts in parameters.

To evaluate the effect of different sources of error on the algorithm, we simulate the algorithm with only one error source turned on at a time, repeated for 500 iterations each. Table 8.4 lists the major error sources and their contributions to the infidelity for each of the distance measures evaluated in Figure 8-7. The table suggests that addressing error is a dominant source of decoherence in our implementation of the QFT algorithm, even given the choice of an operator set that seeks to minimize it. The last 3 entries, initialization error, spontaneous decay, and heating have minimal contribution, though their entries in

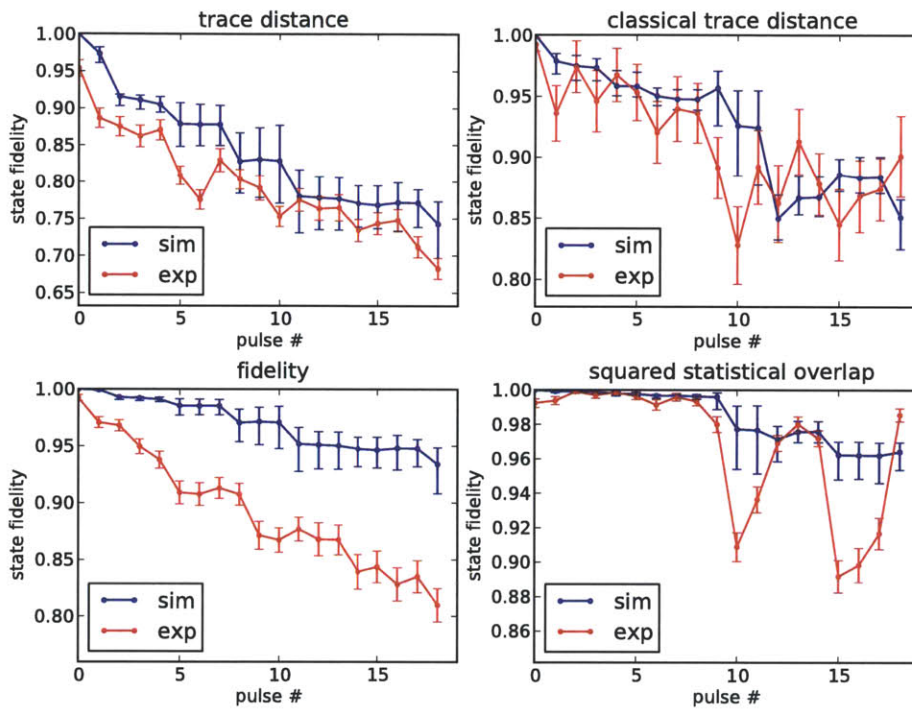


Figure 8-7: State fidelity vs pulse for the QFT sequence with input state  $\psi = |000\rangle$ , for 4 distance measures: trace distance (population and density matrix), squared statistical overlap, and Jozsa (Uhlmann) fidelity.

	trace distance (density matrix)	trace distance (population)	squared statistical overlap	Jozsa fidelity
all	$26.6 \pm 5.5$	$15.4 \pm 2.9$	$4.0 \pm 1.5$	$7.4 \pm 3.2$
addressing	$21.4 \pm 0.0$	$12.0 \pm 0.0$	$2.7 \pm 0.0$	$4.6 \pm 0.0$
dephase	$11.6 \pm 7.3$	$6.1 \pm 3.7$	$0.8 \pm 0.9$	$1.9 \pm 2.2$
specmode	$6.4 \pm 4.9$	$4.5 \pm 3.6$	$0.5 \pm 0.8$	$0.6 \pm 1.0$
intensity	$4.2 \pm 1.4$	$2.0 \pm 1.0$	$0.1 \pm 0.1$	$0.2 \pm 0.1$
init err	$1.8 \pm 7.4$	$1.3 \pm 4.9$	$0.5 \pm 5.0$	$0.6 \pm 5.7$
spont decay	$1.1 \pm 0.0$	$0.8 \pm 0.0$	$0.0 \pm 0.0$	$0.0 \pm 0.0$
heating	$1.1 \pm 0.0$	$0.8 \pm 0.0$	$0.0 \pm 0.0$	$0.0 \pm 0.0$

Table 8.4: Error budget for the QFT pulse sequence. Errors are derived from a simulation of 500 Monte Carlo instances.

the table should be taken as approximations due to small sampling error: a simple estimate for a heating rate of 1 quanta/sec in a pulse sequence of 1 ms, for example, is expected to contribute no more than a 0.1% error, which translates to an expected 0.5 occurrences in 500 simulations. Addressing error is a technical issue and several options exist to further minimize its effect. The current implementation of addressing consists of electro-optic modulators (EOMs) that deflect the beam path through a series of optics, in which the degree of freedom associated with the position of a lens is coupled to the angle of deflection. Improvements to the addressing optics should suppress this problem [Mon11]. Pulse-based addressing error corrections have also been investigated and shown to minimize addressing error, though at the cost of a longer pulse sequence and thus increased sensitivity to decoherence [Ben08].

### 8.3.4 Semiclassical version

In the original QFT circuit as shown in Figure 8-3, the Hadamard and controlled phase gates are sequentially applied to each qubit. Effectively, each qubit is only used once, and afterward its (unchanged) state serves as control to subsequent qubits. In this case, one can replace the (quantum) controlled gates with classically controlled gates, namely by making a measurement on the qubit and using the resulting classical bit to control the subsequent gate [GN96]. The resulting “semiclassical” QFT, also known as the Kitaev version, is shown in Figure 8-8.

Furthermore, we can carry out the QFT by only using one qubit, repeatedly measuring and reinitializing it. This requires two additional operations in the experiment: feedforwarding and recooling. The pulse sequence used, in the TIQC simulator notation, is included in Appendix D.

We experimentally realized the semiclassical QFT, using the same input states as for the fully quantum version described above. A single ion is measured 3 times and recooled twice, to take the place of 3 virtual ions. The various input states are generated by adding additional Hadamard gates to create product states. For example, an additional Hadamard gate in the very beginning of the pulse sequence generates the state  $|000\rangle + |100\rangle$ , which has period = 4. Figure 8-9 shows the results of the combined populations of the 3 virtual ions

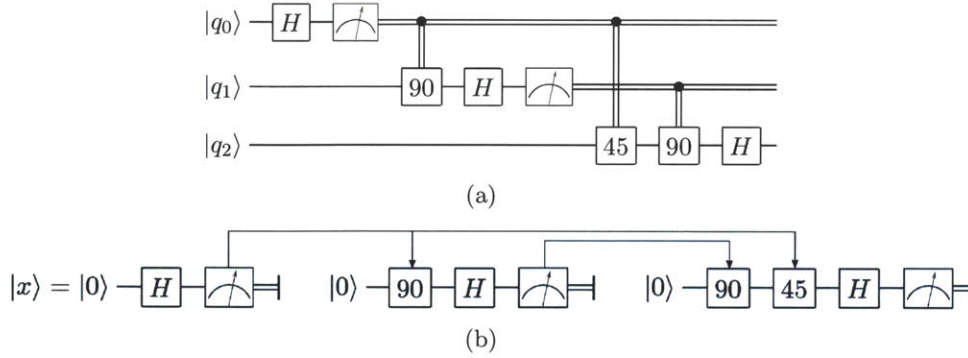


Figure 8-8: (a) Semiclassical QFT, replacing quantum controlled gates by classical controlled gates. (b) Kitaev version of the semiclassical QFT, where 1 qubit act as 3 virtual qubits.

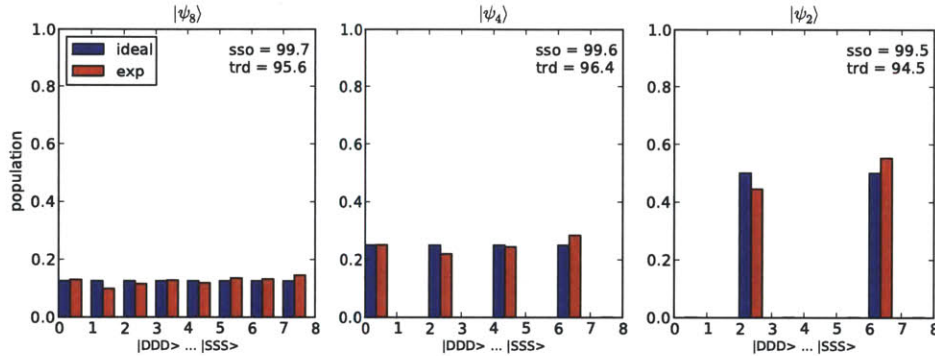


Figure 8-9: Measured state populations in the semiclassical QFT pulse sequence for input states of various periods.

and shows that the semiclassical QFT produces the expected periods. In the next section, we describe how to use this technique to simplify the implementation of the order-finding algorithm.

## 8.4 Order-finding

The order-finding problem is a general case of Shor’s algorithm for factoring. The order-finding problem is to find the order  $r$  when given  $\pi$  and  $y$ , by making queries of the type “make  $x$  applications of  $\pi(y)$  and check if the result is  $y$ ”. The problem can be illustrated as in Figure 8-10. In this figure, there are 8 elements in a set  $M$ , and the permutations are represented by cycles containing one-way arrows connecting a subset of  $M$ . The order  $r$  is 3 if  $y \in \{2, 4, 5\}$ ,  $r = 1$  if  $y = 6$ , and  $r = 4$  if  $y \in \{0, 1, 3, 7\}$ .

The classical complexity of order-finding, defined as the number of queries to  $\pi^x(y)$  necessary to obtain a given probability of success, is lower-bounded to be exponential in the problem size  $M$  [Cle00]. The quantum algorithm, by contrast, requires a constant number of queries for a fixed probability. Intuitively, it’s as if all transitions are made in parallel, followed by a QFT step which identifies the correct order through interference.

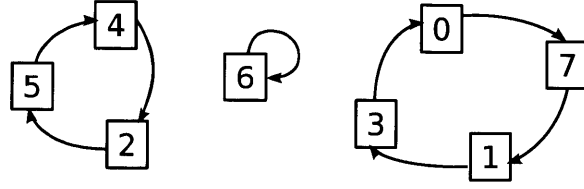


Figure 8-10: Example of a permutation on 8 elements.

The quantum order-finding algorithm was first demonstrated in NMR systems using 5 qubits and 50-200 pulses [VSB<sup>+</sup>00] and the discussion of the algorithm here closely follows that described in [Van01]. A very recent implementation was achieved with photons [LLL<sup>+</sup>11] using 2 photons encoding up to 5 qubits, but required the use of post-selection.

### 8.4.1 Algorithm

The steps of the order-finding algorithm can be illustrated with the following example using  $M = 4$  states, labeled 0, 1, 2, 3, and the permutation as follows:  $\pi(0) = 0$ ,  $\pi(1) = 3$ ,  $\pi(2) = 2$ ,  $\pi(3) = 1$ . We will use a register  $|y\rangle$  of  $m = \log_2 M = 2$  qubits to store the input, and a register  $|x\rangle$  of  $n = 3$  qubits to store the answer. The maximum order for a set of  $M = 4$  states is 4, so in principle  $n = \log_2 M = 2$  qubits would be sufficient; however, in general  $r$  may not be a power of 2 or even divide  $M$ , and in such cases a continued fraction expansion is necessary to estimate  $r$  from  $cN/r$  [Kob94] where  $N = 2^n$ . In our small example, we will use one additional qubit to provide the extra digit of precision on determining  $N/r$  where  $N$  and  $r$  are co-prime.

Suppose we wish to find the order of  $\pi(y = 1)$ . We initialize the system to the state  $|x\rangle |y\rangle = |000\rangle |01\rangle$ . Then perform a Hadamard gate on the first register  $x$ . Now the state becomes:

$$|\psi_1\rangle = (|0\rangle + |1\rangle + |2\rangle + |3\rangle + |4\rangle + |5\rangle + |6\rangle + |7\rangle) |y\rangle. \quad (8.3)$$

Next we evaluate  $\pi^x(1)$  for every  $x \in [0, 7]$ .  $\pi^x$  can be represented by  $\pi^x = \pi^{x_0} \pi^{2x_1} \pi^{4x_2}$ , where  $x_2 x_1 x_0$  is the binary representation of  $x$ . In circuit terms, these represent controlled permutations on  $y$  controlled by qubits  $x_2$ ,  $x_1$ , and  $x_0$  respectively. The resulting state is

$$\begin{aligned} |\psi_2\rangle &= |0\rangle |1\rangle + |1\rangle |3\rangle + |2\rangle |1\rangle + |3\rangle |3\rangle + |4\rangle |1\rangle + |5\rangle |3\rangle + |6\rangle |1\rangle + |7\rangle |3\rangle \\ &= (|0\rangle + |2\rangle + |4\rangle + |6\rangle) |1\rangle + (|1\rangle + |3\rangle + |5\rangle + |7\rangle) |3\rangle. \end{aligned} \quad (8.4)$$

At this point, a single measurement of register  $x$  would give 0..7 with equal probability and thus provide no useful information. The answer we are looking for is actually the periodicity of  $x$ , which corresponds to the number of elements in the set 1,3 contained in this permutation, which gives the desired order. Therefore we perform a QFT on register  $x$ . The resulting state is

$$|\psi_3\rangle = (|0\rangle + |4\rangle) |1\rangle + (|0\rangle - |4\rangle) |3\rangle. \quad (8.5)$$

Now a measurement on register  $x$  would return either 0 or 4, which are multiples of  $N/r$  as

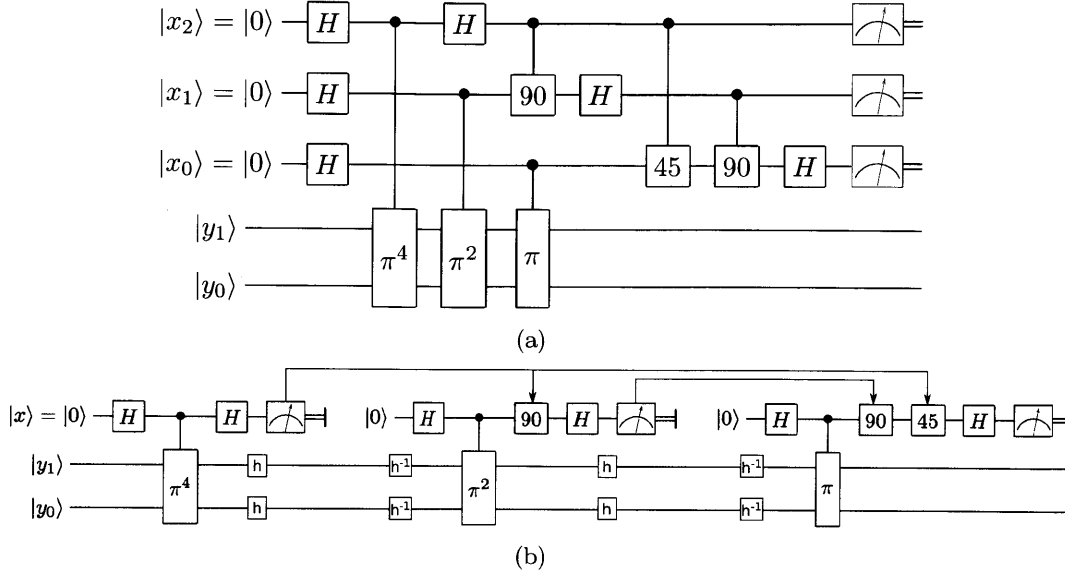


Figure 8-11: (a) Quantum circuit for order-finding in the case of  $M = 4$ .  $\pi$  represents the permutation for which the order we seek to find,  $H$  = Hadamard gate,  $Z$  = phase shift gates. The top register of 3 qubits are measured at the end of the circuit. (b) Order-finding implemented with the semiclassical QFT. One qubit is recycled by repeated measuring and reinitialization, acting as 3 virtual qubits. Here,  $h$  = hiding,  $h^{-1}$  = unhiding.

expected. The algorithm as stated is represented by the quantum circuit shown in Figure 8-11(a).

The circuit for the order-finding algorithm is similar to the QFT circuit in that the controlled operations are applied sequentially to each qubit. Thus we can apply the same logic and reduce the number of qubits required by recycling only one qubit to perform the measurement, essentially using the semiclassical QFT. This one qubit serves as 3 virtual qubits for the QFT. The circuit for this Kitaev order-finding is shown in Figure 8-11(b). This implementation also applies to Shor’s algorithm and is described in Ref. [PP00].

### 8.4.2 Pulse sequence

For an experimental realization of order-finding with ions, we choose a representative set of permutations that was used in the NMR demonstration [VSB<sup>+</sup>00]. The 4 permutations used are listed in Table 8.5 including the transition matrices corresponding to the unitary operators. The necessary task is to convert the controlled permutations to pulse sequences. For  $y = 4$  input states, each permutation can be written as a transition matrix where the  $(i, j)$ -th entry is 1 if  $\pi(i) = j$ , otherwise the entry is 0. The controlled version of these matrices, of size  $8 \times 8$ , are the desired controlled unitaries that we wish to implement. The pulse compiler uses these matrices to derive optimized pulse sequences for the unitary operations. The full pulse sequences are listed in Appendix D.

The Hadamard gates on the first qubit are implemented by a sequence of refocusing pulses, since the individually addressed carrier rotation pulse is not within our operator set. The pulse sequence necessary to implement an addressed Hadamard gate on the first qubit



$y$	$\pi_1(y)$	$\pi_2(y)$	$\pi_3(y)$	$\pi_4(y)$
0	0	1	0	3
1	3	0	3	0
2	2	3	1	1
3	1	2	2	2
$\max(r)$	2	2	3	4
unitaries	$\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \\ & & & 1 \end{pmatrix}$
pulse sequences:				
$\pi$	3Z,5R,3M	4Z,4R,2M	9R,10Z,4M	11Z,9R,5M
duration	1.2 ms	0.96 ms	2.4 ms	2.4 ms
$\pi^2$	-	-	7Z,6R,4M	4Z,4R,2M
duration	-	-	1.9 ms	0.76 ms

Table 8.5: Representative permutations for order-finding on 4 elements, showing the output after one application of the permutation on the inputs 0,1,2,3. Pulse sequence compositions and durations used in the experiment are also listed; Z: AC stark pulse, R: carrier pulse, M: MS gate. Note that  $\pi_3^4(y) = \pi_3(y)$ .

is  $R(\pi/4, \pi)Z(\pi, 0)R(\pi/4, 0)$ .

For planning future experiments, we investigated whether the frequent use of this refocusing gate for the single-qubit Hadamard gate is preferable over an addressed carrier pulse on the single ion. Generally the operator set consisting of addressed AC Stark pulses instead of addressed carrier pulses are preferred, since the Z rotation caused by an addressed AC Stark pulse is less sensitive to addressing error. However, in this case we directly trade off this advantage by requiring 2 extra pulses. We simulate both cases using the standard addressing error listed in Table 7.6 and found the expected gate fidelity of the addressed gate to be 99.9% while the expected gate fidelity of the refocused Hadamard gate to be 99.6%. This is a small difference and likely not relevant for the types of pulse sequences considered here.

The intermediate measurements of the qubit forming register  $|x\rangle$  for feedforwarding is done via the PMT, which offers no ability to distinguish light from different ions. Therefore it is necessary to hide the qubits of the other register before performing these measurements. Hiding ions is accomplished by driving transitions between Zeeman sublevels that don't compose the main qubit transition. The  $S$  qubit level is hidden to the  $D'(m = -5/2)$  and the  $D$  qubit level is hidden to the  $S'(m = 1/2)$  level. Applying a  $\pi$  pulse on the appropriate transition transfers the population to the  $D'$  or  $S'$  levels where the qubit laser does not couple to the ion. For detection, only the pulse to hide the  $S$  level to the  $D'$  level is needed.

### 8.4.3 Experiment and simulation

We implement the order-finding algorithm in the ion trap experiment using 3 ions, by applying each of the pulse sequences for the  $\pi(y)$ 's and measuring the output of the QFT on the first register. Following Doppler and sideband cooling, we apply a pulse sequence that



realizes the circuit in Figure 8-11(b). Prior to both the first and second measurement, a  $\pi$  pulse is applied on the  $S \leftrightarrow D'$  transition to hide qubits  $|y_0\rangle$  and  $|y_1\rangle$  from the PMT detection. The subsequent phase rotations on the qubit  $|x\rangle$  is conditionally applied depending on the outcome of these measurements. The final measurement is performed with the CCD camera, which is capable of individual ion resolution, so the qubits in the bottom register are not hidden for the final measurement.

For all of the  $\pi(y)$ 's defined in Table 8.5 except for  $\pi_3(y)$ , the  $\pi^4$  unitary is the identity operation. For  $\pi_1(y)$  and  $\pi_2(y)$ , similarly the  $\pi^2$  unitary is identity. If  $\pi^4 = \text{identity}$ , then the first qubit will be found in the state  $|0\rangle$  after the first measurement. Similarly if  $\pi^2 = \text{identity}$ , the second qubit will always be found in the state  $|0\rangle$ . Knowing this in advance, we could shorten the pulse sequence considerably by omitting these virtual qubits. This would be effectively equivalent to omitting the extra bits of precision on the QFT result, knowing they would be 0. However, we choose to always carry out the full sequence including the hiding and unhiding pulses, since there are noise sources in the experiment and the measurement result on the first virtual qubit might not return  $|0\rangle$  with 100% probability. The full sequence thus allows us to have a better characterization of the errors in the experiment.

The entire pulse sequence is repeated 4000 times per order operation to measure the probabilities representing the QFT output. Errors on the measured probabilities are due to projection noise and are at the level of  $< 1\%$  for 4000 measurements. The result of each experiment is 3 pairs of population measurements, all on the first qubit, corresponding to each of the 3 virtual qubits. Converting the measurement results to the state populations for the QFT is done as follows. The outcome of a single measurement can be expressed in terms of a concatenation of binary bits: for example, if the first virtual qubit =  $|1\rangle$ , second virtual qubit =  $|0\rangle$ , third virtual qubit =  $|1\rangle$ , then this measurement represents a vote for the population of the state  $|101\rangle$ . Repeating the experiment and tabulating all the populations give the outcome of the QFT similar to the state populations as obtained by the fully quantum QFT experiment described in Section 8.3.3.

Figure 8-12 shows the ideal and experimentally measured outcomes of the QFT for each of the permutations. Table 8.6 shows the trace distance and sso for the experimental and simulated data. For the order = 1,2,4 permutations, the measured populations agree well with the expected order. However, the order = 3 results are consistent with a random mixed state, suggesting that the pulse sequence is realized with low fidelity in this case. The order = 3 permutation is the only one that requires all of the  $\pi$ ,  $\pi^2$ , and  $\pi^4$  operations to be executed and the total pulse sequence is significantly longer than the rest.

We simulate the pulse sequences for the 3-qubit semiclassical order-finding algorithm as realized in the experiment. The semiclassical pulse sequences require simulating the feed-forward. This is accomplished by simulating all possible combinations of pulse sequences, storing the measured populations of the  $|D\rangle$  and  $|S\rangle$  states, and combining the populations appropriately afterward to obtain the desired probabilities as the outcome of the QFT. Since the order-finding pulse sequence contains two measurement and feedforwards, 4 simulations are needed to capture all the possible pulse sequence combinations.

The set of decoherence parameters we have used to simulate order-finding is the same

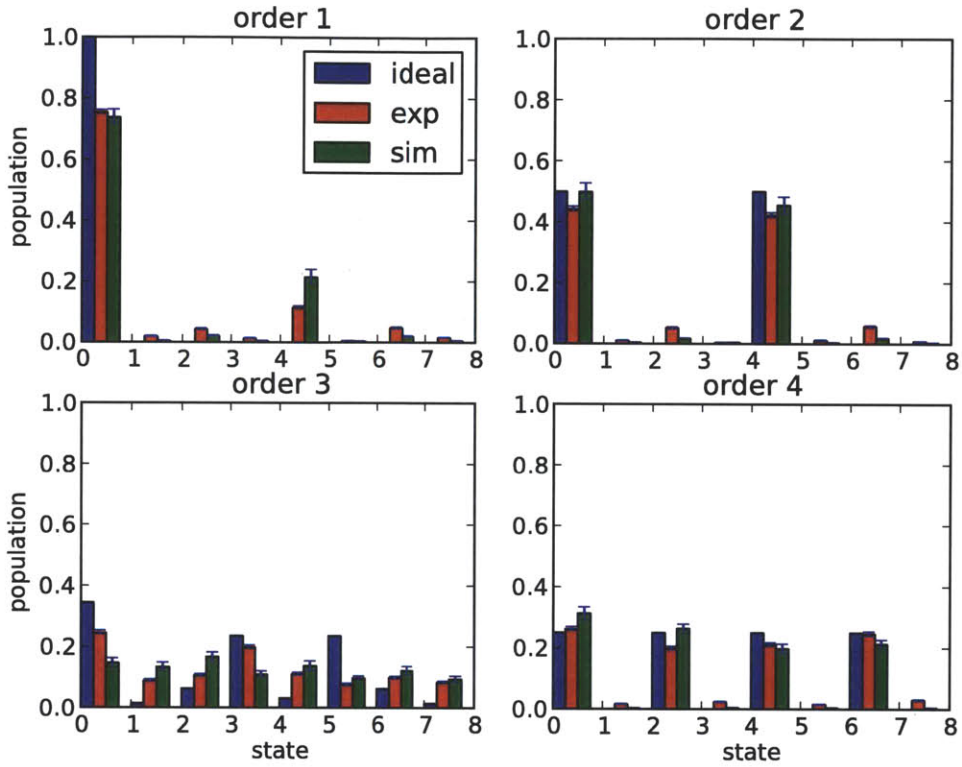


Figure 8-12: Ideal, experimental, and simulated state populations for the order finding algorithm on 3 qubits.

order	data	sso	trace distance
1	sim	$73.3 \pm 3.4$	$73.3 \pm 3.4$
	exp	$75.3 \pm 0.7$	$75.3 \pm 0.7$
2	sim	$93.6 \pm 0.5$	$87.5 \pm 1.5$
	exp	$86.4 \pm 0.6$	$86.5 \pm 0.6$
3	sim	$74.3 \pm 1.1$	$55.2 \pm 1.1$
	exp	$85.9 \pm 0.6$	$70.3 \pm 0.8$
4	sim	$96.1 \pm 0.5$	$86.6 \pm 1.1$
	exp	$91.6 \pm 0.5$	$90.7 \pm 0.6$

Table 8.6: Experimental and simulated fidelities (sso and trace distance) for the population outputs of the order-finding algorithm.

as for the QFT, as summarized in Table 7.6, except that the most recently measured coherence time is closer to 15 ms. In addition, there are errors resulting from the hiding and unhiding operations. These operations are implemented by carrier pulses on other transitions, which may both be more sensitive to magnetic field fluctuations, and are not calibrated as completely as the main qubit transition. These transitions are also susceptible to the same error sources as the rest of the pulse sequence. Rather than modeling the full Hilbert space now consisting of 4 levels (which would increase the Hilbert space dimension from  $2^N$  to  $4^N$ ), we use the following two simple models for hiding errors. First, each hide/unhide cycle has a probability of the ion not returning to the original state, the net effect being that the ion is effectively decoupled from the rest of the pulse sequence. Second, during the detection due to the PMT, if an ion in the  $|S\rangle$  state was not hidden successfully, then fluorescence from the wrong ion is detected and is a spurious signal. Both of these errors can occur in the case of an offset in the Rabi frequency of the other carrier pulses. We model the first error as a probability that an ion is decoupled from the computational space at the 1% level. We model the second error as a (fixed) probability that the PMT will detect extra population in the  $|S\rangle$  state, at the level of 0.5% per ion per hiding or unhiding operation. These levels are estimated by performing experiments in which only the hiding and unhiding pulses are applied, and observing the fidelity of the measured populations compared to the ideal populations.

The outcome of the simulations are as shown in Figure 8-12. Here the errorbars are at 1 sigma, for a total of 24 simulations for each case. The simulations take between 20 min for the  $\pi_2(y)$  sequence and 2.2 hours for the  $\pi_3(y)$  sequence, on a Core-i7 2.67GHz machine. The simulation has few free parameters – most of the error parameters can be related directly to experiments designed to specifically measure the error. For the  $\pi_1(y)$ ,  $\pi_2(y)$ , and  $\pi_4(y)$  permutations, the simulated populations and fidelities are comparable or better than the experimental data. The simulated fidelities should be considered an upper bound, because systematic effects such as offsets in a parameter in the experiment are not included in the simulation. In the experiment, these offsets can have large effects on the output fidelity. The  $\pi_3(y)$  permutation does not match the ideal or experimental data. One possible cause is that the QFT output for this permutation operator is sensitive to phases of the individual pulses, which was not explicitly checked in the experiment. In principle this could be done by performing state tomography after each pulse, as was done in the QFT experiment.

The successful demonstration of order-finding for all orders up to 4 validates the choices that we have made in implementing this algorithm. In particular, for example, the semi-classical QFT using measurement and feedforward enabled the use of 2 less qubits but with no extra pulses applied. The reasonable agreement between the measured and simulated data is evidence that the most important error sources are included in our model. In the next section, we describe an application of order-finding for a famous quantum algorithm: Shor’s algorithm for factoring integers.

## 8.5 Shor's algorithm

Shor's algorithm for factoring prime numbers using a quantum computer was one of the first quantum algorithms proposed that showed a significant speedup over classical computers. The first demonstration of the algorithm in an NMR quantum computer was widely considered to be an important landmark [VSB<sup>+</sup>01]. Subsequently many other physical systems ([PMO09, XZL<sup>+</sup>11, LBY<sup>+</sup>11]) followed in performing this algorithm as a milestone toward scalable systems. In this section, we describe the experimental implementation of Shor's algorithm with an ion trap quantum computer using 5 ions. The choice of circuits here closely follows the NMR implementation as described in Ref. [Van01].

### 8.5.1 Algorithm

Shor's algorithm for factoring is closely related to the order-finding problem. It consists of finding the order of a particular permutation:

$$\pi(y) = ay \bmod N \tag{8.6}$$

for  $y = 1$ .  $N$  is the integer we want to factor and  $a$  is any integer  $< N$  which is coprime with  $N$  (does not have common factors other than 1). For a random choice of  $a$ , with high probability at least one prime factor of  $N$  is given by

$$\gcd(a^{r/2} - 1, N) \quad \text{or} \quad \gcd(a^{r/2} + 1, N). \tag{8.7}$$

Euclid's algorithm can be used to compute the gcd of two integers on a classical computer. This version of the factoring problem is discussed in number theory [Kob94]. To implement order-finding for this particular permutation, we need to implement the permutations  $\pi^x(y) = a^x y \bmod N$  for each  $x$ . These are known as the modular exponentiation functions.

The smallest meaningful instance of Shor's algorithm is  $N = 15$ , since the algorithm fails if  $N$  is even or is the power of a prime number. For a given  $N$ , the register storing  $N$  must have  $m = \lceil \log_2 N \rceil$  qubits and the register for performing the QFT needs  $2m$  qubits. The modular exponentiation is converted to  $2 \log_2 N$  multiplications by the same method as for the order-finding.

For  $N = 15$ , the table below calculates two examples of  $a^x \bmod N$ :  $a = 7$  and  $a = 11$ . It turns out that for all  $a$  coprime with  $N$ , the periods  $r$  are less than 4, and we can leave out all multiplications by  $a^x \bmod N$  where  $x \geq 4$ . Thus we can use only 2 qubits to store the result of the modular exponentiations for the QFT. We add an extra qubit to represent an extra digit of precision for  $x$ .

$a$	$x$					$r$	$a^{r/2} \pm 1$	$\gcd(a^{r/2} \pm 1, 15)$
	0	1	2	3	4			
7	1	7	4	13	1	4	48,50	3,5
11	1	11	1	11	1	2	10,12	3,5

The simplified circuit for Shor's algorithm for  $N = 15$  is shown in Figure 8-13. For our proposed experimental implementation we make one additional simplification: the quantum

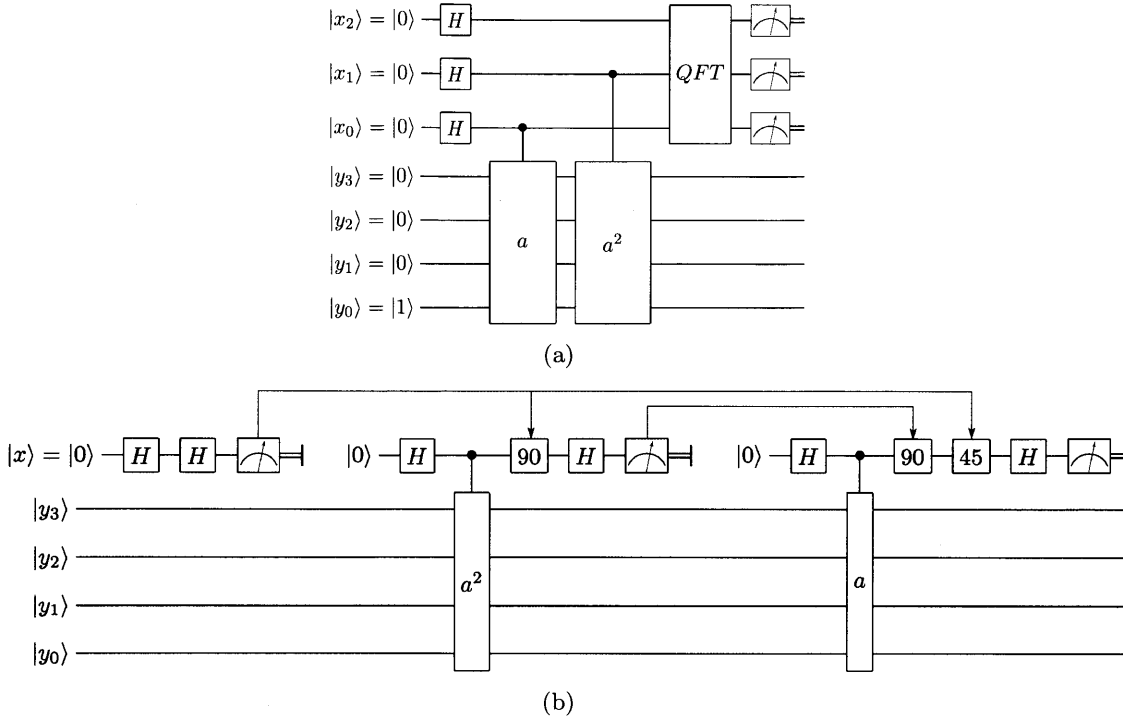


Figure 8-13: (a) Simplified quantum circuit for Shor's algorithm for factoring  $N = 15$ . (b) Semiclassical version.

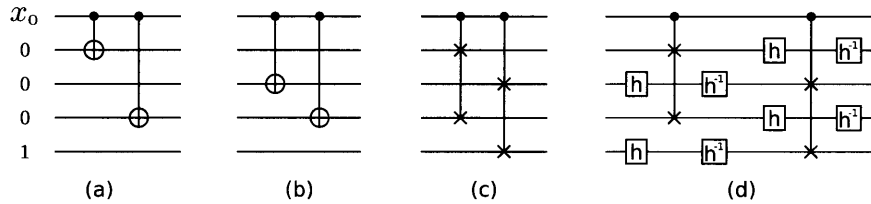


Figure 8-14: Circuits for the modular exponentiation: (a)  $a = 11$ , (b)  $a = 7, x = 1$ , (c)  $a = 7, x = 2$ , (d)  $a = 7, x = 2$  with hiding (h) and unhiding ( $h^{-1}$ ) operations shown. (a) and (b) uses 2 CNOT gates. (c) and (d) uses 2 Fredkin (controlled-swap) gates.

Fourier transform can be carried out semiclassically, just like for the order-finding in the previous section. Thus the final version of our factoring circuit requires 5 qubits, including 2 classical registers to store the intermediate results.

### 8.5.2 Pulse sequence

The circuit in 8-13(b) is identical to the circuit used for the order-finding algorithm, except that the controlled- $\pi(y)$  operations are replaced by the controlled- $a^x \bmod N$  operations. As in the NMR case, we choose to run the algorithm for the case  $a = 11$  ("easy" case, because the expected order = 2 and requires only one controlled-  $a \bmod N$  operation) and the case  $a = 7$  ("hard" case, because the expected order = 2 and requires both controlled-  $a \bmod N$  and controlled-  $a^2 \bmod N$ ). The circuit representation for the controlled multiplications are shown in Figure 8-14.

In this case, rather than providing the unitaries generated by the 2 CNOTs to our pulse compiler, we used rules defined in Ref. [Neb08] for producing pulse sequences consisting of consecutive CNOTs for arbitrary number of qubits where the control qubit stays the same. The pulse sequences for 2 CNOTs, applicable to  $11 \bmod 15$  and  $7 \bmod 15$ , consists of 16 pulses and is listed in Appendix D. The Fredkin gate in  $7^2 \bmod 15$  was converted to a pulse sequence consisting of 18 pulses by the pulse compiler. However, the resulting pulse sequences only apply to the 3-qubit Fredkin gate. To use these sequences for our 5-qubit register, hiding and unhiding pulses are necessary before and after each Fredkin gate to apply it to the correct ions. Figure 8-14(d) shows the complete sequence for the  $7^2 \bmod 15$  operation with hiding pulses included.

### 8.5.3 Experiment and simulation

We implement Shor’s algorithm by running the pulse sequences that implement the circuit in Figure 8-13, for both the  $a = 11$  and  $a = 7$  cases. Here it is also apparent that the first virtual qubit always returns the same result and therefore in principle can be omitted. However as before we retain this as an extra qubit of precision on the QFT outcome and use it as an indicator for experimental errors.

The measurement and feedforward, hiding of ions before each intermediate measurement, and combining the measurement results to a QFT output are all carried out in the same manner as for the order-finding experiment. The one main difference is that the  $7^2 \bmod N$  pulse sequence with the 2 Fredkin gates and the associated hiding/unhiding requires decoupling the hidden qubits completely from the computational space, not just hiding them from measurement. This is accomplished by hiding qubits in state  $D$  to  $S'$  and qubits in state  $S$  to  $D'$  (see Figure 8-2). The transitions needed to implement hiding are also implemented by global pulses and addressed refocusing pulses, and thus a large number of pulses are required for the 8 hiding and unhiding operations within the  $7^2 \bmod N$  sequence. The hiding pulses we implemented have been manually optimized to remove redundant pulses given knowledge of the initial state of the sequence (but not making any assumptions about the outcomes of the Fredkin gates). The manually optimized hiding sequence is given in Appendix D.

The outcome of the Shor’s algorithm as a result of running the pulse sequences for the  $a = 11$  and  $a = 7$  cases are shown in Figure 8-15. The QFT output for  $a = 11$  agrees with the ideal outputs for order = 2 with a classical fidelity (sso) of 91% and a trace distance of 86%. The QFT output for  $a = 7$  agrees with the ideal outputs for order = 4 with a classical fidelity of 93% and a trace distance of 89%. Obtaining the order  $r = 2$  in the  $a = 11$  case gives  $\gcd(11^{2/2} \pm 1, 15) = 3,5$ , the prime factorization of 15. Similarly for  $r = 4$  in the  $a = 7$  case,  $\gcd(7^{4/2} \pm 1, 15) = 3,5$ .

We simulate Shor’s algorithm for the cases  $a = 11$  and  $a = 7$ , using the same parameters and methods as for the order-finding. Table 8.7 lists the sso and trace distances for the population outputs of Shor’s algorithm.

To evaluate the effects of various error sources on the algorithm output and fidelity, we produce an error budget for the experiment by simulating the pulse sequence with only one error source turned off at a time. Table 8.8 lists the major error sources and

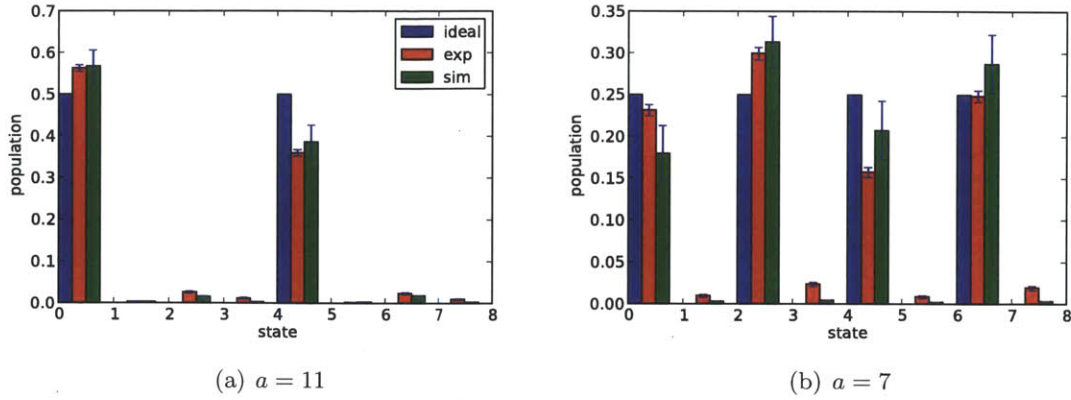


Figure 8-15: State populations measured as the output of Shor's algorithm on  $N = 15$  for (a)  $a = 11$  and (b)  $a = 7$ .

case	data	sso	trace distance
$a = 11$	exp	91.1(4)	86.0(7)
$a = 11$	sim	93(2)	86(7)
$a = 7$	exp	92.6(4)	88.8(7)
$a = 7$	sim	94(1)	81(3)

Table 8.7: Measured and simulated fidelities (sso and trace distance) of the QFT output of Shor's algorithm for  $N = 15$ ,  $a = 11, 7$ .

their contributions to the infidelity in the population output of the QFT at the end of the algorithm. The results suggest that errors due to hiding, mostly in the partial measurement step, can be a major contribution to measuring the wrong order at the end of the QFT step. In the Shor experiments we did not perform state tomography after each pulse or at the end of the sequence, but this information is available from the simulation. The final state fidelity caused by each error source is summarized in Table 8.9. In particular, in Table 8.9(a) the error sources were turned ON one at a time; in Table 8.9(b) the error sources were turned OFF one at a time and the infidelities listed is the difference between this and having all the error sources on. (The same two approaches were done for the QFT and no difference in the resulting error budget was found.) In the first case, the simulation results suggest that addressing and dephasing are the major contributions to the total gate infidelity, similar as what was found for the QFT. However, just turning off addressing error while leaving all other errors on has little effect. This suggests that in the Shor pulse sequences, there is some coupling between error sources, and the effect of any particular error source would be difficult to isolate.

The obtained population fidelities are relatively high compared to the fidelities of some of the sub-components of this pulse sequence. Each Fredkin gate is estimated to have a population fidelity of around 80% (obtained by comparing the measured and expected state population after applying the gate to one fixed input state). Each MS gate has a fidelity of 94%, determined by measuring the population and parity of GHZ states. Based on these estimates, the entire algorithm fidelity should only be around 50%. The unexpected high



	sso	trace dist
all errors	5.2 +/- 0.4	8.2 +/- 1.6
hiding	4.1 +/- 1.2	2.0 +/- 0.6
dephase	0.8 +/- 0.1	3.7 +/- 0.7
addressing	0.4 +/- 0.0	0.9 +/- 0.2
intensity	0.3 +/- 0.0	0.4 +/- 0.1
specmode	0.2 +/- 0.0	0.5 +/- 0.1

Table 8.8: Error budget on final population measurement for the Shor  $a = 11$  pulse sequence. Errors are derived from a simulation of 24 Monte-Carlo instances.

(a)

	jozsa fidelity	trace dist ( $\rho$ )	sso	trace dist ( $p$ )
all	39.8 +/- 17.6	63.1 +/- 12.6	26.4 +/- 12.3	26.6 +/- 12.4
dephase	29.3 +/- 16.4	30.6 +/- 7.8	17.0 +/- 8.9	15.2 +/- 8.2
addressing	0.7 +/- 0.6	0.6 +/- 0.2	4.5 +/- 4.1	3.2 +/- 2.5
hiding	2.2 +/- 1.5	1.7 +/- 0.6	5.5 +/- 3.8	0.2 +/- 0.1
intensity	1.2 +/- 0.9	1.0 +/- 0.3	2.5 +/- 1.5	0.4 +/- 0.3
specmode	0.7 +/- 0.5	0.8 +/- 0.3	1.1 +/- 0.7	3.0 +/- 1.7

(b)

	jozsa fidelity	trace dist ( $\rho$ )	sso	trace dist ( $p$ )
all	32.0 +/- 14.0	56.3 +/- 12.2	21.5 +/- 12.0	25.9 +/- 10.4
addressing	9.9 +/- 0.0	31.5 +/- 0.0	8.3 +/- 0.0	8.3 +/- 0.0
dephase	6.5 +/- 10.2	25.3 +/- 15.3	2.8 +/- 4.3	4.5 +/- 6.0
intensity	0.4 +/- 0.1	6.2 +/- 1.2	0.3 +/- 0.1	2.6 +/- 1.3
specmode	0.4 +/- 1.8	6.1 +/- 6.8	0.3 +/- 1.2	3.7 +/- 4.2
hiding	0.0 +/- 0.0	1.3 +/- 0.0	0.0 +/- 0.0	0.1 +/- 0.0

Table 8.9: Error budget on the final state fidelity for the Shor  $a = 11$  pulse sequence. Errors are derived from a simulation of 24 Monte-Carlo instances. In (a), error sources are turned ON one at a time, and in (b), error sources are turned OFF one at a time.

fidelities can be linked intuitively to the nature of the Kitaev measurement: the ordering of the virtual qubits is such that the least significant bit is measured first. In the case that the expected order is 2 (for both  $\pi_2(y)$  in the order-finding and  $11 \bmod N$  in Shor's algorithm),  $\pi^2$  and  $\pi^4$  are both equivalent to the identity operator, so both the first and second virtual qubits should return 0. The final permutation should generate the state  $|0\rangle + |1\rangle$  for the virtual qubit. What happens when we have errors that accumulate over time? Measurement on the first virtual qubit will return 0 with high probability, since it is early in the pulse sequence. The second virtual qubit will similarly return 0 upon measurement, with slightly lower probability. This is supported by our data: the intermediate population measurements follow this pattern. Finally the last virtual qubit suffers from all errors during the full length of the sequence; but here the expected qubit state is  $|0\rangle + |1\rangle$ , and if only the population and not the density matrix is measured, then the measurement result does not distinguish it from a mixed state.

Thus in the order = 2 case, the measurement result of  $|000\rangle + |100\rangle$  is relatively robust against noise that accumulates monotonically through the sequence. The order = 4 permutation is similar. The order = 1 permutation should actually be more susceptible to noise and result in a lower fidelity, because the final virtual qubit is expected to be in the state  $|0\rangle$ , incompatible with a mixed state. Finally in the order = 3 case, every virtual qubit follows a non-trivial sequence that performs a permutation, so its fidelity is the lowest of all. This logic explains the general trend of both the measured and simulated fidelities for the order-finding algorithm in Section 8.4, as well as the relatively high fidelities for Shor's algorithm despite the lengths of the pulse sequences we used.

The primary difference between this demonstration of Shor's algorithm using trapped ions and some of the recent results using other systems (including photons and superconducting qubits so far) is that the current implementation does not do any compiling that assumes a specific form of the controlled-  $a \bmod N$  operations. For example, in the  $a = 11$  case, we do not discard the second and last qubit of the bottom 4-qubit register even though they're entirely unaffected by the permutation operation. We do make use of the knowledge that no choice of  $a$  for  $N = 15$  requires more than 2 order operations, so only 3 virtual qubits are used to store the QFT output. If other orders were possible and more bits of precision on the QFT is needed, we could easily append additional virtual qubits to the beginning of the circuit, without having to use additional real qubits. Thus the Kitaev method that we implemented is completely general and can be applied to realize Shor's algorithm for any  $N$  up to the maximum value allowed by the bottom storage register.

In ion traps, fast measurement and recoling enabled the use of feedforwarding as a control technique, which we have demonstrated. The ability to perform feedforwarding is key in implementing quantum error correction which is required for fault-tolerant quantum computing. Thus more important than realizing Shor's algorithm, our experiment here suggests the feasibility of a future ion trap fault-tolerant quantum computer.

## 8.6 Summary and future work

The successful demonstration of various quantum algorithms on up to 5 ions, and the ability to evaluate and predict the fidelity of such algorithms, represent a small milestone for trapped ion quantum computing. These demonstrations show that we have excellent quantum control over the physical system, even as the complexity of some of these algorithms are pushing the capabilities of the current experiment. In addition, the outputs of our numerical simulations were in reasonable agreement with experimental results, indicating that noise sources in the experiment were well-understood, and this knowledge could be used to design and evaluate more complex algorithms.

However, much work remains to be done in the near future. As the next step, we might consider how one might factor  $N = 21$  using the same circuit and methods that we have implemented. Factoring 21 requires 6 ions: 1 for the Kitaev qubit and 5 to store the number in the bottom register. For a choice of  $a = 11$ , the order of  $a^x \bmod N$  is 6; since this is not a simple power of 2, we would need to carry out all powers of the permutation to get sufficient digits of precision on the QFT output. This requires  $2\lceil\log_2(N)\rceil$ , or 10, controlled- $a^x \bmod N$  operations. This is a factor of 5 more than that required to factor 15, so we estimate that  $\sim 500$  pulses might be required, with a total time of  $\sim 50$  ms. The duration of such a pulse sequence is comparable to the coherence time, and the cumulative effects of all error sources will likely make the total algorithm fidelity very low.

How do we overcome the limits on algorithm fidelity? Higher fidelity gates and operations, especially the MS gate on large qubit registers, will be important in improving overall algorithm fidelity. Noise-insensitive gates, such as those which generate states in decoherence-free subspaces, would be preferable. Other gate schemes such as ultrafast gates using pulsed lasers [CMQ<sup>+</sup>10] and microwave gates requiring no narrow-linewidth lasers [OWC<sup>+</sup>11] are being investigated. Better pulse sequence optimization might help shorten the lengths of the sequences necessary to implement operations and perhaps design sequences that are naturally robust against errors. Clearly technical sources of error such as addressing error, dephasing caused by magnetic field fluctuations, as well as imperfect control leading to systematic errors are major contributions to infidelity. While some technical aspects can be improved in principle, a degree of imperfection will always be present in practice. Thus quantum error correction is key to overcoming such practical noise sources and achieving fault-tolerance. The basic principles of quantum error correction has been demonstrated with ions [CLS<sup>+</sup>04, SBM<sup>+</sup>11].

Apart from improving gate fidelities, the number of ions in a 3D Paul trap likely presents the next limitation to scaling up algorithms with trapped ions. With the 3D trap architecture, all ions in the register must be trapped in a linear chain to enable addressing. The number of motional modes for  $n$  ions in a chain is  $3n$ , so as the trap gets denser, it becomes increasingly difficult to cool all modes. Also, recent experiments [MSB<sup>+</sup>11] have shown that the gate fidelity due to collective noise sources scale as  $1/n^2$ , and even for as few as 10 ions, the fidelity of GHZ states created with a single entangling gate is reduced to  $\sim 60\%$ . The prospects for scaling the trap itself, by moving to a 2D architecture and incorporating new capabilities such as shuttling, is the subject of much ongoing work in the community including the first two parts of this thesis. The ability to trap an arbitrary number of ions

in multiple zones, combined with the capability of shuttling and enabling a small subset of the ions to interact at a time, would overcome some of the limitations of the 3D Paul trap. The challenge then is to map our existing algorithms onto the new architecture, using and understanding the new capabilities in order to realize algorithms both more powerful and more resistant against known error sources. The principles of TIQC-SPICE can be generalized to include such systems, for example to design pulse/operation sequences optimized for a 2D trap where the contributions of various noise sources might be different. Projects are underway within the ion trapping community to develop such tools for 2D traps with complex architectures [QMP12].



## Chapter 9

# Integration with optics

While a large-scale quantum computer would enable the implementation of quantum algorithms with exponential speedups, the required number of qubits needed to be competitive with classical computers still represents a seemingly insurmountable challenge today. On the other hand, applications such as secure quantum communication, which require fewer qubits, can potentially enable new capabilities in the near term. Taking advantage of the photonic qubit’s ability to travel long distances without loss over commercially available fibers, basic quantum communication protocols have already been demonstrated [JSW<sup>+</sup>00, UTSM<sup>+</sup>07].

Quantum networks is an exciting potential application of quantum information processing. In such systems, it is often proposed that a “stationary qubit” such as trapped ions form the nodes of the network, performing computations and storing results, while “flying qubits” such as photons carry information between nodes of the network [Kim08]. Such a hybrid quantum system can overcome the technical size limitations of individual systems. However, efficient conversion between information stored in ions and transmitted via photons is needed.

Integration of optical elements such as mirrors, lenses, and fibers with trapped atoms and ions to facilitate such conversion has been intensively pursued in recent years. Highly efficient coupling between optical cavities and atoms have been demonstrated [BBM<sup>+</sup>07], but the same system poses a challenge for ions: optical elements typically consist of dielectrics, which can contain stray charges generating stray electric fields that disturb the ion’s trapping potential.

In this chapter, we explore three potential approaches to couple photons with trapped ions. Section 9.1 describes the fabrication and characterization of a surface-electrode ion trap integrated with a high-reflectivity mirror, which is designed to be part of a high-finesse cavity capable of achieving strong-coupling<sup>1</sup> between ion and photons. Section 9.2.1 describes a transparent surface-electrode ion trap, which enables light collection through the trap with an integrated photon detector. Section 9.2.2 is a proposal for integrating surface-electrode ion traps with superconducting nanowire single-photon detectors.

---

<sup>1</sup>The strong coupling regime refers to the case where the rate of excitation exchange between the ion and the cavity (vacuum Rabi flopping rate) is greater than the geometric mean of the spontaneous decay rates of the ion and the cavity [CZKM97].

## 9.1 Mirror trap

The design of a planar ion trap fabricated on a high-finesse mirror, plus a bulk mirror to form a cavity, was first proposed in [Lei09]. It was shown via CPO simulations that the presence of a hole in the center electrode does not significantly change the trapping potential, and some suggested parameters were presented which would enable the resulting system to achieve strong-coupling. The major questions are, (1) whether the presence of the exposed dielectric directly below the ion will cause too much anomalous heating and laser-induced charging to allow for stable trapping, and (2) will it be possible to fabricate the surface-electrode ion trap without significantly degrading the quality of the mirror. This section will mainly focus on the experimental approach we used to investigate the first question. The second question is addressed in Ref. [HWS<sup>+</sup>11] and in future theses.

### 9.1.1 Trap design and fabrication

The design of the planar trap is the standard 5-rod design described in Chapter 3, except for a circular aperture in the center ground electrode exposing the mirror to the ion. The width of the notched part of the ground electrode is  $150\ \mu\text{m}$ , and the aperture is  $50\ \mu\text{m}$  in diameter. This gives an expected ion height of  $165\ \mu\text{m}$  from CPO simulations. Temperature sensors identical to those used in [LGL<sup>+</sup>08] were mounted on the trap to measure the increase in trap temperature due to the RF drive.

Several traps were fabricated on 1.6 mm-thick fused silica substrates that have a highly reflective dielectric coating optimized for light near 422 nm. The coating consists of alternating layers of  $\text{Ta}_2\text{O}_5$  and  $\text{SiO}_2$  for a total thickness of  $2.2\ \mu\text{m}$ , and was deposited by Advanced Thin Films (ATFilms) using ion beam sputtering. The general fabrication procedure is as follows, with two slight variations as marked by (\*).

#### Wafer cleaning

- submerge the substrate in a beaker of acetone. Clean the substrate by mechanical rubbing with cotton swabs. Rinse with isopropanol alcohol (IPA) and blow dry. Submerge again in IPA and clean by rubbing, then rinse with IPA and blow dry. Check with a light probe for remaining dust and imperfections.
- (\*) further clean the substrate with a Piranha solution (3:1  $\text{H}_2\text{SO}_4$ : $\text{H}_2\text{O}_2$ ), submerging for 10 minutes.

#### Lithography

- prebake at 110C for 5 min
- spin NR9-3000PY photoresist at 3000rpm for 60 s
- bake at 110C for 5 min
- (\*a) Exposure with a chrome mask for 2 min at  $3300\ \mu\text{W}/\text{cm}^2$
- (\*b) Expose with Heidelberg direct-write system at 18mW, 50%, 4x energy mode
- bake at 110C for 2 min
- develop the photoresist in RD6 developer for 17 s
- rinse with deionized water for at least 30 s, blow dry



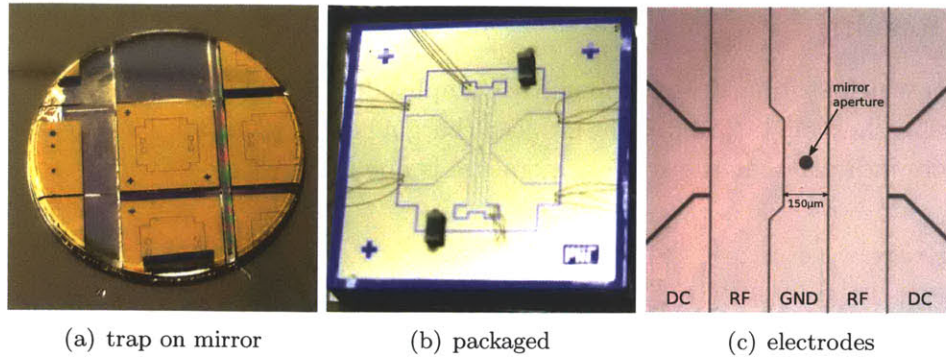


Figure 9-1: (a) Photo of MirrorTrap fabricated on a high-reflectivity mirror after diesawing. The yellow color is due to the protective layer of photoresist applied before cutting. (b) Photograph of MirrorTrap packaged in a CPGA. RuO temperature sensors are soldered onto two of the electrodes. (c) Diagram of trap electrodes showing the aperture in the ground electrode. The ion is trapped  $169 \mu\text{m}$  above the center of the aperture.

### Deposition

- E-beam evaporate 10 nm of Ti followed by 400 nm Ag, at a rate of  $5 \text{ \AA/s}$
- soak in acetone for 25 min until liftoff is completed
- rinse in methanol, blow dry

### Dicing and packaging

- spin NR9-3000PY photoresist and bake at 110C for 3 min
- use the same photoresist to glue the substrate onto a silicon wafer
- cut into traps using a diesaw and a diamond blade, making 3 cuts of up to 0.7 mm deep each.

The trap is typically stored under a protective layer of photoresist in a gelpak in normal atmosphere conditions. Before packaging, the trap is rinsed with acetone to remove the photoresist. Traps are glued in a CPGA with epoxy (Varian TorrSeal) and wirebonded using gold wire. RuO resistors acting as temperature sensors were mounted on the trap electrodes using conductive epoxy (Epo-Tek H20E). Figure 9-1(b) shows a picture of the finished trap.

Eight traps were fabricated, packaged, and tested in the cryostat. Some variation in the processing of these traps were as indicated above: substrates were cleaned with and without the Piranha solution step, and exposure was performed using either a UV light source and a direct-contact chrome mask (Tamarack Scientific) or a direct-write lithography system (Heidelberg Instruments uPG 101). All combinations of these options were done on at least one trap each. Of the 8 traps fabricated, many of them exhibited RF breakdown at voltages below that necessary to trap. Only two of the traps survived breakdown testing: both were not cleaned with the Piranha solution, but one was exposed with a mask and another with the direct-write system. There is insufficient data to determine the root cause of low breakdown voltages.

### 9.1.2 Results

The traps which sustained high enough RF amplitudes to enable trapping were tested for trap stability, heating rate, and laser-induced charging in the cryostat. The surface-mounted thermistors indicate that the trap temperature is about 15 K.

#### Heating rate

Measurements of heating rate are done at a secular frequency of  $\omega_z = 2\pi \times 0.7$  MHz. and the lowest heating rate observed is  $109 \pm 44$  quanta/s. The spectral density  $S$  of the fluctuating fields driving the heating is  $\omega_z S(\omega_z) = (3.6 \pm 0.4) \times 10^{-10}$  V<sup>2</sup>/cm<sup>2</sup>, which is comparable to or lower than traps of similar dimensions operated at room temperature [ESL<sup>+</sup>07]. The result is about an order of magnitude higher than the lowest heating rate obtained previously for silver traps with the same geometry at cryogenic temperatures, but without the aperture and the dielectric mirror coating [LGA<sup>+</sup>08].

The slightly higher heating rate is consistent with the observation that increased oxide layer thicknesses contribute to generally higher heating rates in our aluminum traps (Section 5.2.2). Non-contact friction measurements using cantilevers have similarly observed about an order of magnitude increase for a bare fused silica substrate relative to a gold surface [SMS<sup>+</sup>01]. However, relatively few heating rate measurements were made in this trap and only over one day. Considering the large variations observed in our other traps over different days and varying parameters such as RF amplitude and compensation voltages, more controlled experiments are required to determine if the heating rate observed here is influenced by the exposed dielectric.

#### Laser-induced charging

We studied laser-induced charging in this trap by performing charging tests with light at 405 nm, 460 nm, and 674 nm, as described in Chapter 6. About 200  $\mu$ W of power focused to a  $\sim 50$   $\mu$ m radius spot is incident at grazing angle across the trap surface under the ion to simulate the effect of misaligned laser beams. The ion displacement, as a result of charge buildup, can be measured via the induced micromotion as the ion is displaced from the node of the RF field and quantified in terms of the adjustment of the trap voltages of the MID electrode required to compensate this effect. For the wavelengths we studied, the required changes in the DC voltages after 10 min of continuous exposure due to charging are at the level of 5-50 mV. The strongest effect is observed with light at 405 nm and corresponds to an induced field at the ion location of about 20 V/m. With no excess laser light incident on the electrodes and mirror, trapping is observed to be stable without the need for adjustment of DC voltages over a time span of 1 hr.

#### Mirror quality

The second major question to be addressed by this work is whether the fabrication of the trap on the high-reflectivity mirror damages the mirror itself. This is evaluated by performing ring-down spectroscopy in a near-confocal Fabry-Perot cavity setup before and after

fabrication. The testing is done on a substrate that has an array of apertures, fabricated in parallel with the trap using the same recipe. Before fabrication, we find the cavity losses to be in agreement with the vendor specifications of a 45 ppm transmission coefficient and scattering and absorption losses of 25 ppm. The average increase in losses after fabrication is  $130\pm 10$  ppm. Further details of the testing setup and measurement method will be described in a future thesis.

Although the geometry of our design does not result in increased detection of ion fluorescence, we can demonstrate the principle of light collection using the mirror located directly below the ion. By adjusting a bulk lens located outside the cryostat to select the imaging plane, either the primary image of the ion or the secondary image as reflected by the mirror can be imaged. By measuring the displacement of the bulk lens outside the cryostat, the ion height was measured to be  $169\pm 4$   $\mu\text{m}$ , in good agreement with the CPO prediction.

### 9.1.3 Summary

We have demonstrated integration of ion traps and optics by microfabricating a planar ion trap on top of a high-reflectivity mirror. We showed that stable trapping of a single ion is possible, with lifetimes on the order of hours and compensation voltages stable for at least an hour, despite the close presence of exposed dielectrics. Laser-induced charging effects are minimal, and heating rates are in the typical range of other traps tested at cryogenic temperatures. The quality of the mirror is not significantly reduced by the fabrication procedure.

The natural next step of this design is to form an ion-cavity system by adding a second concave mirror above the trap. The low ion height enables a cavity length of potentially less than 1 mm, potentially a small enough mode volume to achieve strong coupling. Furthermore, our design and fabrication procedure is automatically self-aligning between the ion and the mirror, and can be scalable to larger number of hybrid ion-optics systems.

## 9.2 Photon detectors

The fluorescence collection of trapped ions is a major source of inefficiency in many experiments. While research into integrating elements such as mirrors, MEMS elements, and control electronics with ions are actively pursued, the collection of photons from trapped ions for quantum state detection is mostly implemented with bulk optics. This limits the available collection angle and the total collection efficiency. In a recent experiment demonstrating quantum teleportation, the low detection efficiency was the primary cause of a loss of fidelity in the experiment [OMM<sup>+</sup>09]. In the experiment described in Chapter 8, state detection with a CCD camera takes at least 8 ms, which is significantly longer than most of the pulse sequences.

In recent years several approaches have been investigated to increase the detection efficiency. A bulk ion trap was recently integrated with a parabolic bulk mirror which demonstrated increased light collection efficiency with 25% solid angle [SCK<sup>+</sup>11]. Fresnel lenses have also been combined with ion traps to improve imaging resolution [JSN<sup>+</sup>11]. For planar traps, the trap itself naturally limits optical access significantly; this has been overcome

with options such as a microfabricated concave reflective mirror on the trap [MVL<sup>+</sup>11] and integrated optical fibers [VCA<sup>+</sup>10].

In the following sections, we explore two additional options for increasing the photon collection efficiency in microfabricated ion traps. Section 9.2.1 describes the design and implementation of a transparent ion trap integrated with a large-area photodetector. Section 9.2.2 describes a proposed design for integrating a superconducting photon detector with a superconducting ion trap.

### 9.2.1 Indium-tin oxide traps

Indium-tin oxide (ITO) is a well-known transparent conductor with a wide range of applications such as touch displays and solar cells. We fabricate a transparent ion trap with ITO as the electrode material, and place a large-area photodetector directly below the trap, so that fluorescence of the ion is collected through the trap by the photodiode. If the photodiode size is large compared to the thickness of the trap and the ion height, then nearly  $2\pi$  of solid angle is collected.

A large variety of photodiodes with various operating environments and performances are commercially available. For our demonstration we chose to use a widely available PIN photodiode (Advanced Photonix PDB-C613-2) with a 9 mm<sup>2</sup> active area. The transparent trap is fabricated and packaged with the photodiode in place of a spacer in a standard CPGA. Here we describe the fabrication and packaging procedure, the optical characterization of the deposited ITO, and a first signal of an ion cloud collected by the photodiode.

#### Fabrication

All of the ITO traps fabricated and tested used the standard design described in Chapter 3, with an expected ion height of 100  $\mu\text{m}$ . Several different methods of deposition and lithographic patterning of ITO traps were attempted. Deposition methods included e-beam evaporation and sputtering. Initial attempts were made to pattern the traps on commercial ITO films (Sigma-Aldrich) deposited to 100 nm on glass slides, using photolithography and wet etch (36.46% HCl). The resulting patterns had uneven edges and sharp corners which will likely result in breakdown when RF voltage is applied. E-beam evaporation resulted in opaque films. Possibly further refinement of the fabrication recipe would have mitigated these issues. The successful recipe we used for fabrication of ITO traps consisted of the usual photolithography patterning step, followed by RF sputtering with argon gas at a rate of 0.5  $\text{\AA}/\text{s}$ . The resist pattern is transferred to the ITO via lift-off.

Separate sputtering runs were performed on three batches of ITO traps. The first batch had a measured sheet resistance of 120  $\Omega/\text{square}$ , which would result in a total resistance of 3200  $\Omega$  for our RF electrode with cross-sectional area of 400 nm  $\times$  150  $\mu\text{m}$ . This would likely result in too high RF power dissipation on the trap for practical purposes. Subsequent batches had sheet resistances of 50  $\Omega/\text{square}$  and 20  $\Omega/\text{square}$  respectively. To increase the conductivity of the RF paths only, an additional lithography and deposition step was carried out to deposit a gold layer on the RF electrode. First samples had 50-70 nm of gold deposited, which should increase the conductivity to be comparable to typical traps.

Trap	ITO thickness (nm)	RF gold thickness (nm)	comments
Ia	450	70	stable trapping; charging data
Ib	450	none	no ions trapped
II	400	50	no ions trapped
III	400	4	photodiode data
IV	150	4	RF shorted

Table 9.1: Summary of fabrication parameters and results for all tested ITO traps.

Subsequent samples had 4 nm of gold deposited, which should still be approximately 60% transparent, but significantly increase the conductivity of the RF path [SNTK03].

A list of traps and fabrication-related characteristics are listed in Table 9.1. Only traps with the successful recipe, packaged and trapping attempted are included. Figure 9-2 shows a few representative microscope and SEM images of fabricated ITO traps. Figure 9-3 shows packaged photos of trap Ia, in which stable trapping was demonstrated and charging data was collected.

For tests of photon detection, a standard silicon PIN photodiode was mounted in the CPGA directly below the trap. The trap is glued on top of the photodiode with Krazy glue (cyanoacrylate) on the edges and the center area is kept clear. The trap is wirebonded to the CPGA with aluminum wire.

## Characterization

The optical transmission of some of the ITO traps were measured with a PIN photodiode at room temperature with a 422 nm light source, by comparing the measured light intensity with and without the presence of a trap in the beam path. The laser source is modulated by either switching an AOM on/off or using a mechanical chopper in the beam path, in order to distinguish the signal from background light. The measured transmission with 400 nm ITO and 5 nm of gold, averaged over an area of 4 mm<sup>2</sup> was found to be 57% and 64% for two different samples. A third sample with 150 nm of ITO and 4 nm of gold had a transmission of 76%. These values are slightly lower than commercially available ITO films (170 nm, 90%) [KPH<sup>+</sup>99]. Similar to the sheet resistance, the transmission varies significantly over traps fabricated in different batches. Better control of the fabrication process should result in more consistent optical transmission.

PIN photodiodes are not designed to operate at cryogenic temperatures; at low temperatures, carrier freeze-out is expected to significantly degrade the detection efficiency. To characterize this, we installed a photodiode inside the cryostat without a trap or an amplifier circuit and measured its response to the 422 nm laser as a function of temperature. At room temperature, the responsivity of the photodiode is 0.1 A/W. At ~40 K, the responsivity of the photodiode drops rapidly such that at 4K the response is 0.01 A/W. Including geometric factors, the combined detection efficiency of our transparent trap and photodiode system is about 1%, which is comparable to the existing bulk optics and PMT system. Local heating of the photodiode should improve this efficiency.



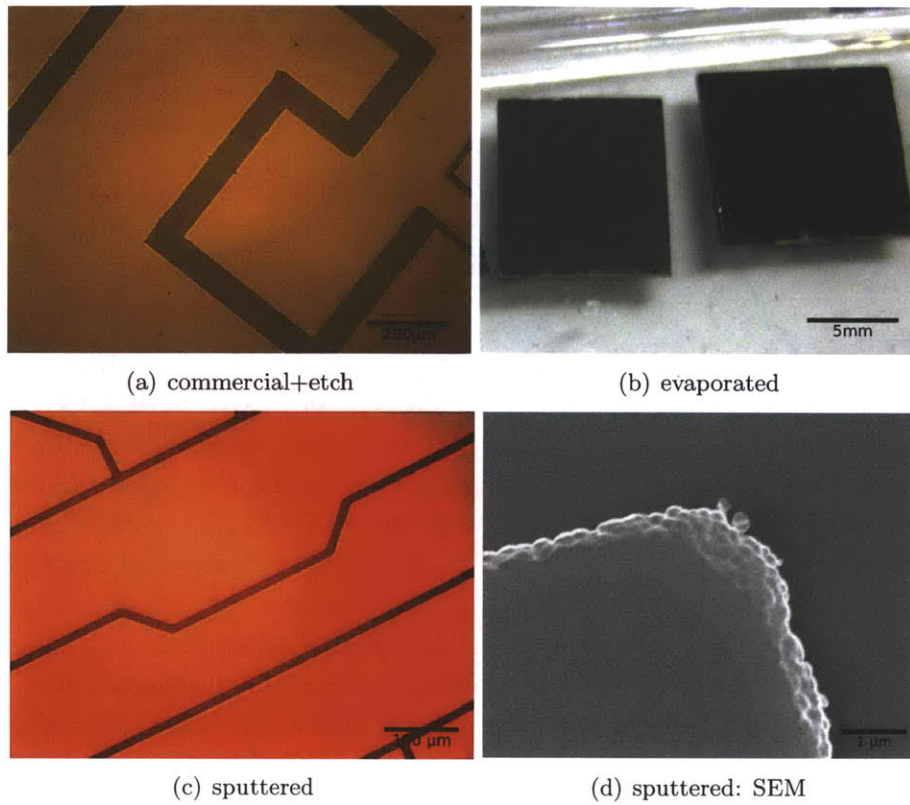


Figure 9-2: (a) Microscope image of an ITO trap made by wet-etching a commercially available ITO film on glass slides. Sharp uneven electrode edges are visible. (b) Photo of evaporated ITO films which are opaque. (c) Microscope image of sputtered ITO trap showing smooth edges and no sharp corners. (d) SEM image of sputtered ITO trap.

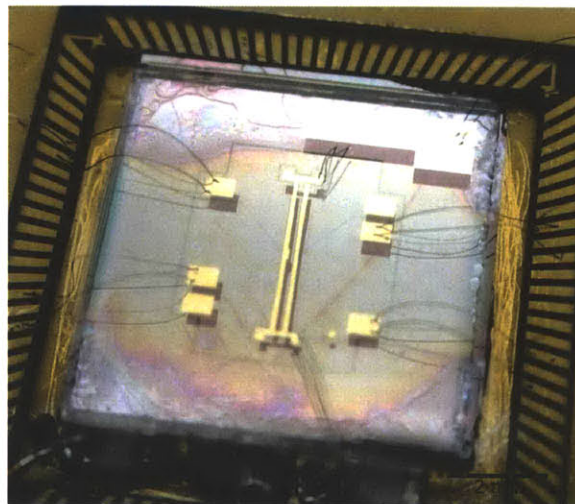


Figure 9-3: Photograph of trap ITO Ia mounted in a CPGA for testing.

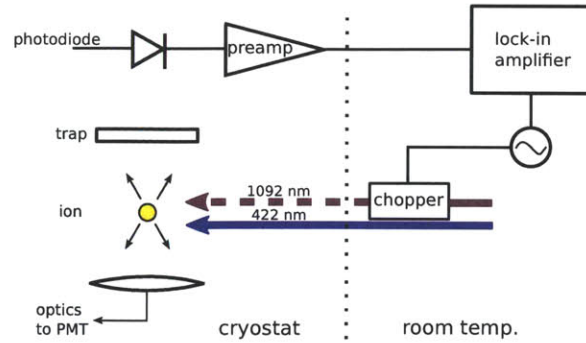


Figure 9-4: Ion fluorescence detection setup.

A single trapped ion scatters  $\sim 10^7$  photons/s. With a solid angle of  $2\pi$ , this corresponds to 2 pW of power, or 24 fA of current on the photodiode with the measured responsivity at 4 K. A two-stage preamplifier with a transimpedance gain of  $\sim 3 \times 10^6$  was added to the 77 K stage of the cryostat to amplify the ion signal from the photodiode. The resulting signal is detected using a lock-in amplifier.

The first ITO trap, with 70 nm gold deposited on the RF electrode, was tested in the cryostat and able to load and trap single ions reliably with lifetimes of several hours. Measurements of light-induced charging, made by grazing the trap surface with the 405 nm and 460 nm lasers for up to 10 min, showed no significant charging effects.

A subsequent trap with 400 nm of ITO and 4 nm of gold was used to test detection of the ion signal with the photodiode. A large cloud of ions was continuously loaded at 77 K; the PMT measured the cloud's fluorescence to be 150,000 photons/s, whereas typical measured fluorescence of a single cooled ion would be 50,000 photons/s. A mechanical chopper was inserted in the beam path of the 1092 nm IR laser such that the ion signal is modulated at 237 Hz, for lock-in detection. The photodiode is also sensitive to the IR laser, so the signal from the ion cloud is determined by subtracting the background of the lock-in signal without an ion cloud present. Figure 9-4 shows the detection setup.

Figure 9-5 shows an example of the raw lock-in signal with and without an ion cloud, as well as histograms of the signal, averaged over 5-10 minutes. In the three sets of data taken, the ion signal from the lockin detection with background subtracted was measured to be 0.11(7), 0.09(8), and 0.15(6) V. While the errorbars are large, the signals are beyond 1 sigma from zero, indicating that signal from the ion is detected by the photodetector. Further work is needed to quantify and compare the signal with expected values. Unfortunately the trap was shorted before further testing with stably trapped ions at 4 K could be performed.

## Summary

Traps made of ITO, a transparent conductor, were fabricated and tested in a 4 K cryostat. Ions were trapped stably and minimal sensitivity to laser-induced charging was observed, indicating that ITO is a viable material to use for microfabricated traps. A standard silicon PIN photodiode was placed under the trap, and preliminary evidence of ion fluorescence detection was observed.



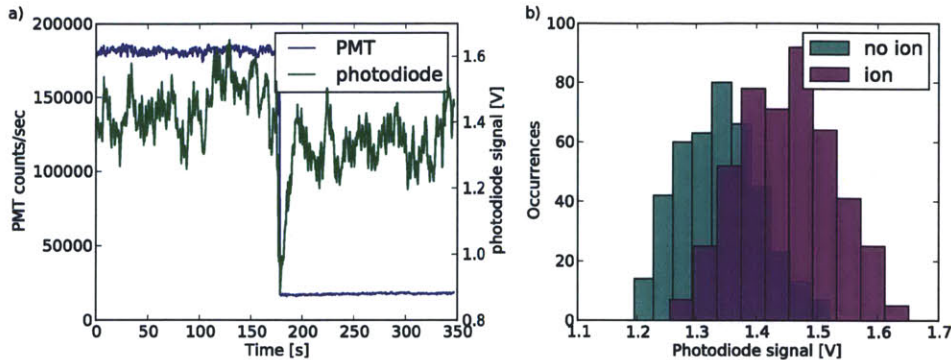


Figure 9-5: (a) Example of photodetector’s lock-in signal (green) and PMT signal (blue) with (up to Time = 175 s) ions and without (after Time = 175 s). (b) Histogram of photodetector signal with and without ions.

The cryogenic operating environment used in our experiment is a significant disadvantage for the silicon PIN photodiode, due to its decreased detection efficiency at low temperatures. However, our proposed scheme of a transparent trap coupled with a photodetector is not restricted to a cryogenic environment, and the system could equally well be operated in a room-temperature setup. Alternatively, many other types of photodiodes exist with better performance at cryogenic temperatures and could be integrated with the transparent trap in the same manner. Photodetectors with internal gain mechanisms would be especially beneficial. For example, Visible Light Photon Counters (VLPCs) have been demonstrated with a quantum efficiency of 88% at 694 nm at 4 K [MKH09a], but are unfortunately not commercially available. Another alternative is CMOS avalanche photodiodes (APDs), which only show a  $2\times$  drop in quantum efficiency at 4 K [JCM<sup>+</sup>09]. Another disadvantage of the silicon PIN photodiode is that it does not allow for single-photon detection. In the next section, we describe a proposal to integrate ion traps with superconducting nanowire photon detectors, which has both high detection efficiency at cryogenic temperatures and single-photon resolution capabilities.

### 9.2.2 Superconducting nanowire single photon detectors

Single photon detectors which are capable of providing timing information on single photon events are crucial in optical quantum information applications such as quantum key distribution. A wide range of technologies in all stages of development exist [Had09]. Among them, superconducting nanowire single-photon detectors (SNSPDs) exhibit fast timing response, broad wavelength response, and high detection efficiency. The need for a cryogenic operating environment limits them to specific and demanding applications, and photon detection for trapped ions may be one possibility.

In this section, we propose to integrate microfabricated ion traps with SNSPDs directly on-chip. Section 5.3 presented a superconducting trap that is capable of sustaining the required currents to trap ions, has low heating rates, and demonstrated basic sensitivity to blue light. In pursuing integration of other devices with ion traps, a superconducting

photon detector a part of such a trap is a natural candidate. In the following we discuss the idea and investigate key issues including detection efficiency, fabrication, and device operation challenges.

### **Proposed implementation**

SNSPDs consist of a narrow (10-100 nm) meandering wire made with niobium nitride (NbN), patterned by electron-beam lithography. They are operated in a temperature range of 1.5-4 K, well below the typical superconducting transition temperature of NbN of above 10 K. The wire is biased to near its critical current. When a photon arrives at the wire, it triggers a local hotspot and causes it to transition to the normal state, resulting in a voltage spike which is then amplified and measured.

The largest size of a single SNSPD element reported is  $20 \times 20 \mu\text{m}$  [MFS<sup>+</sup>08]. A more typical size of the detector is  $10 \times 10 \mu\text{m}$  or  $3.3 \times 3.3 \mu\text{m}$  [DKR<sup>+</sup>09]. The size is limited primarily by the kinetic inductance of the device: a longer nanowire has greater kinetic inductance, which leads to a slower response time. Another limit to the size is constriction current due to fabrication and material defect: a single defect along any length of the wire lowers the critical current and detection efficiency of the whole device. However, the total detection area can be increased by interlacing multiple detectors in a meandering pattern [DKR<sup>+</sup>09].

The envisioned ion trap + SNSPD system consist of a single or an array of SNSPDs fabricated on the same substrate as the trap and centered directly below the ion. The trap geometry can either be modified to have a hole in the ground electrode to allow photons through, as in the MirrorTrap design described in Section 9.1, or be made entirely of a transparent material such as ITO. The vastly different dimensions of the SNSPD vs the ion trap presents a challenge. A possible fabrication scheme and estimates of system efficiency is described below.

### **Fabrication**

Superconducting ion traps are fabricated with Nb or NbN on a sapphire substrate. The typical minimum feature size of the traps is 5-10  $\mu\text{m}$  (gaps between electrodes). The Nb or NbN film is 400 nm thick. The fabrication process involves sputtering a Nb or NbN film onto the substrate, defining the pattern via optical lithography, then transferring the pattern by reactive ion etching (RIE) for 40 minutes. Each sample occupies a  $1\text{cm}^2$  area.

SNSPDs are fabricated with NbN and have typical feature size 10-100 nm corresponding to the width of the wires. The film is usually 4-12 nm thick (thinner film corresponds to higher detection efficiency). The typical fabrication process is described in e.g. [DKR<sup>+</sup>09] and is briefly summarized here. It begins by sputtering a thin film of NbN; the quality of the film must be high to minimize defects that can cause constrictions limiting the critical current. Contact pads, consist of 10 nm of Ti and 50 nm of gold, are first patterned using photolithography. Then the substrate is cleaned with TMAH, followed by spinning 90-100 nm of hydrogen silsesquioxane (HSQ) resist. Patterns are defined using E-beam lithography. The pattern is transferred to the NbN film by RIE for 1 min. Typical total

size of the sample is  $1 \text{ mm}^2$ .

To integrate ion traps and SNSPDs onto a single device, a more sophisticated fabrication process is likely necessary due to two major incompatibilities between these two devices: the orders of magnitude difference in film thickness, and the different patterning methods. In particular, e-beam lithography is slow and unsuitable for patterning large areas due to possible step mismatch, while optical lithography lacks the required resolution for 10-100 nm features.

One possibility is to fabricate a single or an array of SNSPDs on a  $1 \text{ cm}^2$  substrate via the standard method, followed by depositing or spinning a transparent material on top of which the usual ion trap fabrication procedure can be carried out. SU8 is a common resist used for such purpose. For example, a multi-layer point-Paul trap has been fabricated using SU8 as an insulating material between layers [Kim11]. Another option is to develop a process for performing e-beam lithography in a small area and optical lithography for the large area of the ion trap. Such a hybrid process has been demonstrated recently for CMOS devices and magnetic recording heads [GCF<sup>+</sup>09, Yan05]. A single-layer device would force the ion trap to have a thickness of 4-12 nm, which is too resistive to sustain the necessary RF amplitude required for trapping. However, as in the case of ITO traps, this can be mitigated by depositing additional metal on the RF electrodes.

## Geometry and efficiency

The typical ion height in a microfabricated ion trap is 75-150  $\mu\text{m}$ . The smallest traps demonstrated have an ion height of 30  $\mu\text{m}$  [VCA<sup>+</sup>10]. Assuming the ion emits photons isotropically, the solid angle covered by a single square SNSPD with side  $a$  is given by  $4 \arcsin(a^2/(4d^2 + a^2))$ , where  $d$  is the ion height. A 20  $\mu\text{m}^2$  detector with 40  $\mu\text{m}$  ion height would have a solid angle of 1.9%; a less demanding geometry consisting of a 10  $\mu\text{m}^2$  detector and 75  $\mu\text{m}$  ion height would have a solid angle of 0.14%.

One of the first reported quantum efficiencies of the SNSPD at 400 nm is 10% [SVG<sup>+</sup>03]. More recently, higher efficiency at 1550 nm has been achieved with cavity enhancement to 50% [RYD<sup>+</sup>06]. The SNSPD community have concentrated on enhancing the detection efficiency (DE) in the IR range due to telecommunication applications, but since blue photons have higher energy than IR photons, it should be relatively easier to make SNSPDs that have higher efficiency in the 400 nm region. For example, 50% DE at 400 nm without cavity should be possible [Naj11]. Note that there is large variation in DE in detectors fabricated with the same process due to the sensitivity of constriction current to single defects [KDY<sup>+</sup>07].

The  $\text{Sr}^+$  ion's Doppler cooling transition at 422 nm, for our experimental parameters, scatters  $\sim 1 \times 10^7$  photons/s. A typical imaging system consists of a PMT with quantum efficiency  $\sim 20\%$ , and the collection optics has a total efficiency  $\sim 6\%$ , determined by the NA of the first lens and losses of the optical elements. The total system efficiency is  $\sim 1\%$ . In practice we measure  $\sim 1 \times 10^5$  photons/s at the maximum scatter rate of the Doppler cooling transition and operate at half-detuning, giving us 48000 photons/s.

The intensity distribution of the ion's emission is isotropic. Along the quantization axis defined by the applied field, the emitted photons are circularly polarized. Since the SNSPDs

are polarization-selective, this introduces another factor of 1/2 in the total system efficiency. The polarization of the emission can be made linear by redefining the quantization axis (changing the direction of the applied bias field). The best possible case (20  $\mu\text{m}^2$  detector, 40  $\mu\text{m}$  ion height, 50% DE at 400nm), would result in 0.45% system efficiency, or 45000 photons/s, similar to our current system performance. A more conservative case (10  $\mu\text{m}^2$  detector, 75  $\mu\text{m}$  ion height, 20% DE at 400 nm) would result in only 0.07% system efficiency, or 7000 photons/s.

### Dark counts

Typical background counts for trapped ions with conventional optics/PMT is 5000 counts/s. The intrinsic dark count rate of SNSPDs can be very low (10 counts/s, [Had09]), but dark counts increases along with detection efficiency as bias current is increased. A greater concern is background due to laser scatter and trapping lasers at other wavelengths. SNSPDs are not meant to be wavelength-selective and typical detectors are very broadband [SVG<sup>+</sup>03]. Some possibilities exist for narrowing the wavelength bandwidth: adding a cavity, which would also enhance detection efficiency, or fabricate an interference-type filter on top of the detector.

Thus for “typical” operating parameters it appears that the signal from SNSPDs would not be competitive with conventional optics, and may in fact be barely above the noise floor to be detected. However there are a number of ways that the system efficiency can be improved. For example, adding a mirror on top of the trap would greatly enhance the collection solid angle. Making an array of multiple detectors would also increase the available detection area.

### Additional challenges

A number of issues remain in the practical realization of a scheme to combine microfabricated ion traps with SNSPDs, all related to the question of whether the large electric field needed for trapping is compatible with the electric signals needed to drive & detect SNSPD signals. Specifically, the electric fields of the devices could interact with several mechanisms:

- RF current creates AC field which reduces the critical current
- RF pickup on superconducting wires
- large electric fields created by DC/RF electrodes
- effect of (non-signal) light on superconductors: effect of bandgap, quasiparticle creation, thermal heating
- noise on DC/RF electrodes inducing currents on superconducting wires
- other technical noise (pickup on detection chain, 20dB gain electronics)

The expected net effect is fluctuation and uncertainty in the applied current to the SNSPD. Typically the applied current needs to be close to the critical current to maximize detection efficiency. To compensate for the fluctuation and maintain superconductivity in the absence of incident photons, the bias current would have to be decreased, at expense of

detection efficiency. Typically the SNSPDs need  $< 1 \mu\text{A}$  current noise to operate, and the detection electronics need  $> 1 \text{ GHz}$  bandwidth.

As an example, we estimate amount of RF pickup on a single superconducting wire from geometry and practical operating values. In the superconducting ion traps, the impedance of the RF electrode is dominated by the capacitance of the lines; typical capacitance of the trap is  $C = 8 \text{ pF}$ . With an RF frequency of  $\omega = 2\pi \times 36 \text{ MHz}$ , this yields an approximate impedance of  $Z_{RF} = \frac{1}{j\omega C} = -550j \Omega$ . Thus, for an RF amplitude of  $150 \text{ V}$  on the trap, the current amplitude is  $|I_{RF}| = 0.27 \text{ A}$  in the RF electrodes.

Modeling the RF electrodes as long, straight wires, the magnitude of the magnetic field at the position of the center of the trap (where the SNSPD would be located) is given by  $B = 2\frac{\mu_0|I_{RF}|}{2\pi r} \cos(\omega t)$ . The field direction is orthogonal to the plane of the trap. At the center of the trap,  $r = 125 \mu\text{m}$ , so we obtain  $B = 8.7 \cos(\omega t) \text{ G}$ .

The voltage induced in the SNSPD nanowire is then given by  $V_{\text{induced}} = -A\frac{dB}{dt} = A\omega(8.7 \text{ G}) \sin(\omega t)$ , where  $A$  is the area of the loop enclosed by the meander nanowire and the bias current supply. Approximating  $A$  by the active area of the SNSPD, taken to be  $15 \times 15 \mu\text{m}^2$  results in  $V_{\text{induced}} = 44 \cos(\omega t) \mu\text{V}$ .

Assume the width of the nanowire to be  $100 \text{ nm}$ , the impedance of the nanowire can be approximated by:  $Z_{\text{SNSPD}} = j\omega L_k = j14 \Omega$ . This is a lower bound on the impedance, assuming that the entire nanowire is in the superconducting state; however, during normal operation an SNSPD would contain a segment switched to the normal state, which contributes a significant resistance, as much as  $1.5 \text{ k}\Omega$  [Yan09]. Using the lower bound impedance yields an induced (AC) current on the superconducting nanowires of  $\sim I_{\text{induced}} = 3 \mu\text{A}$  as an upper bound. The current picked up on the superconducting nanowire is directly proportional to the RF frequency  $\omega$ , the voltage applied to the RF electrodes, and to the area enclosed by the SNSPD meander. This induced current is a significant fraction of the typical critical current of SNSPDs of  $\sim 10 \mu\text{A}$  [Dau09, YKD<sup>+</sup>07] and will significantly impact the operation of the SNSPD in terms of detection efficiency. However, the induced current from the RF source at detector vary in time (synchronized to RF source) and so will the critical current. This suggests synchronizing the bias current to RF source as a possible solution.

## Summary

In the previous section we have thus far only computed one possible detrimental effect of bringing an operational ion trap close to an SNSPD. Other effects can be estimated similarly. It would be important to quantify these effects in a potential design. Reducing and minimizing the effects of electric field pickup, along with the vastly different device scales demanding a multi-step fabrication process, are examples of many technical challenges in realizing this ion trap + SNSPD system. Given the typical current estimates of device efficiency and geometry, it appears that the overall system efficiency would be uncompetitive with conventional bulk optics + PMT detection schemes. However, this integrated design might become more relevant for potential systems with extremely limited optical access, such as integrating ions with very small optical cavities, or operation inside a dilution refrigerator for integration with superconducting qubits.

# Chapter 10

## Conclusion

This thesis has investigated a few examples of challenges facing the prospects of scaling up trapped ion quantum information processing. In this brief conclusion, I will summarize some key results in this work and provide some outlook for future directions.

Part I described the demonstration of a basic quantum logic gate in a surface-electrode ion trap. In previous work performed in this group, low decoherence rates were achieved via cryogenic cooling and implementation of the Cirac-Zoller CNOT gate was demonstrated by population measurements. Here, we extended this work by implementing full quantum process tomography on the two qubit system and thus determined the mean and process fidelities of the gate to be 91% and 85%, respectively. We also studied the effects of decoherence mechanisms and imperfect control on the gate fidelity via numerical simulations and concluded that the gate infidelity is reasonably accounted for by known sources of error. There appears to be no additional fundamental sources of errors that might prevent the many successful algorithms that were demonstrated with 3D ion traps from being repeated with surface-electrode ion traps. In addition, the techniques and tools that were developed to precisely control the atomic and motional states for a single ion might be applicable to studies of other quantum mechanical systems.

The Cirac-Zoller CNOT gate was one of the first gates proposed for realizing entanglement in trapped ions, but it is not without its drawbacks, most notably its requirement for ground state cooling. From the time of this work, the community has shifted its attention to other kinds of gates which are less sensitive to the motional state occupation. The Mølmer-Sørensen gate in particular has been performed in 3D ion traps with high fidelity [BKRB08]. Other groups have reported implementations of this gate in surface-electrode traps as a straightforward extension of the techniques used in 3D traps. Other gates such as the geometric phase gate have also been demonstrated in microfabricated traps with high fidelity [LDM<sup>+</sup>03, HHJ<sup>+</sup>09b].

Part II examined two key issues associated with the small scale of surface-electrode traps: electric field noise leading to heating and motional state decoherence of the ion, and laser-induced charging of the trap surface. Continuing the study of material dependence of heating rate previously performed in this group, we measured the heating rate in aluminum and superconducting traps. Aluminum traps showed a wide variation in measured heating rate between different samples, similar to previous materials (gold and silver) that were

tested in the same experimental setup. Superconducting traps performed comparably to traps made of these metals. In addition, the measured heating rates above and below the superconducting transition were the same within error bars, strongly indicating that heating is caused by surface, not bulk, effects.

This and previous heating rate studies suggest a different experimental approach for investigating this problem: instead of fabricating many different traps and testing them individually, subjecting them to variations in fabrication and processing, any method that varies trap properties in situ would be preferred. Recently the community has shifted to exactly this kind of experiments, for example by performing laser cleaning and comparing the before-and-after heating rates. Some very recent results [AHJ<sup>+</sup>11, HCW<sup>+</sup>11] indicated improved heating rates after laser cleaning, from a factor of 2 to orders of magnitude, suggesting that surface contamination plays a major factor. There is ongoing work in our group and other groups to study graphene traps, chosen for their passivity to surface adsorbates, in the hopes that they might offer more consistently low heating rates in both cryogenic and room-temperature environments.

Laser-induced charging has been an issue for ion trap experiments, and we studied it systematically by varying materials and laser wavelengths. Consistent with a physical model based on photoemission, we found that charging is worst for short wavelengths, and worse in aluminum, which both has a native oxide layer and a lower work function than other metals including copper and gold. While we were not able to quantitatively relate the measured charging rates with material properties such as modifications of the work function by the oxide layer, this work nevertheless provided some insight on the importance of material selection for building scalable ion traps.

Part III examined the question of scalability at a system level: what are the next steps we need to go beyond a handful of qubits? We addressed three specific technical needs: a modeling system to aid with design and implementation of practical algorithms, demonstration of such algorithms, and techniques for integrating with other quantum systems.

A modeling system that evaluates practical pulse sequences for implementing specific algorithms, and can predict experimental outcomes, would be a useful tool for evaluating design tradeoffs in practice. We wrote TIQC-SPICE, a program to simulate a complex trapped ion experiment, and incorporated various models of noise using experimentally determined parameters and evaluated their impact on real pulse sequences using Monte Carlo methods. The simulator is an update of a decade-old tool for simulating a 3D ion trap experiment at the University of Innsbruck. Other parts of the modeling system are actively being developed by our collaborators and other groups. As ion traps, especially surface-electrode ion traps, grow in complexity, such a tool is integral to their design and implementation, just as SPICE was for the classical computer. A short-term goal for TIQC-SPICE would be to take a role in the experimental realization of a significant quantum algorithm, perhaps even with error correction, in a surface-electrode ion trap. We also hope that a general-purpose trapped ion simulator may be useful for other kinds of experiments using trapped ions. The quantum computing community as a whole is beginning to address system architecture related issues, and a number of large collaborations have just recently begun. Many of the approaches being explored today are strongly inspired by conventions in clas-



sical computer architecture, which progressed alongside the physical hardware development to achieve present-day success. Perhaps we can imagine TIQC-SPICE as being one small beginning to such grand visions.

Central to the validation of a system model and design framework is the demonstration of real quantum algorithms using the quantum system. We realized the quantum Fourier transform, order-finding, and Shor's algorithm using a trapped ion quantum computer. Our result on Shor's algorithm for factoring 15 using 5 ions is one of the most complex quantum algorithms realized in trapped ion quantum computing to date. The success of these experiments required using capabilities currently unique to the ion trap system: ability to selectively hide qubits, fast initialization and measurement, and ability to perform feed-forward in a circuit. We also simulated these experiments and obtained results that showed reasonable agreement with the data, suggesting the validity of our noise model. In addition, the simulation was used to identify addressing error, dephasing, and hiding as the main technical sources of error that contributed to algorithm infidelity, providing useful guidelines for improving future experiments. While these experiments are small milestones for ion trap quantum computing, they are more important in improving our understanding of the physical system and recognizing limitations in scaling to larger algorithms. Keys to implementing larger algorithms in the near future include noise-insensitive gates, better pulse sequence optimization, and practical quantum error correction.

The ultimate quantum information processor may well require the integration of many different qubit systems. To this end, it is desirable to integrate ions with photons in order to improve collection efficiency for state detection and facilitate the interchange of atomic qubits to photon qubits. We have investigated two variations on this theme. One is to microfabricate a surface-electrode ion trap on top of a high-reflectivity mirror. We showed stable trapping of ions above the mirror with typical heating rates and low sensitivity to light-induced charging, and that the fabrication does not significantly degrade the quality of the mirror. The next step of this project, currently in progress in this group, is to integrate such a planar mirror trap with a small concave mirror to form a micro cavity. The aim is to achieve strong coupling with a single trapped ion, which is necessary for efficient transfer of quantum information from atomic to photonic qubits. The other goal is to achieve higher light collection efficiency by directly integrating photon detectors with surface-electrode traps. We have fabricated and tested a trap made of indium-tin oxide, a transparent conductor, and showed first signs of photon collection through the trap by a conventional Si PIN photodiode. This proof of principle suggests the viability of this approach, but the intrinsic collection efficiency of the photodetector used was a major limiting factor. Several alternatives are currently being investigated; we discussed one possible alternative, a superconducting nanowire single-photon detector, which is expected to have greater than 50% quantum efficiency in the desired wavelength. However, their sensitivity to stray fields is a major technical challenge in this case.

This thesis is a part of a long history of theses and related works all aiming to address the same questions: is it possible, and what would it take, to build a trapped ion quantum computer? While there appears to be no fundamental limit to building a quantum computer, undoubtedly the technical challenges to be overcome are immense. For ion traps,

ongoing efforts in the community exist to increase coherence times, develop scalable control schemes, and design system architectures that would be fault-tolerant and technically feasible. Similar efforts are happening in parallel in almost all of the other physical qubit implementations available today.

No one knows whether ion traps or any other physical qubits we know of today might be found in a future quantum information processor. We might be reminded that Charles Babbage's Difference Engine bears no physical resemblance to the classical computers we now have, computers which took nearly 200 years to develop their present form. As hard as it might be to imagine such radical developments in a technology as quantum computing, it is just as hard to envision what potential benefits and new capabilities that revolutionary technologies will bring. It is hoped that this thesis might contribute to perhaps unforeseen developments in quantum information processing and beyond.

# Appendix A

## Ion trap history

This appendix lists all traps tested within the duration of this thesis. All traps have the same design as described in Chapter 3.

Table A.1: Gold

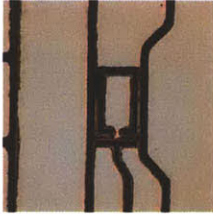
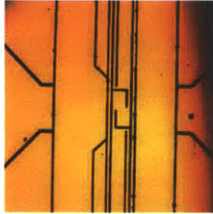

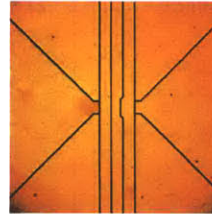
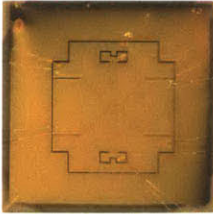
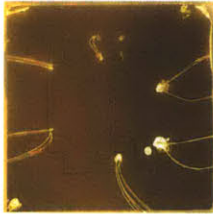
Trap	GPT v3 I	GPT v4 I	GPT v4 II	GPT v2 IV
ion height fabrication data	 100 $\mu\text{m}$ 2007/12/11 2008/01/22-24, 02/26, 03/03	 100 $\mu\text{m}$ 2008/04/04	 100 $\mu\text{m}$ 2008/04/04 2008/06/04- 05, 06/10-11, 06/16-17, 07/13, 07/16, 07/25-26, 09/05, 09/11-13, 09/18-19	 75 $\mu\text{m}$ 2008/08/27
Trap	GPT v5 I	GVT-75		
ion height fabrication data	 75 $\mu\text{m}$ 2010/01/08 2010/08/17-18	 75 $\mu\text{m}$ 2010/05		

Table A.2: Aluminum, Aluminum Oxide, Copper

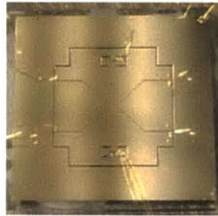





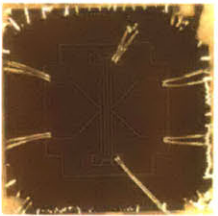
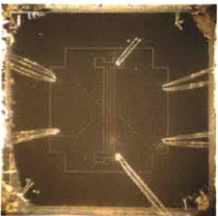

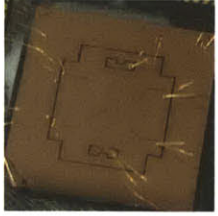
Trap	Al I 	Al IIa 	Al IIb 	Al IIIa 
ion height fabrication data	75 $\mu\text{m}$ 2009/07/10 2009/07/15-21, 08/29-30	75 $\mu\text{m}$ 2009/07/10 2009/09/09-10	75 $\mu\text{m}$ 2009/07/10	100 $\mu\text{m}$ 2009/10/19 2009/11/08-11
Trap	Al IIIb 	Al IIIc 	AlO-10 	AlO-5 
ion height fabrication data	100 $\mu\text{m}$ 2009/10/19 2009/12/18-21	100 $\mu\text{m}$ 2009/10/19 2010/01/05- 07, 01/21-22, 02/26-27	100 $\mu\text{m}$ 2010/03/04 2010/03/10-12	100 $\mu\text{m}$ 2010/03/11 2010/03/15-18, 2011/04/07
Trap	AlO-20 	Cu Ia 		
ion height fabrication data	100 $\mu\text{m}$ 2010/04 2010/04/23, 05/10	100 $\mu\text{m}$ 2010/02 2010/02/08-11		

Table A.3: Niobium and Niobium Nitride

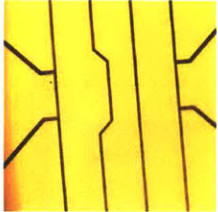

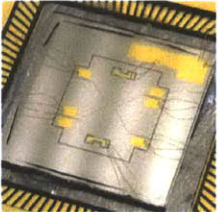
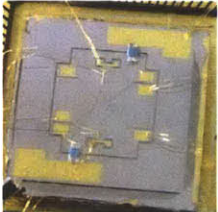
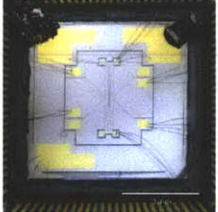

Trap	NbN I	NbN II	Nb I	Nb II
ion height fabrication data	 <p>100 <math>\mu\text{m}</math> 2007/11-12 2008/02/08</p>	 <p>100 <math>\mu\text{m}</math> 2008/03/03</p>	 <p>100 <math>\mu\text{m}</math> 2008/04/11 2008/05/07, 07/02-03, 07/07, 10/22-25, 12/14- 15, 2009/03/01, 04/04-06, 04/19- 20, 06/26-29, 09/19-25</p>	 <p>100 <math>\mu\text{m}</math> 2008/10/24 2008/11/05, 12/07</p>
Trap	Nbg Ia	Nbg Ib		
ion height fabrication data	 <p>100 <math>\mu\text{m}</math> 2010/04-06 2010/06/14-16, 10/18, 10/28- 30, 11/02, 2011/04/01, 10/27</p>	 <p>100 <math>\mu\text{m}</math> 2010/04-06 2010/06/05</p>		

Table A.4: MirrorTrap

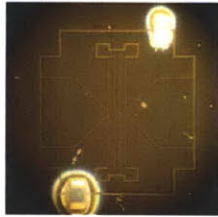
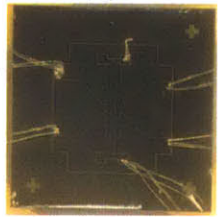
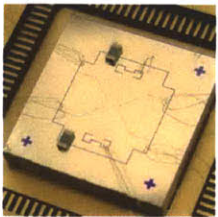
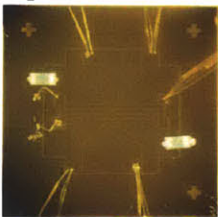
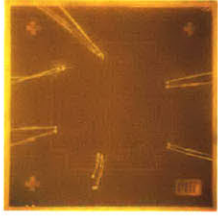
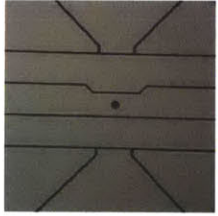
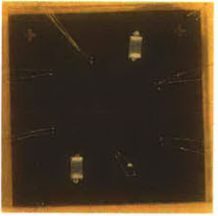
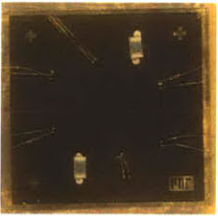
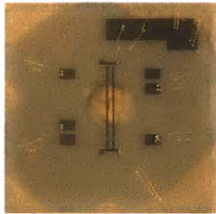

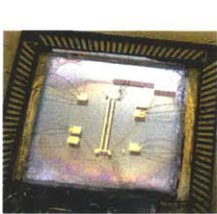

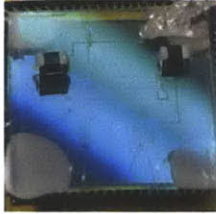
Trap	alpha-yge1	alpha-yge2	alpha-sxw1	alpha-sxw2
ion height fabrication data	 150 $\mu\text{m}$ 2010/05	 150 $\mu\text{m}$ 2010/05 2010/05/24-30, 06/01-03, 06/06, 06/15-16	 150 $\mu\text{m}$ 2010/04/30	 150 $\mu\text{m}$ 2010/04/30
Trap	beta-1	beta-2	beta-IIa	beta-IIb
ion height fabrication data	 150 $\mu\text{m}$ 2010/08/24	 150 $\mu\text{m}$ 2010/08/24	 150 $\mu\text{m}$ 2010/09/10	 150 $\mu\text{m}$ 2010/09/10 2010/09/23-24, 09/27-28, 10/01- 03, 10/06-08



Table A.5: Indium-tin oxide

Trap	ITO Ia	ITO Ib	ITO IIa	ITO IIIa
ion height fabrication data	 100 $\mu\text{m}$ 2010/04 2010/11/22-25	 100 $\mu\text{m}$ 2010/06	 100 $\mu\text{m}$ 2010/06	 100 $\mu\text{m}$ 2011/01 2011/04/25
Trap	ITO IVa			
ion height fabrication data	 100 $\mu\text{m}$ 2011/02-06			





# Appendix B

## Pulse sequencer

This appendix describes the architecture and instruction set of the Direct-Digital Synthesis (DDS) -based pulse sequencer in the MIT lab. A number of frequently used pulse sequences in ion trap experiments are included.

### B.1 Architecture

Memory

- 4000× 32-bit memory
- 32-bit W register
- 12-bit INDF register
- Command: 8-bit opcode + 24-bit data

Internal registers

- `f_phasecorr`
- `globalphase`
- `globaltime`

### B.2 Instruction set

Table B.1: Instruction Set – DDS controls

Opcode	Command	Operation
01	DDSFRQ #, addr	set the DDS# frequency to value stored at addr
02	DDSAMP #, addr	set DDS# amplitude to value stored at addr
03	DDSPHS #, addr	set DDS# phase to value stored at addr, + <code>f_phasecorr</code> × <code>globaltime</code> – <code>globalphase</code>
04	DDSCOR #, addr	set DDS# phase correction register to value stored at addr
05	DDSTZR	zero time reference counter
06	DIGOUT #	turn on/off digital outputs. use a 4 bit mask as input.

Table B.2: Instruction Set – Timing controls

Opcode	Command	Operation
00	NOP	do nothing
07	COUNT addr	fetch x from addr, count photons for x * clock period (max x = 2 <sup>24</sup> )
08	DELAY addr	fetch x from addr, delay execution for x * clock period (max x = 2 <sup>24</sup> )
09	TRIG addr	fetch x from addr, delay until trigger or value

Table B.3: Instruction Set – Logic and arithmetic

Opcode	Command	Operation
0A	LDWR addr	load value from addr into W
0B	LDWI addr	load value from memory cell pointed to by INDF into W
0C	STWR addr	store W into memory at addr
0D	STWI	store W into memory cell pointed to by INDF
0E	LDINDF	load value from addr into INDF
0F	SWAPWI	swap the values in W and INDF
10	ANDW addr	fetch x from addr, W = W & x
11	ADDW addr	fetch x from addr, W = W + x
12	INC addr	fetch x from addr, W = x + 1
13	DEC addr	fetch x from addr, W = x - 1
14	CLRW	set W = 0
15	CMP addr	fetch x from addr, if W > x, W = W else W = 0

Table B.4: Instruction Set – Flow control

Opcode	Command	Operation
16	JMP addr	jump to address addr
17	JMPZ addr	jump to address addr only if W == 0
18	JMPNZ addr	jump to address addr only if W != 0
19	JMPOFF addr	jump to address addr + W
1A	FJMP addr	push registers W, INDF, and PC onto stack, jump to address addr
1B	FRET	pop W, INDF, and PC from stack, PC = PC + 1
FF	END	finish execution

## B.3 Higher-level constructs

Table B.5: Higher-level instructions

Instruction	syntax	expanded low-level code
Inserts	<code>#insert code.pp</code>	recursively parses code. works like “ <code>#include</code> ” in C.
If-else	<code>(cond) ? YES : NO</code>	<pre> CMP addr JMPZ label_no yes statements JMP label_yes label_no: no statements label_yes: NOP </pre>
function def	<pre> function f(arg0, arg1) { body } </pre>	<pre> var arg_f_0 var arg_f_1 f: body FRET </pre>
function call	<code>f(arg1, arg2)</code>	<pre> LDWR arg1 STWR arg1 LDWR arg2 STWR arg2 FJMP f </pre>

## B.4 Standard pulse sequences

### SimpleGates.pp

Defines the qubit laser and repumping laser pulses

note: the first 6 lines in this file appears in all subsequent files.

```
#define REDDDS      0
#define BLUEDDS    1
#define IRDIGOUT   7
#define IRON       8
#define IROFF      0

var ZERO 0
var NS_IRDELAY 4000

global function redpulse(freq, amp, phase, length, phcorr) {
    DDSAMP REDDDS, ZERO

    DDSCOR REDDDS, phcorr
    DDSFRQ REDDDS, freq
    DDSPHS REDDDS, phase
    DDSAMP REDDDS, amp

    DELAY length

    DDSAMP REDDDS, ZERO
    DDSFRQ REDDDS, ZERO
    NOP
}

global function 1033repump() {
    DIGOUT IRDIGOUT, IRON
    DELAY NS_IRDELAY
    DIGOUT IRDIGOUT, IROFF
}
```

### Readout.pp

Performs a fluorescence measurement

```
global function SimpleReadout(pointer) {

    DDSFRQ BLUEDDS, F_BlueOn
    DDSAMP BLUEDDS, A_BlueOn

    COUNT us_MeasTime
    LDINDF pointer
    STWI

    DDSAMP BLUEDDS, ZERO
    DDSFRQ BLUEDDS, ZERO
}
```

### InitStateS0.pp

Initialize the qubit, optionally with sideband cooling

```
# SC parameters
#define F_BSB (F_RedCenter - F_StarkSec)
#define F_MeasRabi 1/(2.*us_PiTime)
#define F_SCRabi 1/(2.*us_PiTime)
#define F_Sec math.sqrt(4.*F_StarkSec**2 + F_MeasRabi**2)
```

```

#define F_SCSec    math.sqrt(F_Sec**2 - F_SCRabi**2)
#define eta       2*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)

#define SCTO      us_PiTime/eta
#define f_Sctime  lambda N, T0: map(lambda x: max(T0/math.sqrt(x), T0/8.), range(int(SCloops), 1, -1))
# #define f_Sctime lambda N, T0: map(lambda x: T0-0.25*x, range(int(SCloops),1,-1))

# registers that don't get manipulated:
var ZERO          0
var US_CHECKTIME  us_MeasTime/3.
var CHECKTHOLD    2
var US_POLARTIME  us_PiTime*math.sqrt(5./2.)
var POLARDIV      3
var F_RedPL       (F_RedCenter + F_PolarOffset)
var F_RedSC       (F_RedCenter + F_SCSec/2.)

var us_Sctimelist f_Sctime(SCloops, SCTO) + [0.]
var counter_init  SCloops
var Sctimelist_p  &us_Sctimelist

# registers that do get manipulated:
var us_sctime 0
var counter 0
var pointer 0

global function Initialize_S0() {
init:  1033repump()
      DDSFRQ  BLUEDDS, F_Blue0n
      DDSAMP  BLUEDDS, A_Blue0n
      COUNT   US_CHECKTIME
      CMP     CHECKTHOLD
      JMPZ    init

      DDSAMP  BLUEDDS, ZERO
      DDSFRQ  BLUEDDS, ZERO

      LDWR    counter_init
      STWR    counter
      JMPZ    done
      LDWR    Sctimelist_p
      STWR    pointer

scool: LDINDF  pointer
      INC     pointer
      STWR    pointer
      LDWI
      STWR    us_sctime

      redpulse(F_RedSC, A_Red0n, ZERO, us_sctime, ZERO)
      1033repump()

      LDWR    counter
      ANDW    POLARDIV
      JMPNZ   nopol

      redpulse(F_RedPL, A_Red0n, ZERO, US_POLARTIME, ZERO)
      1033repump()

nopol: DEC     counter
      STWR    counter
done:  JMPNZ   scool
      1033repump()

```

```

    redpulse(F_RedPL, A_RedOn, ZERO, US_POLARTIME, ZERO)
    1033repump()
    redpulse(F_RedPL, A_RedOn, ZERO, US_POLARTIME, ZERO)
    1033repump()
}

```

## Shelving.pp

Initialize the qubit, apply a variable length qubit pulse and measure.

```

#insert Init_StateS0.pp
#insert Readout.pp
#insert SimpleGates.pp

var ZERO      0
var memstart  3900
var memend    4000
var mempointer 0

function main() {
    DIGOUT  IRDIGOUT, IRON

    LDWR    memstart
    STWR    mempointer
loop:      Initialize_S0()

    DELAY   ms_ReadoutDly

    redpulse(F_RedOn, A_RedOn, ZERO, us_RedTime, ZERO)
#    1033repump()
    SimpleReadout(mempointer)

    INC     mempointer
    STWR    mempointer
    (CMP memend) ? NOP : JMP loop

    DIGOUT  IRDIGOUT, IRON
    DDSFRQ  BLUEDDS, F_BlueHi
    DDSAMP  BLUEDDS, A_BlueHi
}

```

## Ramsey.pp

Perform a Ramsey measurement with variable delay; optionally include spin-echo pulses

```

#define F_Rabi    1/(2.*us_PiTime)
#define F_Sec     math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta       2.*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)

var memstart  3900
var memend    4000
var mempointer 0

var us_HalfPi  us_PiTime/2.
var us_Tfixed  500

#insert Init_StateS0.pp
#insert Readout.pp
#insert SimpleGates.pp

function echo() {
    DELAY  us_Tfixed
    redpulse(F_RedOn, A_RedOn, ZERO, us_PiTime, ZERO)
}

```



```

}
function echo2() {
    DELAY us_Tfixed
    DELAY us_Tfixed
    redpulse(F_RedOn, A_RedOn, ZERO, us_PiTime, ZERO)
}

function main() {
    LDWR memstart
    STWR mempointer
loop: Initialize_S0()

    DDSTZR

    redpulse(F_RedOn, A_RedOn, ZERO, us_HalfPi, ZERO)
#    echo()
#    echo2()
    DELAY us_RamseyDly
    redpulse(F_RedOn, A_RedOn, Ph_Ramsey, us_HalfPi, ZERO)

    SimpleReadout(mempointer)

    INC mempointer
    STWR mempointer
    (CMP memend) ? NOP : JMP loop

    DIGOUT IRDIGOUT, IRON
    DDSFRQ BLUEDDS, F_BlueHi
    DDSAMP BLUEDDS, A_BlueHi
}

```

## StatePrep.pp

### State preparation pulses for tomography

```

#define F_uconv 1e6/1e9*(2**32)
#define ph_uconv (2**14)/360.
#define us_uconv 1e-6/16e-9

#define F_Rabi 1/(2*us_PiTime)
#define F_Sec math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta 2.*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)

#define us_Pi us_PiTime * us_uconv
#define us_BSBPi (us_PiTime/eta) * us_uconv
#define us_HalfPi us_PiTime/2. * us_uconv
#define us_BSBHalfPi us_PiTime/(2*eta) * us_uconv

#define ph_HalfPi (45.) * ph_uconv
#define ph_Pi (90.) * ph_uconv
#define ph_negHalfPi (-45.) * ph_uconv

#define OtherStark OtherStarkCorr*(F_Rabi**2)
#define RedStark (2*F_StarkSec - math.sqrt(4*F_StarkSec**2 + F_Rabi**2))

#define F_starkCarr 0*OtherStark/2. * F_uconv
#define F_starkBSB (RedStark + OtherStark)/2. * F_uconv

#define F_Carrier (F_RedCenter - F_starkCarr) * F_uconv
#define F_BSB (F_RedCenter - F_StarkSec) * F_uconv

var data \
[F_Carrier, ph_negHalfPi, us_Pi , F_starkCarr, ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
F_BSB , ZERO , us_BSBPi , F_starkBSB, ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
F_BSB , ZERO , us_BSBPi , F_starkBSB , F_Carrier, ph_Pi , us_Pi , F_starkCarr, ZERO, ZERO, ZERO, ZERO, \
F_Carrier, ph_negHalfPi, us_HalfPi , F_starkCarr, F_BSB , ZERO , us_BSBPi, F_starkBSB, ZERO, ZERO, ZERO, ZERO, \
F_Carrier, ph_negHalfPi, us_HalfPi , F_starkCarr, F_BSB , ph_negHalfPi, us_BSBPi , F_starkBSB , ZERO, ZERO, ZERO, ZERO, \
F_Carrier, ph_negHalfPi, us_HalfPi , F_starkCarr, ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
F_Carrier, ZERO , us_HalfPi , F_starkCarr, ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \

```

```

F_BSB , ph_Pi , us_BSBHalfPi, F_starkBSB , F_Carrier, ph_negHalfPi, us_Pi , F_starkCarr, ZERO, ZERO, ZERO, ZERO, \
F_BSB , ph_HalfPi , us_BSBHalfPi, F_starkBSB , F_Carrier, ph_negHalfPi, us_Pi , F_starkCarr, ZERO, ZERO, ZERO, ZERO, \
F_BSB , ZERO , us_BSBHalfPi, F_starkBSB , ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
F_BSB , ph_HalfPi , us_BSBHalfPi, F_starkBSB , ZERO , ZERO , ZERO , ZERO , ZERO, ZERO, ZERO, ZERO, \
F_BSB , ZERO , us_BSBPi , F_starkBSB , F_Carrier, ph_HalfPi , us_HalfPi, F_starkCarr, ZERO, ZERO, ZERO, ZERO, \
F_BSB , ZERO , us_BSBPi , F_starkBSB , F_Carrier, ph_Pi , us_HalfPi, F_starkCarr, ZERO, ZERO, ZERO, ZERO, \
F_Carrier, ph_HalfPi , us_HalfPi , F_starkCarr, F_BSB , ph_Pi , us_BSBPi , F_starkBSB , F_Carrier, ph_negHalfPi, us_Pi, F_starkCarr, \
F_Carrier, ph_HalfPi , us_HalfPi , F_starkCarr, F_BSB , ph_HalfPi , us_BSBPi , F_starkBSB , F_Carrier, ph_negHalfPi, us_Pi, F_starkCarr ]

var data_p &data
var addr8 map(lambda x: 12*x, range(16))
var addr8_p &addr8
var datarow 0
var params_p 0

global function StatePrep(initState) {

    LDWR    addr8_p[initState]
    STWR    datarow

    LDWR    data_p
    ADDW    datarow
    STWR    params_p

# Prep UVV

    LDWR    params_p[n0]
    JMPZ    skip1
    redpulse(params_p[n0], A_RedOn, params_p[n1], params_p[n2], params_p[n3])

skip1: LDWR    params_p
    ADDW    n4
    STWR    params_p
    LDWR    params_p[n0]
    JMPZ    skip2
    redpulse(params_p[n0], A_RedOn, params_p[n1], params_p[n2], params_p[n3])

skip2: LDWR    params_p
    ADDW    n4
    STWR    params_p
    LDWR    params_p[n0]
    JMPZ    r_done
    redpulse(params_p[n0], A_RedOn, params_p[n1], params_p[n2], params_p[n3])

r_done: NOP
}

```

## StateMeasure.pp

### State measurement pulses for tomography

```

#define F_uconv      1e6/1e9*(2**32)
#define ph_uconv    (2**14)/360.
#define us_uconv    1e-6/16e-9

#define F_Rabi      1/(2*us_PiTime)
#define F_Sec      math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta        2.*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)

var IONONTHOLD    3

#define us_Pi      us_PiTime * us_uconv
#define us_BSBPi  (us_PiTime/eta) * us_uconv
#define us_HalfPi us_PiTime/2. * us_uconv
#define us_BSBHalfPi us_PiTime/(2*eta) * us_uconv

#define ph_HalfPi (45.) * ph_uconv

#define OtherStark OtherStarkCorr*(F_Rabi**2)
#define RedStark   (2*F_StarkSec - math.sqrt(4*F_StarkSec**2 + F_Rabi**2))

#define F_starkCarr 0*OtherStark/2. * F_uconv
#define F_starkBSB  (RedStark + OtherStark)/2. * F_uconv

#define F_Carrier  (F_RedCenter - F_starkCarr) * F_uconv
#define F_BSB     (F_RedCenter - F_StarkSec) * F_uconv

var data \
[ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO, \

```

```

ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , \
ZERO , ZERO , ZERO , ZERO , F_BSB, ph_HalfPi, us_BSBPi , F_starkBSB, ZERO , ZERO , ZERO , ZERO , \
F_Carrier, ph_HalfPi, us_Pi , F_starkCarr, F_BSB, ph_HalfPi, us_BSBPi , F_starkBSB, ZERO , ZERO , ZERO , ZERO , \
F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , \
F_Carrier, ZERO , us_HalfPi, F_starkCarr, ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , ZERO , \
ZERO , ZERO , ZERO , ZERO , F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, \
F_Carrier, ph_HalfPi, us_Pi , F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, \
F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, \
F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ZERO , us_HalfPi, F_starkCarr, \
ZERO , ZERO , ZERO , ZERO , F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ZERO , us_HalfPi, F_starkCarr, \
F_Carrier, ZERO , us_Pi , F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ZERO , us_HalfPi, F_starkCarr, \
F_Carrier, ZERO , us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ZERO , us_HalfPi, F_starkCarr, \
F_Carrier, ZERO , us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ZERO , us_HalfPi, F_starkCarr, \
F_Carrier, ZERO , us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, \
F_Carrier, ph_HalfPi, us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, ZERO , ZERO , ZERO , ZERO , \
F_Carrier, ZERO , us_HalfPi, F_starkCarr, F_BSB, ph_HalfPi, us_BSBHalfPi, F_starkBSB, ZERO , ZERO , ZERO , ZERO , \

```

```

var data_p      &data
var addr      map(lambda x: 12*x, range(16))
var addr_p     &addr
var datarow    0
var params_p   0

```

```

global function StateMeasure(measState, mempointer) {

```

```

    LDWR    addr_p[measState]
    STWR    datarow

    LDWR    data_p
    ADDW    datarow
    STWR    params_p

# MeasUVV

    LDWR    params_p[n0]
    JMPZ    meas1
    redpulse(params_p[n0], A_Red0n, params_p[n1], params_p[n2], params_p[n3])
meas1: SimpleReadout(mempointer)

    LDWR    params_p
    ADDW    n4
    STWR    params_p
    LDWR    params_p[n0]
    JMPZ    r_done

    LDINDF mempointer
    LDWI
    (CMP IONONTHOLD) ? NOP : JMP dark
bright: CLRW
    STWI
    JMP    r_done
dark: redpulse(params_p[n0], A_Red0n, params_p[n1], params_p[n2], params_p[n3])

    LDWR    params_p
    ADDW    n4
    STWR    params_p
    LDWR    params_p[n0]
    JMPZ    meas2
    redpulse(params_p[n0], A_Red0n, params_p[n1], params_p[n2], params_p[n3])

meas2: SimpleReadout(mempointer)

r_done: NOP
}

```

## CNOT.pp

CNOT gate between atomic and motional qubits of one ion

```

### frequency defs
#define F_Rabi      1/(2*us_PiTime)
#define F_Sec      math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta        2.*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)
#define us_BSBPiTime (us_PiTime/eta)
#define us_bPiBySqrt2 (us_BSBPiTime/math.sqrt(2))

### phase defs REMEMBER FACTORS OF 2
#define OtherStark (OtherStarkCorr*(F_Rabi**2))
#define RedStark (2*F_StarkSec - math.sqrt(4*F_StarkSec**2 + F_Rabi**2))

```

```

#define F_delstarkCarr 0*OtherStark/2.
#define F_delstarkBSB (RedStark + OtherStark)/2.

var F_starkCarr      F_delstarkCarr
global var F_starkBSB F_delstarkBSB

var F_Carrier  (F_RedCenter - F_delstarkCarr)
var F_BSB      (F_RedCenter - F_StarkSec)

### time defs. a lot of unnecessary defs here, for readability.
var us_Tg0     us_PiTime/2.
var us_Tg1     us_BSBPiTime
var us_Tg2     us_bPiBySqrt2
var us_Tg3     us_BSBPiTime
var us_Tg4     us_bPiBySqrt2
var us_Tg5     us_PiTime/2.

var Ph_0       (-45.)
var Ph_1       (0.)
var Ph_2       (45.)
var Ph_3       (0.)
var Ph_4       (45.)
var Ph_5       (-45.)

global function cnot() {
# Uc: R(pi/2,0)
    redpulse(F_RedCenter, A_RedOn, Ph_0, us_Tg0, F_starkCarr)
# Um(0): R+(pi,0)
    redpulse(F_BSB, A_RedOn, Ph_1, us_Tg1, F_starkBSB)
# Um(pi/2): R+(pi/sqrt2,pi/2)
    redpulse(F_BSB, A_RedOn, Ph_2, us_Tg2, F_starkBSB)
# Um(0): R+(pi,0)
    redpulse(F_BSB, A_RedOn, Ph_3, us_Tg3, F_starkBSB)
# Um(pi/2): R+(pi/sqrt2,pi/2)
    redpulse(F_BSB, A_RedOn, Ph_4, us_Tg4, F_starkBSB)
# Uc: R(pi/2,pi)
    redpulse(F_RedCenter, A_RedOn, Ph_5, us_Tg5, F_starkCarr)
}

```

## Tomography.pp

State tomography on one ion (atomic + motional qubit)

```

#define F_Rabi 1/(2*us_PiTime)
#define F_Sec  math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta    2*3.1415*math.sqrt(1.055e-7/(4*math.pi*F_Sec*88*1.676))/(674.0e-6)

#insert Init_StateS0.pp
#insert Readout.pp
#insert CNOT.pp
#insert SimpleGates.pp

var memstart 3900
var memend 4000
var mempointer 0

var F_BSB      (F_RedCenter - F_StarkSec)
var us_BSBPiTime (us_PiTime/eta)

var iDoBSBPi 1
var iDoCarrierPi 2

```

```

var rDoCarrierPi 4
var rDoZero      8
var IONONTHOLD  3

var Ph_Pi 90.

function main() {
    LDWR    memstart
    STWR    mempointer
loop:      Initialize_S0()

    DDSTZR

    LDWR    StateMask
    (ANDW iDoBSBPi) ? redpulse(F_BSB, A_RedOn, ZERO, us_BSBPiTime, F_starkBSB) : NOP

    LDWR    StateMask
    (ANDW iDoCarrierPi) ? redpulse(F_RedCenter, A_RedOn, ZERO, us_PiTime, ZERO) : NOP

# Stick Gate Here
    cnot()
# End of gate

    LDWR    StateMask
    (ANDW rDoCarrierPi) ? redpulse(F_RedCenter, A_RedOn, Ph_Pi, us_PiTime, ZERO) : NOP

# First measurement
    SimpleReadout(mempointer)
    LDINDF  mempointer
    LDWI
    CMP     IONONTHOLD
    JMPZ    r_inDstate

    LDWR    StateMask
    ANDW    rDoZero
    JMPZ    r_done
    CLRW
    STWI
    JMP     r_done

r_inDstate: NOP
    redpulse(F_BSB, A_RedOn, Ph_Pi, us_BSBPiTime, F_starkBSB)
    SimpleReadout(mempointer)

r_done: INC    mempointer
    STWR    mempointer
    (CMP memend) ? NOP : JMP loop

    DIGOUT  IRDIGOUT, IRON
    DDSFRQ  BLUEDDS, F_BlueHi
    DDSAMP  BLUEDDS, A_BlueHi
}

```

## ProcessTomo.pp

Process tomography on one ion (atomic + motional qubit)

```

#####
# partial Process tomography: initialize in one of 4 states, measure 1/16 of the full density matrix
# in DDScon: define StateMask
# calls file/function StatePrep(), StateMeasure()
#####

```

```

#define F_Rabi      1/(2*us_PiTime)
#define F_Sec      math.sqrt(4*F_StarkSec**2 + F_Rabi**2)
#define eta        2.*math.pi*math.sqrt(1.055e-7/(4.*math.pi*F_Sec*88.*1.676))/(674.0e-6)

#define OtherStark OtherStarkCorr*(F_Rabi**2)
#define RedStark   (2*F_StarkSec - math.sqrt(4*F_StarkSec**2 + F_Rabi**2))

#define F_starkCarr 0*OtherStark/2.
var F_starkBSB     (RedStark + OtherStark)/2.

#define F_Carrier   (F_RedCenter - F_starkCarr)
var F_BSB          (F_RedCenter - F_StarkSec)
var us_BSBPi       (us_PiTime/eta)
var us_BSBHalfPi   us_PiTime/(2*eta)
var us_HalfPi       us_PiTime/2.
var ph_HalfPi       (45)%180.

var us_wait        300

var ZERO           0
var memstart       3900
var memend         4000
var mempointer     0

var initState      int(StateMask) / 16
var measState      int(StateMask) % 16

#insert Init_StateS0.pp
#insert Readout.pp
#insert CNOT.pp
#insert SimpleGates.pp
#insert StatePrep.pp
#insert StateMeasure.pp

function main() {
    LDWR    memstart
    STWR    mempointer
loop:      Initialize_S0()

           DDSTZR

           StatePrep(initState)

# Stick Gate Here
           cnot()
# End of gate

           StateMeasure(measState, mempointer)

r_done:   INC    mempointer
           STWR    mempointer
           (CMP memend) ? NOP : JMP loop

           DIGOUT  IRDIGOUT, IRON
           DDSFRQ  BLUEDDS, F_BlueHi
           DDSAMP  BLUEDDS, A_BlueHi
}

```

# Appendix C

## TIQC-SPICE code

### C.1 TIQC-SPICE module list

Table C.1 lists the core modules, classes, and functions in the TIQC simulator.

### C.2 Example simulations

Rabi with detuning

```
import numpy as np
import matplotlib.pyplot as pl
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc

pi = np.pi

NumberOfIons = 1
NumberOfPhonons = 10

hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)
params = sim.parameters(hspace)
dec = sim.decoherence(params)

# list of detunings to try
detunings = np.array([0, 0.02, 0.05, 0.1]) * params.omz
detuning_labels = ['$\omega_z$', '0.02$\omega_z$', '0.05$\omega_z$', '0.1$\omega_z$']
Ylist = []

# pulse sequence: a single carrier pulse
pulseseq = sim.PulseSequence( [ \
    sim.Rcar(params, 8*pi, 0,0), \
    ] )

# simulate each detuning
for det in detunings:
    pulseseq[0].detuning = det
    data = qc.simulateevolution(pulseseq, params, dec)
    data.tracedpopulation(0)

T = data.T
Ylist.append(data.YtrN[:,0])
```



```

# make plot
for i,Y in enumerate(Ylist):
    pl.plot(T, Y, label=detuning_labels[i])
pl.xlabel('time [ $\mu$ s]')
pl.ylabel('D state population')
pl.legend()
pl.show()

```

## Rabi on thermal state

```

import numpy as np
import matplotlib.pyplot as pl
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc

pi = np.pi

NumberOfIons = 1
NumberOfPhonons = 10

hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)
hspace.initial_state('thermal', nBar=5)
params = sim.parameters(hspace)
params.eta = np.array([0.3,0.3])
dec = sim.decoherence(params)

params.LDapproximation = False

pulseseq = sim.PulseSequence( [ \
    sim.Rcar(params, 10*pi, 0, 0), \
    ] )
pulseseq[0].dotimedepPulse = True

data = qc.simulateevolution(pulseseq, params, dec)

data.tracedpopulation(2)

```

## Spectrum

```

import numpy as np
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc
import PyTIQC.tools.progressbar as progbar
import matplotlib.pyplot as pl

pi = np.pi

NumberOfIons = 1
NumberOfPhonons = 1
hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)
hspace.initial_state("thermal", nBar=0.5)
params = sim.parameters(hspace, lab='MIT')
dec = sim.decoherence(params)

params.stepsize = 600

# scan over this frequency range and store the results in YR
detuning = np.arange(-1.2*params.omz, 1.2*params.omz, params.omz/300)
YR = np.zeros([len(detuning),hspace.dimensions], np.complex128)

# create progressbar
widgets = [progbar.Percentage(), ' ', progbar.Bar(), ' ', progbar.ETA()]

```

```

pbar = progbar.ProgressBar(widgets=widgets).start()

for i in range(len(detuning)):
    params.detuning = detuning[i]
    # do a carrier transition with duration close to a pi-pulse on the sideband
    pulseseq = sim.PulseSequence( [ \
        sim.Rcar(params, pi/params.eta[0], 0), \
        ] )

    data = qc.simulateevolution(pulseseq, params, dec)
    YR[i,:] = data.YR[-1,:] # get last data point

    pbar.update(int(1.*(i+1)*100/(len(detuning))))

data1 = sim.database(detuning, YR, hspace, pulseseq=None, statesvalid = False)
data1.tracedpopulation(0)
pl.plot(detuning/2/pi, 1-data1.YtrI)
pl.xlabel('relative frequency (2$\pi$ MHz)')
pl.ylabel('D state population')
pl.show()

```

## MS gate (2 ions)

```

import numpy as np
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc
import matplotlib.pyplot as pl

pi = np.pi

NumberOfIons = 2
PhononOverhead = 7
hspace = sim.hspace(NumberOfIons,2,NumberOfIons+PhononOverhead,0)
params = sim.parameters(hspace)
dec = sim.decoherence(params)

# use pi/4 as the smallest MS unit
#params.shortestMS=4.
#params.calcPulseParams()

pulseseq = sim.PulseSequence( [ \
    sim.RMS(params, pi, 0),
    ] )

data = qc.simulateevolution(pulseseq, params, dec)

data.displayPMTpopulations(1, plotlegend=True)

```

## Ramsey oscillations

```

import numpy as np
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc
import PyTIQC.tools.progressbar as progbar
from scipy.optimize import leastsq
import matplotlib.pyplot as pl

pi = np.pi

NumberOfIons = 1
NumberOfPhonons = 2
hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)

```

```

params = sim.parameters(hspace)
dec = sim.decoherence(params)

dec.doRandNtimes = 40
dec.dict['dephase'] = True
dec.progbar = False

params.detuning = 2*pi*0.01
params.progbar = False
params.coherenceTime = 1000

# parameters for pp
dec.doPP = True
params.use_servers( ['local'])
params.pplog = False
dec.doPPprintstats = False

def doRamseyRun(hspace, params, dec):
    tdelay = np.linspace(0,1000,200)
    YR = np.zeros([len(tdelay),hspace.dimensions], np.complex64)
    pop = np.zeros([len(tdelay),hspace.dimensions], np.complex64)

    widgets = [progbar.Percentage(), ' ', progbar.Bar(), ' ', progbar.ETA()]
    pbar = progbar.ProgressBar(widgets=widgets).start()

    for i in range(len(tdelay)):
        pulseseq = sim.PulseSequence( [ \
            sim.Rcar(params, pi/2, 0), \
            sim.Delay(params, tdelay[i]), \
            sim.Rcar(params, pi/2, pi/2) \
        ] )

        data = qc.simulateevolution(pulseseq, params, dec)

        YR[i,:] = data.YR[-1,:]

        pbar.update(int(1.*(i+1)*100/(len(tdelay))))

    data1 = sim.database(tdelay, YR, hspace, pulseseq=None, statesvalid = False)
    data1.tracedpopulation(0)

    return data1

data1 = doRamseyRun(hspace, params, dec)
pl.plot(data1.T, data1.YtrI)
pl.xlabel('time [μs]')
pl.ylabel('D state population')
pl.show()

```

## Spontaneous decay

```

import numpy as np
import PyTIQC.core.simtools as sim
import PyTIQC.core.qctools as qc

pi = np.pi

NumberOfIons = 1
NumberOfPhonons = 1
hspace = sim.hspace(NumberOfIons,2,NumberOfPhonons,0)
hspace.initial_state(state='quantum', qstate='D')
params = sim.parameters(hspace)

```

```
dec = sim.decoherence(params)

params.lifetime = 300
params.progbar = False

dec.doRandNtimes = 100
dec.dict['spontdecay'] = True
dec.doPP = True

pulseseq = sim.PulseSequence( [ \
    sim.Delay(params, 1000), \
    ] )

data = qc.simulateevolution(pulseseq, params, dec)

data.tracedpopulation(1)
```

<i>simtools</i>	
<b>hspace</b>	
initial_state	Define initial state for a pulse sequence
<b>parameters</b>	
initial_state	wrapper for the equivalent function in <b>hspace</b>
use_servers	select servers to be used for parallel processing
calcPulseParams	calculate parameters for the MS gate, if modified
use_rho0	use a density matrix as the starting state, by initializing multiple simulations, each using an eigenstate
<b>decoherence</b>	
calcDephasing	construct a time-correlated gaussian series
plotcorrelation	plot the sequence and correlation function (for testing)
<b>pulse</b>	
<b>PulseSequence</b>	
prepend	prepend the sequence with another sequence
append	append the sequence with another sequence
extend	like append, but input is a pulse sequence object
changeions	modify the list of ions on which the Z gates in the sequence are applied
getdigitalrepresentation	return a time vector with 1s where a pulse is on and 0s where there is no pulse
counttypes	count the number of types of each pulse
totalUnitary	return the ideal unitary of the whole sequence
<b>database</b>	
statesatpulse	save or plot state populations at the end of each pulse
endpopulation	print the end states, traced over the motional state
tracedpopulation	trace over motional state and plot D state population for each ion
displaypopulation	display final populations in text form and make plots of population for each ion+phonon state
displayPMTpopulations	make plot of populations observed with a PMT ( $n$ lines, each corresponding to $n$ bright ions)
displayPhaseSpace	plot the time evolution in phase space
<hr/>	
<i>qctools</i>	
<b>simulateevolution</b>	top-level, user-accessible function for running a simulation. a wrapper for simulateevolutionOnce, to run simulation for varying starting states
<b>simulateevolutionOnce</b>	a wrapper for simulationCore, for processing and collecting results from multiple MC instances
<b>simulationCore</b>	heart of the computation process. input: pulse sequence, parameters (includes hilbert space def), and decoherence object. output: a database object containing the time evolution of states.
<b>loadRun</b>	load the data from a shelve object. Returns: pulseq, params, dec, data
<b>saveRun</b>	save all info from a simulation run into a shelve object
<hr/>	
<i>qmttools</i>	
<b>Hamilton</b>	
Hamiltonian	defines the Hamiltonian for a given pulse in the laser frame
Hamiltonian_timedep_complete	define the time-dependent Hamiltonian in the full interaction frame of the ions
<b>Noise</b>	
Noise	generate the Noise_dict object to be used by qctools
<b>indexToState</b>	given an index, return the number of phonons and binary representation of state
<b>stateToIndex</b>	return the index given string representation of state; inverse of indexToState()
<hr/>	
<i>gates</i>	
<b>calculateevolution</b>	parallel of qc.simulateevolution
<b>fidelity</b>	fidelity of two density matrices
<b>jozsafid</b>	Jozsa (Uhlmann) fidelity of two density matrices
<b>tracedist</b>	trace distance of two quantum states. Inputs can be density matrices or populations.
<b>sso</b>	squared statistical overlap of two population vectors
<b>plotBar3D</b>	make 3D bar plots of density matrices or unitaries

Table C.1: Core modules, classes, and functions in TIQC-SPICE. Modules are in italics, classes and module functions are in typewriter font, class functions are indented.

# Appendix D

## Pulse sequences for algorithms

This appendix lists the pulse sequences used for the Quantum Fourier Transform, Order-finding, and Shor’s algorithms discussed in Chapter 8.

### Pulse syntax

A global operation is denoted by a box covering all qubits.

An addressed operation is a box covering one qubit.

Pulse types are listed in Table D.1.

$k$ ,  $\gamma$ ,  $\alpha$ , and  $\beta$  are corrections to Stark shifts caused by hiding/unhiding and are manually optimized during the experiment. The values used are  $\gamma = 0.9$ ,  $k = 1.16$ ,  $\beta = 1.09$ .

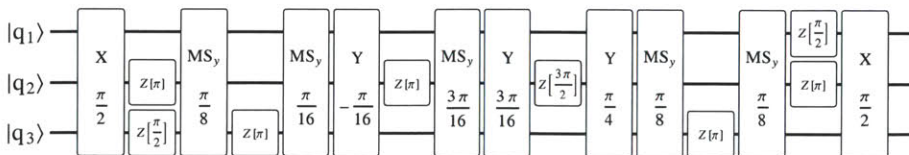
label	pulse	lab	TIQC-SPICE
X, Y	global carrier rotation	Rcar3	Rcar
$MS_x, MS_y$	global Mølmer-Sørensen gate	Rbichro1	RMS
Z	addressed phase shift	Rzred	Rac
$Z_i$	phase shift controlled by qubit $i$	MBRred	-
M,I	measure and recool	MBPMDetection	MeasInit
$H_s$	hide/unhide: carrier on S-D'	Rcar2	Hide
$H_d$	hide/unhide: carrier on S'-D	Rcar4	Hide

Table D.1: List of symbols used to denote pulse types in pulse sequences.

Special sub-sequences are marked by colors as shown in Table D.2.

### Quantum Fourier Transform

QFT sequence



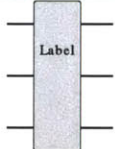
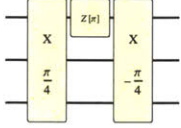
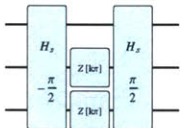
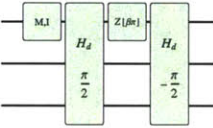
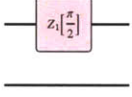
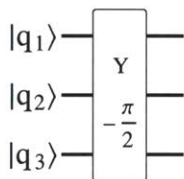
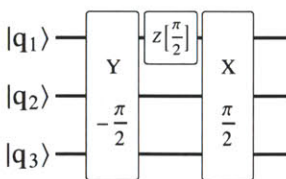
pulse or sub-sequence	Explanation
	Substituted by a pulse sequence with the given label.
	Hadamard gate on qubit 1; implemented by global and refocusing pulses.
	Typical hiding or unhiding sequence, implemented by global rotations on other carrier transitions and refocusing pulses.
	Measurement and reset. After recooling, 3 pulses reset the first qubit to $ S\rangle$ .
	Controlled rotation on qubit 1. Subscript indicates the (virtual) control qubit.

Table D.2: Color legend for pulse sequences.

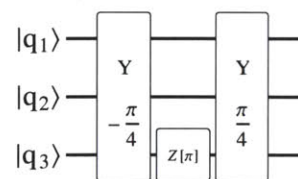
Sequences for preparing input states  $|\psi_n\rangle$  with periods  $n$ .



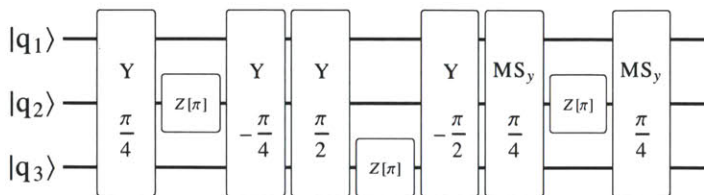
(a)  $|\psi_1\rangle$



(b)  $|\psi_2\rangle$



(c)  $|\psi_4\rangle$

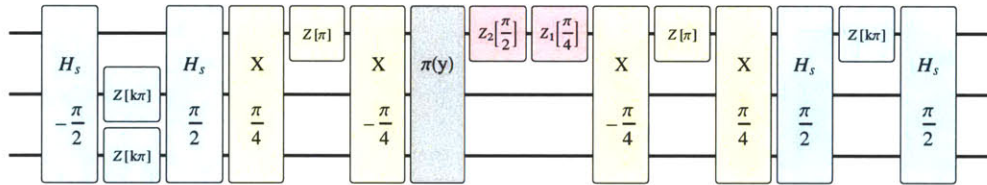
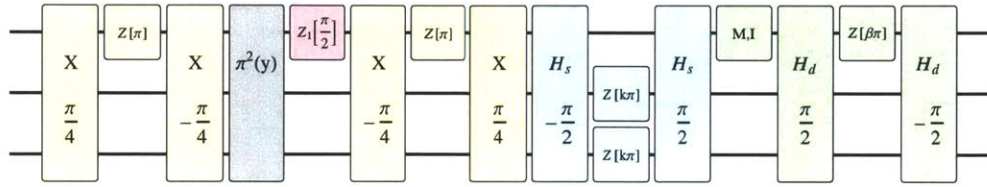
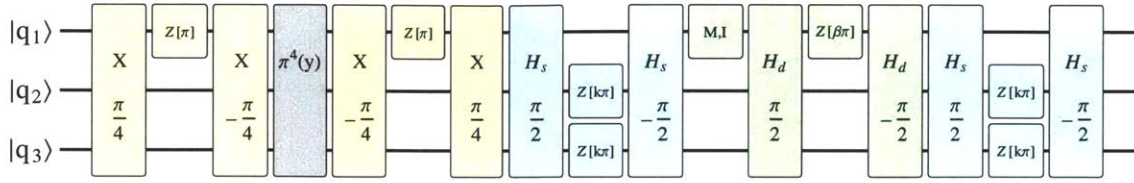


(d)  $|\psi_3\rangle$

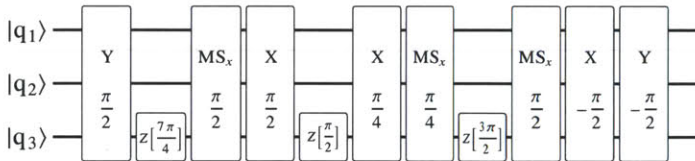


## Order-finding

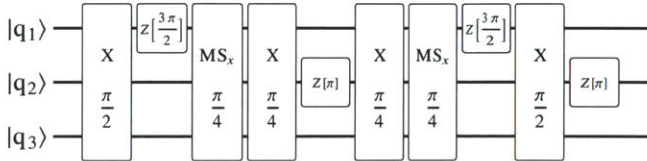
Main sequence.  $\pi^4(y)$ ,  $\pi^2(y)$ ,  $\pi(y)$  are sub-sequences for each choice of permutations, or omitted if identity.



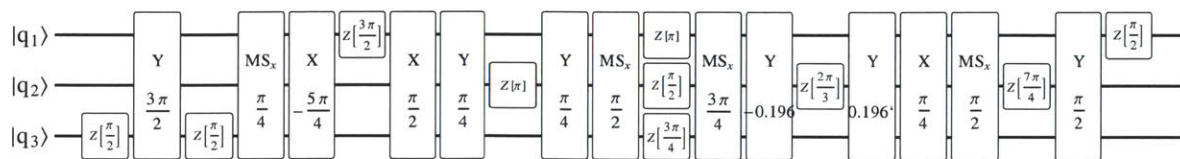
$\pi_1(y)$



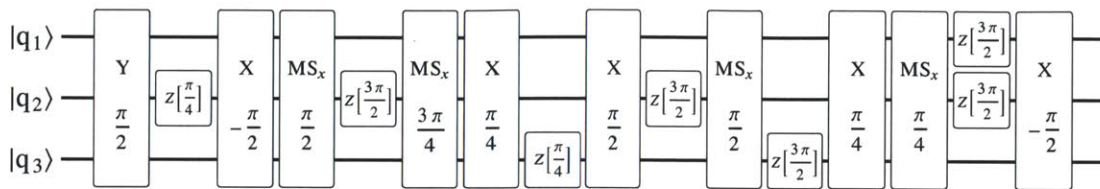
$\pi_2(y)$



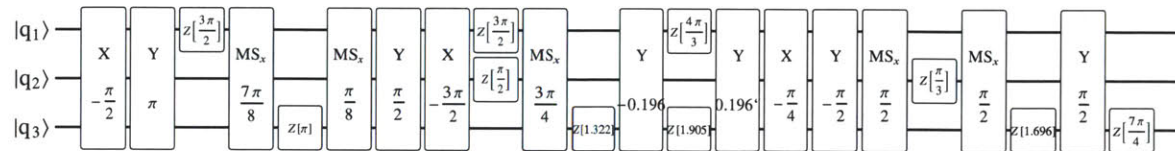
$\pi_3(y)$



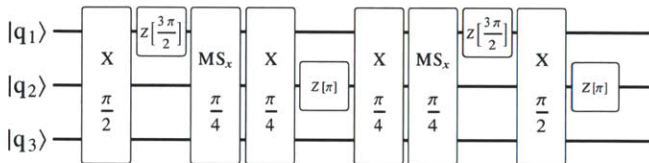
$\pi_3^2(y)$



$\pi_4(y)$

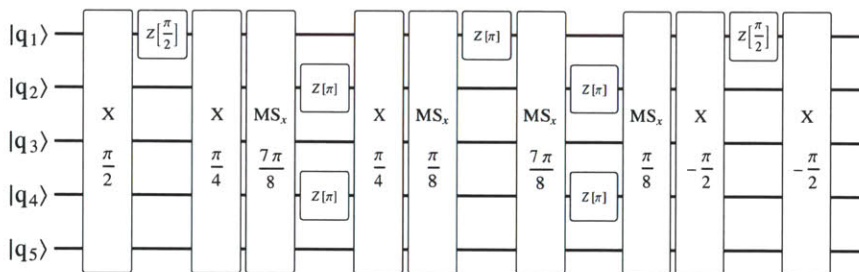


$\pi_4^2(y)$

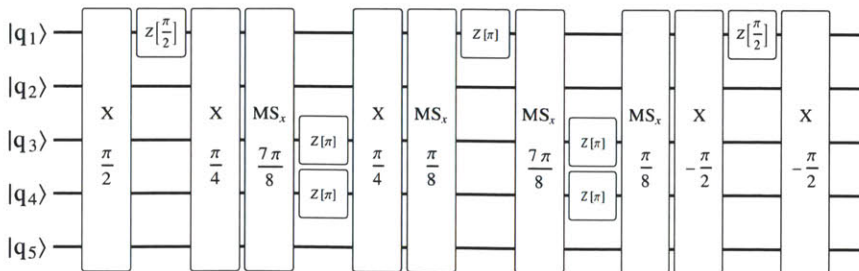


### Shor

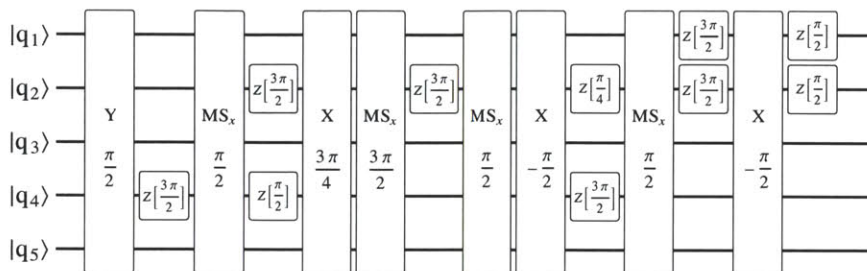
CN(2,4) – CNOT, control ion 1, target ions 2, 4



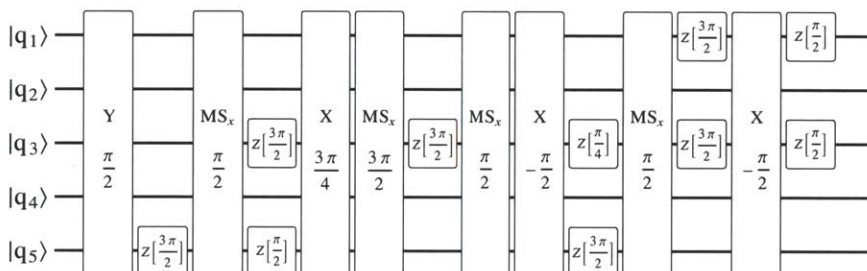
CN(3,4) – CNOT, control ion 1, target ions 3, 4



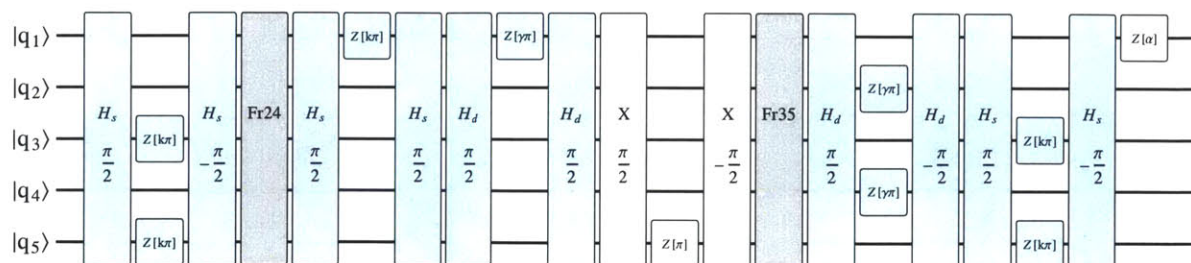
Fr24 – Fredkin, control ion 1, target ions 2, 4



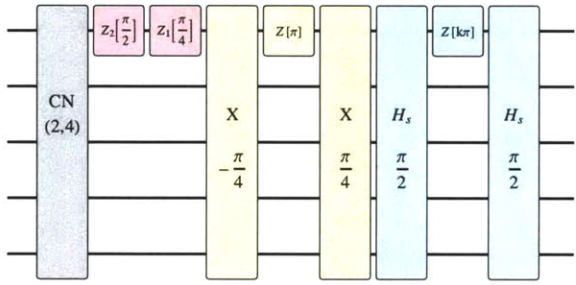
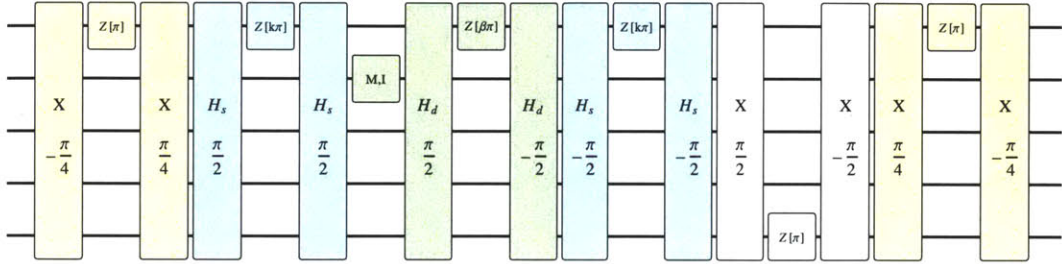
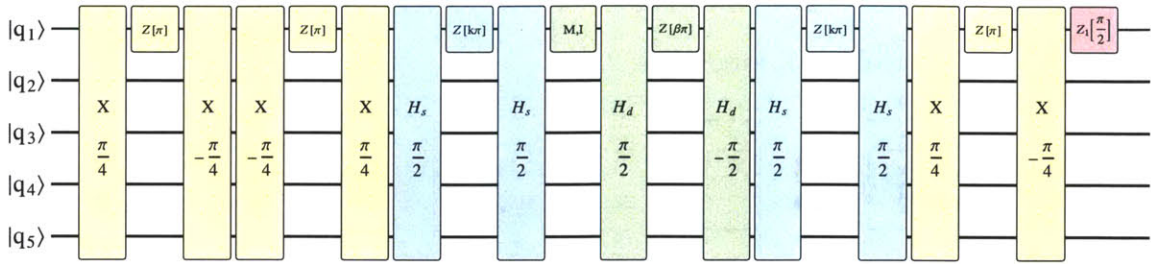
Fr35 – Fredkin, control ion 1, target ions 3, 5



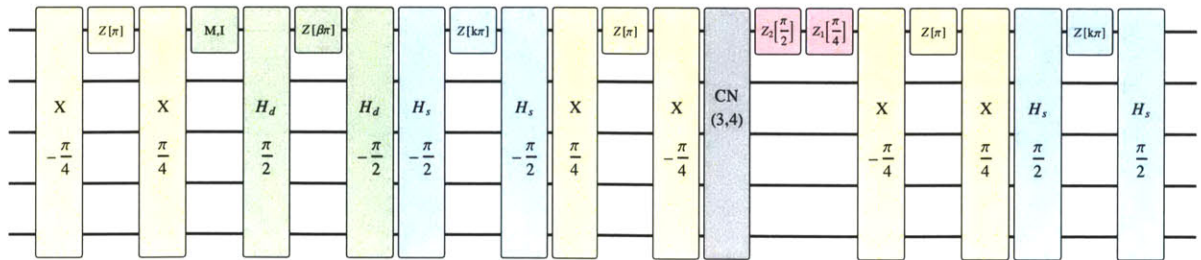
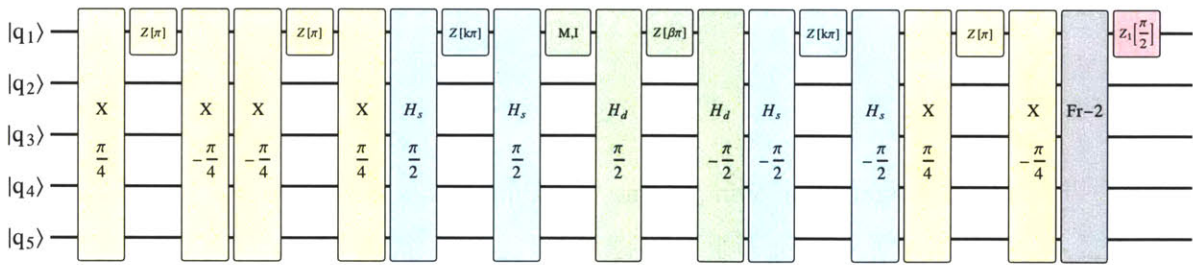
Fr-2 – Both Fredkin gates including hiding/unhiding pulses for the Shor  $a = 7$  sequence



Main sequence for Shor  $a = 11$ .



Main sequence for Shor  $a = 7$ .



# Bibliography

- [AFK<sup>+</sup>07] K. Aoki, J. Franke, T. Kleinjung, A. K. Lenstra, and D. A. Osvik. A kilobit special number field sieve factorization. In *In Advances in Cryptology ASIACRYPT 2007 (2007)*, LNCS, pp. 1–12. SpringerVerlag, 2007.
- [AHJ<sup>+</sup>11] D. Allcock, T. Harty, H. Janacek, N. Linke, C. Ballance, A. Steane, D. Lucas, R. Jarecki, S. Habermehl, M. Blain, D. Stick, and D. Moehring. Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap. *Appl. Phys. B*, pp. 1–7, 2011.
- [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. *Rev. Sci. Instr.*, **80**(1), 013103, 2009.
- [ASS<sup>+</sup>10] D. T. C. Allcock, J. A. Sherman, D. N. Stacey, A. H. Burrell, M. J. Curtis, G. Imreh, N. M. Linke, D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas. Implementation of a symmetric surface-electrode ion trap with field compensation using a modulated raman effect. *New J. Phys.*, **12**(5), 053026, 2010.
- [BB84] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *IEEE International Conference on Computers, Systems and Signal Processing*, p. 175, 1984.
- [BBC<sup>+</sup>95] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, **52**, 3457–3467, 1995.
- [BBM<sup>+</sup>07] A. D. Boozer, A. Boca, R. Miller, T. E. Northup, and H. J. Kimble. Reversible state transfer between light and a single trapped atom. *Phys. Rev. Lett.*, **98**, 193601, 2007.
- [BBT09] N. Bansal, S. Bravyi, and B. M. Terhal. Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs. *Quant. Inf. Comp.*, **9**, 0701, 2009.
- [BC03] A. A. Babaei Brojeny and J. R. Clem. Magnetic-field and current-density distributions in thin-film superconducting rings and disks. *Phys. Rev. B*, **68**, 174514, 2003.
- [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. *Phys. Rev. A*, **75**, 015401, 2007.

- [BDL<sup>+</sup>00] P. A. Barton, C. J. S. Donald, D. M. Lucas, D. A. Stevens, A. M. Steane, and D. N. Stacey. Measurement of the lifetime of the 3d  $D_{5/2}$  state in  $^{40}\text{Ca}^+$ . *Phys. Rev. A*, **62**, 032503, 2000.
- [Ben08] J. Benhelm. *Precision Spectroscopy and quantum information processing with trapped calcium ions*. Ph.D. thesis, Universität Innsbruck, 2008.
- [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. *Phys. Rev. A*, **72**, 050306, 2005.
- [BKRB08] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, **4**, 463, 2008.
- [BLB<sup>+</sup>06] J. Britton, D. Leibfried, J. Beall, R. B. Blakestad, J. J. Bollinger, J. Chiaverini, R. J. Epstein, J. D. Jost, D. Kielpinski, C. Langer, R. Ozeri, R. Reichle, S. Seidelin, N. Shiga, J. H. Wesenberg, and D. J. Wineland. A microfabricated surface-electrode ion trap in silicon. preprint available at <http://arxiv.org/pdf/quant-ph/0605170>, 2006.
- [BM69] W. T. Beall and R. W. Meyerhoff. AC energy losses above and below  $H_{c1}$  in niobium and niobium-25 at.% zirconium. *J. Appl. Phys.*, **40**, 2052, 1969.
- [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. *J. Appl. Phys.*, **83**, 5025, 1998.
- [BMS<sup>+</sup>11] J. T. Barreiro, M. Mueller, 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, 2011.
- [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. *Phys. Rev. Lett.*, **102**(15), 153002, 2009.
- [BPT<sup>+</sup>10] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner. Probing the superfluid-to-mott insulator transition at the single-atom level. *Science*, **329**, 547–550, 2010.
- [Bri08] J. Britton. *Microfabrication techniques for trapped ion quantum information processing*. Ph.D. thesis, University of Colorado, 2008.
- [BSW82] R. Blatt, H. Schnatz, and G. Werth. Ultrahigh-resolution microwave spectroscopy on trapped  $^{171}\text{Yb}^+$  ions. *Phys. Rev. Lett.*, **48**, 1601, 1982.
- [BUB<sup>+</sup>10] M. J. Biercuk, H. Uys, J. W. Britton, A. P. VanDevender, and J. J. Bollinger. Ultrasensitive detection of force and displacement using trapped ions. *Nature Nanotech.*, **5**(9), 646–650, 2010.
- [BW08] R. Blatt and D. Wineland. Entangled states of trapped atomic ions. *Nature*, **453**, 1008, 2008.

- [BWC<sup>+</sup>11] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below  $10^4$  in a trapped ion. *Phys. Rev. A*, **84**, 030303, 2011.
- [CAV94] J. Chang, A. A. Abidi, and C. R. Viswanathan. Flicker noise in CMOS transistors from subthreshold to strong inversion at various temperatures. *Electron Devices, IEEE Transactions on*, **41**(11), 1965–1971, 1994.
- [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 Inf. Comput.*, **5**, 419, 2005.
- [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, 2005.
- [CC00] A. M. Childs and I. L. Chuang. Universal quantum computation with two-level trapped ions. *Phys. Rev. A*, **63**, 012306, 2000.
- [CGT<sup>+</sup>09] J. M. Chow, J. M. Gambetta, L. Tornberg, J. Koch, L. S. Bishop, A. A. Houck, B. R. Johnson, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Randomized benchmarking and process tomography for gate errors in a solid-state qubit. *Phys. Rev. Lett.*, **102**, 090502, 2009.
- [Chw09] M. Chwalla. *Precision spectroscopy with  $^{40}\text{Ca}^+$  ions in a Paul trap*. Ph.D. thesis, Universität Innsbruck, 2009.
- [CKN<sup>+</sup>99] T. Campbell, R. K. Kalia, A. Nakano, P. Vashishta, S. Ogata, and S. Rodgers. Dynamics of oxidation of aluminum nanoclusters using variable charge molecular-dynamics simulations on parallel computers. *Phys. Rev. Lett.*, **82**(24), 4866–4869, 1999.
- [Cle00] R. Cleve. The query complexity of order-finding. *Proceedings, 15th Annual IEEE Conference on Computational Complexity*, 2000.
- [CLS<sup>+</sup>04] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland. Realization of quantum error correction. *Nature*, **432**, 602, 2004.
- [CMQ<sup>+</sup>10] W. C. Campbell, J. Mizrahi, Q. Quraishi, C. Senko, D. Hayes, D. Hucul, D. N. Matsukevich, P. Maunz, and C. Monroe. Ultrafast gates for single atomic qubits. *Phys. Rev. Lett.*, **105**, 090502, 2010.
- [CPO] Charged particle optics program. <http://www.electrooptics.com>.
- [CZ95] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, **74**, 4091, 1995.
- [CZKM97] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Phys. Rev. Lett.*, **78**, 3221, 1997.



- [Dau09] E. Dauler. *Multi-element Superconducting Nanowire Single Photon Detectors*. Ph.D. thesis, Massachusetts Institute of Technology, 2009.
- [DCG<sup>+</sup>09a] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, **460**, 240–244, 2009.
- [DCG09b] R. Dubessy, T. Coudreau, and L. Guidoni. Electric field noise above surfaces: A model for heating-rate scaling law in ion traps. *Phys. Rev. A*, **80**(3), 031402, 2009.
- [DCM92] J. Dalibard, Y. Castin, and K. Mølmer. Wave-function approach to dissipative processes in quantum optics. *Phys. Rev. Lett.*, **68**, 580–583, 1992.
- [DCPS91] F. DiQuarto, C. Centile, S. Piazza, and C. Sunseri. Photoelectrochemical study on anodic aluminum-oxide films – internal photoemission processes at the metal-oxide interface. *Journal of the Electrochemical Society*, **138**, 1956–1861, 1991.
- [Deh67] H. G. Dehmelt. Radiofrequency spectroscopy of stored ions I: Storage. *Adv. At. Mol. Phys.*, **3**, 53, 1967.
- [Des02] M. Deserno. How to generate exponentially correlated gaussian random numbers. [http://www.cmu.edu/biolphys/deserno/pdf/corr\\_gaussian\\_random.pdf](http://www.cmu.edu/biolphys/deserno/pdf/corr_gaussian_random.pdf), 2002.
- [DiV95] D. P. DiVincenzo. Two-bit gates are universal for quantum computation. *Phys. Rev. A*, **51**, 1015, 1995.
- [DiV00] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte Der Physik-Progress of Physics*, **48**, 771, 200.
- [DKM<sup>+</sup>08] M. Debatin, M. Kröner, J. Mikosch, S. Trippel, N. Morrison, M. Reetz-Lamour, P. Woias, R. Wester, and M. Weidemüller. Planar multipole ion trap. *Phys. Rev. A*, **77**(3), 033422, 2008.
- [DKR<sup>+</sup>09] E. A. Dauler, A. J. Kerman, B. S. Robinson, J. K. W. Yang, B. Voronov, G. Goltsman, S. A. Hamilton, and K. K. Berggren. Photon-number-resolution with sub-30-ps timing using multi-element superconducting nanowire single photon detectors. *J. Mod. Optics*, **56**(2-3), 364–373, 2009.
- [DLC<sup>+</sup>09] N. Daniilidis, T. Lee, R. Clark, S. Narayanan, and H. Häffner. Wiring up trapped ions to study aspects of quantum information. *J. Phys. B*, **42**(15), 154012, 2009.
- [DNM<sup>+</sup>11] N. Daniilidis, S. Narayanan, S. A. Möller, R. Clark, T. E. Lee, P. J. Leek, A. Wallraff, St. Schulz, F. Schmidt-Kaler, and H. Häffner. Fabrication and heating rate study of microscopic surface electrode ion traps. *New J. Phys.*, **13**(1), 013032, 2011.



- [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. *Phys. Rev. Lett.*, **97**, 103007, 2006.
- [DPM<sup>+</sup>09] C. L. Degen, M. Poggio, H. J. Mamin, C. T. Rettner, and D. Rugar. Nanoscale magnetic resonance imaging. *Proc. Nat. Acad. Sci. USA*, **106**(5), 1313–1317, 2009.
- [DUB<sup>+</sup>01] S. A. Diddams, Th. Udem, J. C. Bergquist, E. A. Curtis, R. E. Drullinger, L. Hollberg, W. M. Itano, W. D. Lee, C. W. Oates, K. R. Vogel, and D. J. Wineland. An optical clock based on a single trapped  $^{199}\text{Hg}^+$  ion. *Science*, **293**, 825, 2001.
- [EJ96] A. Ekert and R. Jozsa. Quantum computation and Shor’s factoring algorithm. *Rev. Mod. Phys.*, **68**, 733–753, 1996.
- [ESL<sup>+</sup>07] R. J. Epstein, S. Seidelin, D. Leibfried, J. H. Wesenberg, J. J. Bollinger, J. M. Amini, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Simplified motional heating rate measurements of trapped ions. *Phys. Rev. A*, **76**, 033411, 2007.
- [Fey82] R. P. Feynman. Simulating physics with computers. *Intl. J. Theor. Phys.*, **21**, 467, 1982.
- [FL11] S. T. Flammia and Y.-K. Liu. Direct fidelity estimation from few pauli measurements. *Phys. Rev. Lett.*, **106**, 230501, 2011.
- [FMZ<sup>+</sup>10] R. Fermani, T. Müller, B. Zhang, M. J. Lim, and R. Dumke. Heating rate and spin flip lifetime due to near-field noise in layered superconducting atom chips. *J. Phys. B*, **43**(9), 095002, 2010.
- [GCF<sup>+</sup>09] M. Guillorn, J. Chang, N. Fuller, J. Patel, M. Darnon, A. Pyzyna, E. Joseph, S. Engelmann, J. Ott, J. Newbury, D. Klaus, J. Bucchignano, P. Joshi, C. Scerbo, E. Kratschmer, W. Graham, B. To, J. Parisi, Y. Zhang, and W. Haensch. Hydrogen silsesquioxane-based hybrid electron beam and optical lithography for high density circuit prototyping. *J. Vac. Sci. & Tech. B*, **27**(6), 2588–2592, 2009.
- [Gho95] P. K. Ghosh. *Ion Traps*. Clarendon Press, Oxford, 1995.
- [GLN05] A. Gilchrist, N. K. Langford, and M. A. Nielsen. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A*, **71**(6), 062310, 2005.
- [GN96] R. B. Griffiths and C.-S. Niu. Semiclassical fourier transform for quantum computation. *Phys. Rev. Lett.*, **76**, 3228, 1996.
- [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**, 48, 2003.
- [Gro97] L. K. Grover. Quantum computers can search arbitrarily large databases by a single query. *Phys. Rev. Lett.*, **79**, 4709, 1997.

- [GTC<sup>+</sup>91] G. Gabrielse, J. Tan, P. Clateman, L. A. Orozco, S. L. Rolston, C. H. Tseng, and R. L. Tjoelker. A superconducting solenoid system which cancels fluctuations in the ambient magnetic field. *Journal of Magnetic Resonance*, **91**(3), 564–572, 1991.
- [Gul03] S. Gulde. *Experimental realization of quantum gates and the Deutsch-Jozsa algorithm with trapped  $^{40}\text{Ca}^+$  ions*. Ph.D. thesis, Universität Innsbruck, 2003.
- [Gup04] P. Gupta. Pulsed dye laser for excitation of strontium. Master’s thesis, Rice University, 2004.
- [Had09] R. H. Hadfield. Single-photon detectors for optical quantum information applications. *Nat. Photonics*, **3**, 696, 2009.
- [HBHB10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New J. Phys.*, **12**(9), 093035, 2010.
- [HCW<sup>+</sup>11] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, D. P. Pappas, D. Leibfried, and D. J. Wineland. Reduction of anomalous heating in an in-situ-cleaned ion trap. preprint available at <http://arxiv.org/abs/1112.5419>, 2011.
- [HH08] C. Henkel and B. Horovitz. Noise from metallic surfaces: Effects of charge diffusion. *Phys. Rev. A*, **78**(4), 042902, 2008.
- [HHJ<sup>+</sup>09a] D. Hanneke, J. P. Home, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. Realization of a programmable two-qubit quantum processor. *Nature Physics*, **6**, 13–16, 2009.
- [HHJ<sup>+</sup>09b] 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, 2009.
- [HHL09] A. W. Harrow, A. Hassidim, and S. Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, **103**, 150502, 2009.
- [HHR<sup>+</sup>05] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. Scalable multiparticle entanglement of trapped ions. *Nature*, **438**, 643, 2005.
- [HLBH11] M. D. Hughes, B. Lekitsch, J. A. Broersma, and W. K. Hensinger. Microfabricated ion traps. *Contemporary Physics*, **52**(6), 505–529, 2011.
- [Hof05] H. F. Hofmann. Complementary classical fidelities as an efficient criterion for the evaluation of experimentally realized quantum operations. *Phys. Rev. Lett.*, **94**, 160504, 2005.
- [HOS<sup>+</sup>06] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation. *Appl. Phys. Lett.*, **88**, 034101, 2006.

- [Hou08] M. G. House. Analytic model for electrostatic fields in surface-electrode ion traps. *Phys. Rev. A*, **78**, 033402, 2008.
- [HRB08] H. Häffner, C. F. Roos, and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, **469**, 155, 2008.
- [HUG<sup>+</sup>05] W. K. Hensinger, D. W. Utami, H.-S. Goan, K. Schwab, C. Monroe, and G. J. Milburn. Ion trap transducers for quantum electromechanical oscillators. *Phys. Rev. A*, **72**, 041405, 2005.
- [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. *Opt. Lett.*, **36**(16), 3045–3047, 2011.
- [Jam98] D. F. V. James. Quantum dynamics of cold trapped ions with applications to quantum computation. *Appl. Phys. B*, **66**, 181, 1998.
- [JCM<sup>+</sup>09] E. B. Johnson, X. J. Chen, R. Miskimen, D. Von Maluski, C. J. Stapels, S. Mukhopadhyay, F. Augustine, and J. F. Christian. Characteristics of CMOS avalanche photodiodes at cryogenic temperatures. In *Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE*, pp. 2108–2114, 2009.
- [JcvFcvH03] M. Ježek, J. Fiurášek, and Z. Hradil. Quantum inference of states and processes. *Phys. Rev. A*, **68**(1), 012305, 2003.
- [JKMW01] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White. Measurement of qubits. *Phys. Rev. A*, **64**(5), 052312, 2001.
- [JNN11] J. R. Johansson, P. D. Nation, and F. Nori. QuTiP: An open-source python framework for the dynamics of open quantum systems. preprint available at <http://arxiv.org/abs/1110.0573>, 2011.
- [JSN<sup>+</sup>11] A. Jechow, E. W. Streed, B. G. Norton, M. J. Petراسiunas, and D. Kielbinski. Wavelength-scale imaging of trapped ions using a phase fresnel lens. *Opt. Lett.*, **36**(8), 1371–1373, 2011.
- [JSW<sup>+</sup>00] T. Jennewein, C. Simon, G. Weihs, H. Weinfurter, and A. Zeilinger. Quantum cryptography with entangled photons. *Phys. Rev. Lett.*, **84**, 4729–4732, 2000.
- [Kax03] E. Kaxiras. *Atomic and Electronic Structure of Solids*. Cambridge University Press, 2003.
- [KBZ<sup>+</sup>09] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. Deterministic entanglement of ions in thermal states of motion. *New J. Phys.*, **11**(2), 023002, 2009.
- [KDY<sup>+</sup>07] A. J. Kerman, E. A. Dauler, J. K. W. Yang, K. M. Rosfjord, V. Anant, K. K. Berggren, G. N. Gol'tsman, and B. M. Voronov. Constriction-limited detection efficiency of superconducting nanowire single-photon detectors. *Appl. Phys. Lett.*, **90**(10), 101110, 2007.

- [KHC<sup>+</sup>10] B. Kasch, H. Hattermann, D. Cano, T. E. Judd, S. Scheel, C. Zimmermann, R. Kleiner, D. Koelle, and J. Fortágh. Cold atoms near superconductors: atomic spin coherence beyond the Johnson noise limit. *New J. Phys.*, **12**(6), 065024, 2010.
- [KHC11] T. H. Kim, P. F. Herskind, and I. L. Chuang. Surface-electrode ion trap with integrated light source. *Appl. Phys. Lett.*, **98**(21), 214103, 2011.
- [Kim08] H. J. Kimble. The quantum internet. *Nature*, **453**(7198), 1023–1030, 2008.
- [Kim11] T. H. Kim. An optical-fiber interface to a trapped-ion quantum computer. Master’s thesis, Massachusetts Institute of Technology, 2011.
- [KK09] J. Kim and C. Kim. Integrated optical approach to trapped ion quantum computation. *Quantum Inf. Comput.*, **9**(3&4), 0181–0202, 2009.
- [KLM06] S. Kuehn, R. F. Loring, and J. A. Marohn. Dielectric fluctuations and the origins of noncontact friction. *Phys. Rev. Lett.*, **96**, 156103, 2006.
- [KMW02] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture of a large-scale ion trap quantum computer. *Nature*, **417**, 709, 2002.
- [Kni10] E. Knill. News and views Q&A: Physics: Quantum computing. *Nature*, **463**, 441–443, 2010.
- [Kob94] N. Koblitz. *A Course in Number Theory and Cryptography*. Springer-Verlag, New York, 1994.
- [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. *Appl. Phys. Lett.*, **74**(23), 3444–3446, 1999.
- [KPM<sup>+</sup>05] J. Kim, S. Pau, Z. Ma, H. R. McLellan, J. V. Gates, A. Kornblit, R. E. Slusher, R. M. Jopson, I. Kang, and M. Dinu. System design for large-scale ion trap quantum information processor. *Quantum Inf. Comput.*, **5**, 515–537, 2005.
- [Lab08] J. Labaziewicz. *High Fidelity Quantum Gates with Ions in Cryogenic Microfabricated Ion Traps*. Ph.D. thesis, Massachusetts Institute of Technology, 2008.
- [Lac10] N. S. Lachenmyer. Measurements of electric field noise and light-induced charging in cryogenic surface electrode ion traps. B. S. Thesis, Massachusetts Institute of Technology, 2010.
- [LBY<sup>+</sup>11] E. Lucero, R. Barends, Y. Chen, J. Kelly, M. Mariantoni, A. Megrant, P. O’Malley, D. Sank, A. Vainsencher, J. Wenner, T. White, Y. Yin, A. N. Cleland, and J. M. Martinis. Computing prime factors with a Josephson phase qubit quantum processor. preprint available at <http://arxiv.org/abs/1202.5707>, 2011.
- [LDM<sup>+</sup>03] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, **422**, 412, 2003.

- [Lei09] D. Leibbrandt. *Integrated chips and optical cavities for trapped ion quantum information processing*. Ph.D. thesis, Massachusetts Institute of Technology, 2009.
- [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. *Phys. Rev. Lett.*, **100**, 013001, 2008.
- [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. *Phys. Rev. Lett.*, **101**, 180602, 2008.
- [LHC11] G. H. Low, P. F. Herskind, and I. L. Chuang. Finite-geometry models of electric field noise from patch potentials in ion traps. *Phys. Rev. A*, **84**, 053425, 2011.
- [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**(6052), 57–61, 2011.
- [lib] libquantum: the C library for quantum computing and quantum simulation.
- [LJL<sup>+</sup>10] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. *Nature*, **464**(7285), 45–53, 2010.
- [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. Stick, 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 Inf. Comput.*, **9**(11), 0901, 2009.
- [LLL<sup>+</sup>11] E. M. López, A. Laing, T. Lawson, R. Alvarez, X. Q. Zhou, and J. L. O’Brien. Implementation of an iterative quantum order finding algorithm. preprint available at <http://arxiv.org/abs/1111.4147>, 2011.
- [LOJ<sup>+</sup>05] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-lived qubit memory using atomic ions. *Phys. Rev. Lett.*, **95**, 060502, 2005.
- [LRB<sup>+</sup>07] J. Labaziewicz, P. Richerme, K. R. Brown, I. L. Chuang, and K. Hayasaka. Compact, filtered diode laser system for precision spectroscopy. *Opt. Lett.*, **32**, 572, 2007.
- [LWL<sup>+</sup>07] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White. Experimental demonstration of a compiled version of Shor’s algorithm with quantum entanglement. *Phys. Rev. Lett.*, **99**, 250505, 2007.
- [MBH<sup>+</sup>04] H. S. Margolis, G. P. Barwood, G. Huang, H. A. Klein, S. N. Lea, K. Szymaniec, and P. Gill. Hertz-level measurement of the optical clock frequency in a single  $^{88}\text{Sr}^+$  ion. *Science*, **306**, 1355, 2004.

- [MBK95] W. Mende, K. Bartschat, and M. Kock. Near-threshold photoionization from the Sr I (5s5p)  $^1P_1^0$  state. *J. Phys. B*, **28**, 2385, 1995.
- [MCD93] K. Mølmer, Y. Castin, and J. Dalibard. Monte Carlo wave-function method in quantum optics. *J. Opt. Soc. Am. B*, **10**(3), 524–538, 1993.
- [MFS<sup>+</sup>08] S. Miki, M. Fujiwara, M. Sasaki, B. Baek, A. J. Miller, R. H. Hadfield, S. W. Nam, and Z. Wang. Large sensitive-area NbN nanowire superconducting single-photon detectors fabricated on single-crystal MgO substrates. *Appl. Phys. Lett.*, **92**(6), 061116, 2008.
- [MGC<sup>+</sup>92] H. Minami, Q. Geng, K. Chihara, J. Yuyama, and E. Goto. Detection of trapped flux quanta in superconducting films by scanning a SQUID pick-up coil. *Cryogenics*, **32**(7), 648–652, 1992.
- [MHS<sup>+</sup>11] D. L. Moehring, C. Highstrete, D. Stick, K. M. Fortier, R. Haltli, C. Tigges, and M. G. Blain. Design, fabrication and experimental demonstration of junction surface ion traps. *New J. Phys.*, **13**(7), 075018, 2011.
- [MKH09a] K. S. McKay, J. Kim, and H. H. Hogue. Enhanced quantum efficiency of the visible light photon counter in the ultraviolet wavelengths. *Opt. Express*, **17**(9), 7458–7464, 2009.
- [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. *Phys. Rev. Lett.*, **102**(4), 040501, 2009.
- [MLB<sup>+</sup>09] R. Maiwald, D. Leibfried, J. Britton, J. C. Bergquist, G. Leuchs, and D. J. Wineland. Stylus ion trap for enhanced access and sensing. *Nature Physics*, **5**(8), 551–554, 2009.
- [MMK<sup>+</sup>95] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.*, **75**, 4714, 1995.
- [MN73] R. G. Meyer and L. Nagel. Computer simulation of 1/f noise performance of electronic circuits. *Solid-State Circuits, IEEE Journal of*, **8**(3), 237–240, 1973.
- [Mon11] T. Monz. *Quantum Information Processing Beyond Ten Ion-Qubits*. Ph.D. thesis, Universität Innsbruck, 2011.
- [MS99] K. Mølmer and A. Sørensen. Multiparticle entanglement of hot trapped ions. *Phys. Rev. Lett.*, **82**, 1835, 1999.
- [MS12] I. L. Markov and M. Saeedi. Constant-optimized quantum circuits for modular multiplication and exponentiation. *Quantum Inf. Comput.*, **12**, 361–394, 2012.
- [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. *Phys. Rev. Lett.*, **106**, 130506, 2011.

- [MSGW92] A. Mathai, D. Song, Y. Gim, and F. C. Wellstood. One-dimensional magnetic flux microscope based on the dc superconducting quantum interference device. *Appl. Phys. Lett.*, **61**(5), 598–600, 1992.
- [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. *Phys. Rev. Lett.*, **100**, 200502, 2008.
- [MvdS99] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, New York, 1999.
- [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 J. Phys.*, **13**(10), 103005, 2011.
- [MZF<sup>+</sup>10] T. Müller, B. Zhang, R. Fermani, K. S. Chan, Z. W. Wang, C. B. Zhang, M. J. Lim, and R. Dumke. Trapping of ultra-cold atoms with the magnetic field of vortices in a thin-film superconducting micro-structure. *New J. Phys.*, **12**(4), 043016, 2010.
- [Naj11] F. Najafi. personal communication, 2011.
- [NC00] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, UK, 2000.
- [Neb08] V. Nebendahl. Optimierung verschränkender quantengatter für experimente mit ionenfallen. Master’s thesis, Universität Innsbruck, 2008.
- [NHR09] V. Nebendahl, H. Häffner, and C. F. Roos. Optimal control of entangling operations for trapped-ion quantum computing. *Phys. Rev. A*, **79**, 012312, 2009.
- [Nie02] M. A. Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, **303**(4), 249–252, 2002.
- [NP73] L. W. Nagel and D. O. Pederson. SPICE (simulation program with integrated circuit emphasis). Technical Report UCB/ERL M382, EECS Department, University of California, Berkeley, 1973.
- [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.
- [OPG<sup>+</sup>04] J. L. O’Brien, G. J. Pryde, A. Gilchrist, D. F. V. James, N. K. Langford, T. C. Ralph, and A. G. White. Quantum process tomography of a controlled-NOT gate. *Phys. Rev. Lett.*, **93**(8), 080502, 2004.
- [OWC<sup>+</sup>11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, **476**(7359), 181–184, 2011.

- [PFCP07] C. P. Poole, H. A. Farach, R. J. Creswick, and R. Prozorov. *Superconductivity*. Academic Press, London, UK, 2007.
- [PLB<sup>+</sup>06] C. E. Pearson, D. R. Leibbrandt, W. S. Bakr, W. J. Mallard, K. R. Brown, and I. L. Chuang. Experimental investigation of planar ion traps. *Phys. Rev. A*, **73**, 032307, 2006.
- [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, 2009.
- [Pos10] U. G. Poschinger. *Quantum Optics Experiments in a Microstructured Ion Trap*. Ph.D. thesis, Universität Ulm, 2010.
- [pp] Parallel python. <http://www.parallelpython.com/>.
- [PP00] S. Parker and M. B. Plenio. Efficient factorization with a single pure qubit and  $\log N$  mixed qubits. *Phys. Rev. Lett.*, **85**, 3049–3052, 2000.
- [PS53] W. Paul and H. Steinwedel. Ein neues massenspektrometer ohne magnetfeld. *Z. Naturforsch. A*, **8**, 448, 1953.
- [PSB<sup>+</sup>11] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture. *Phys. Rev. Lett.*, **107**, 240501, 2011.
- [QMP12] Quantum machine parametrizer, 2012. <http://www.quantum.gatech.edu/quantumMachineParameterizer.shtml>.
- [Qua89] T. L. Quarles. *Analysis of Performance and Convergence Issues for Circuit Simulation*. Ph.D. thesis, EECS Department, University of California, Berkeley, 1989.
- [RCK<sup>+</sup>06] C. F. Roos, M. Chwalla, K. Kim, M. Riebe, and R. Blatt. ‘Designer atoms’ for quantum metrology. *Nature*, **443**, 316–319, 2006.
- [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, 2012.
- [RHS<sup>+</sup>08] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency ratio of  $\text{Al}^+$  and  $\text{Hg}^+$  single-ion optical clocks; metrology at the 17th decimal place. *Science*, **319**, 1808, 2008.
- [Rie05] M. Riebe. *Preparation of entangled states and quantum teleportation with atomic qubits*. Ph.D. thesis, Universität Innsbruck, 2005.
- [RKS<sup>+</sup>06] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt. Process tomography of ion trap quantum gates. *Phys. Rev. Lett.*, **97**, 220407, 2006.



- [RLR<sup>+</sup>04] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Häffner, W. Hänsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Bell states of atoms with ultralong lifetimes and their tomographic state analysis. *Phys. Rev. Lett.*, **92**(22), 220402, 2004.
- [RMR<sup>+</sup>07] K. De Raedt, K. Michielsen, H. De Raedt, B. Trieu, G. Arnold, M. Richter, Th. Lippert, H. Watanabe, and N. Ito. Massively parallel quantum computer simulator. *Computer Physics Communications*, **176**(2), 121–136, 2007.
- [RNMW71] R. Rohrer, L. Nagel, R. Meyer, and L. Weber. Computationally efficient electronic-circuit noise calculations. *Solid-State Circuits, IEEE Journal of*, **6**(4), 204–213, 1971.
- [Roo00] C. F. Roos. *Controlling the quantum state of trapped ions*. Ph.D. thesis, Universität Innsbruck, 2000.
- [RYD<sup>+</sup>06] K. M. Rosfjord, J. K. W. Yang, E. A. Dauler, A. J. Kerman, V. Anant, B. M. Voronov, G. N. Gol'tsman, and K. K. Berggren. Nanowire single-photon detector with an integrated optical cavity and anti-reflection coating. *Opt. Express*, **14**(2), 527–534, 2006.
- [SBC<sup>+</sup>11] D. I. Schuster, L. S. Bishop, I. L. Chuang, D. DeMille, and R. J. Schoelkopf. Cavity QED in a molecular ion trap. *Phys. Rev. A*, **83**, 012311, 2011.
- [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, 2011.
- [Sch08] P. Schindler. Frequency synthesis and pulse shaping for quantum information processing with trapped ions. Master's thesis, Universität Innsbruck, 2008.
- [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. *J. Opt. Soc. Am. B*, **28**(12), 2865–2870, 2011.
- [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, J.D. Jost, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Phys. Rev. Lett.*, **96**, 253003, 2006.
- [Sem69] Y. L. Semov. Work function of oxidized metal surfaces and estimation of Al<sub>2</sub>O<sub>3</sub> film band structure parameters. *Physica Status Solidi*, **32**, K41, 1969.
- [She08] R. Shewmon. Coherent manipulations of trapped <sup>88</sup>Sr<sup>+</sup> using the 4D<sub>5/2</sub> → 5S<sub>1/2</sub> transition. B. S. Thesis, Massachusetts Institute of Technology, 2008.
- [Sho52] W. Shockley. Transistor electronics: Imperfections, unipolar and analog transistors. *Proceedings of the IRE*, **40**(11), 1289–1313, 1952.
- [Sho94] P. W. Shor. Algorithms for quantum computation: Discrete log and factoring. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on*

*the Foundations of Computer Science*, p. 124, IEEE Computer Society, Los Alamitos, CA, 1994.

- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. on Comp.*, **26**, 1484, 1997.
- [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**, 36, 2006.
- [SKGR<sup>+</sup>03] F. Schmidt-Kaler, S. Gulde, M. Riebe, T. Deuschle, A. Kreuter, G. Lancaster, C. Becher, J. Eschner, H. Häffner, and R. Blatt. The coherence of qubits based on single Ca<sup>+</sup> ions. *J. Phys. B*, **36**(3), 623–636, 2003.
- [SKHR<sup>+</sup>03] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, Ch. Roos, J. Eschner, and Rainer Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, **422**, 408, 2003.
- [SM99] A. Sørensen and K. Mølmer. Quantum computation with ions in thermal motion. *Phys. Rev. Lett.*, **82**, 1971, 1999.
- [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. *Phys. Rev. Lett.*, **87**(9), 096801, 2001.
- [SNRWS11] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour. Microscopic model of electric-field-noise heating in ion traps. *Phys. Rev. A*, **84**, 023412, 2011.
- [SNTK03] A. Stognij, N. Novitskii, S. Tushina, and S. Kalinnikov. Preparation of ultrathin gold films by oxygen-ion sputtering and their optical properties. *Technical Physics*, **48**, 745–748, 2003.
- [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*. *Rev. Mod. Phys.*, **82**, 2609–2632, 2010.
- [SPZSK08] S. A. Schulz, U. Poschinger, F. Ziesel, and F. Schmidt-Kaler. Sideband cooling and coherent dynamics in a microchip multi-segmented ion trap. *New J. Phys.*, **10**(4), 045007, 2008.
- [Ste03a] A. M. Steane. Overhead and noise threshold of fault-tolerant quantum error correction. *Phys. Rev. A*, **68**, 042322, 2003.
- [Ste03b] M. Steffen. *A Prototype Quantum Computer Using Nuclear Spins in Liquid Solution*. Ph.D. thesis, Stanford University, 2003.
- [Ste07] A. M. Steane. How to build a 300 bit, 1 giga-operation quantum computer. *Quantum Inf. Comput.*, **7**, 171, 2007.
- [SVG<sup>+</sup>03] R. Sobolewski, A. Verevkin, G. N. Gol'tsman, A. Lipatov, and K. Wilsher. Ultrafast superconducting single-photon optical detectors and their applications. *Applied Superconductivity, IEEE Transactions on*, **13**(2), 1151–1157, 2003.

- [Tan99] S. M. Tan. A computational toolbox for quantum and atomic optics. *J. Opt. B: Quantum Semiclass. Opt.*, **1**(4), 424, 1999.
- [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 mechanical ground state. *Phys. Rev. A*, **61**, 063418, 2000.
- [TRBZ04] L. Tian, P. Rabl, R. Blatt, and P. Zoller. Interfacing quantum-optical and solid-state qubits. *Phys. Rev. Lett.*, **92**, 247902, 2004.
- [UTSM<sup>+</sup>07] R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Oemer, M. Fuerst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger. Entanglement-based quantum communication over 144km. *Nature Physics*, **3**(7), 481–486, 2007.
- [Van01] L. Vandersypen. *Experimental Quantum Computation with Nuclear Spins in Liquid Solution*. Ph.D. thesis, Stanford University, 2001.
- [VC05] L. M. K. Vandersypen and I. L. Chuang. Nmr techniques for quantum control and computation. *Rev. Mod. Phys.*, **76**, 1037–1069, 2005.
- [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. *Phys. Rev. Lett.*, **105**, 023001, 2010.
- [vdZ70] A. van der Ziel. Noise in solid-state devices and lasers. *Proceedings of the IEEE*, **58**(8), 1178–1206, 1970.
- [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. *Phys. Rev. Lett.*, **85**, 5452–5455, 2000.
- [VSB<sup>+</sup>01] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, R. Cleve, and I. L. Chuang. Experimental realization of Shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature*, **414**, 883, 2001.
- [VSS<sup>+</sup>99] L. M. K. Vandersypen, M. Steffen, M. H. Sherwood, C. S. Yannoni, G. Breyta, and I. L. Chuang. Implementation of a three-quantum bit search algorithm. *Appl. Phys. Lett.*, **76**, 646–648, 1999.
- [Wes08] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Phys. Rev. A*, **78**(6), 063410, 2008.
- [WGL<sup>+</sup>10] S. X. Wang, Y. Ge, J. Labaziewicz, E. Dauler, K. Berggren, and I. L. Chuang. Superconducting microfabricated ion traps. *Appl. Phys. Lett.*, **97**(24), 244102, 2010.
- [WLG<sup>+</sup>09] S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, and I. L. Chuang. Individual addressing of ions using magnetic field gradients in a surface-electrode ion trap. *Appl. Phys. Lett.*, **94**, 094103, 2009.

- [WLG<sup>+</sup>10] S. X. Wang, J. Labaziewicz, Y. Ge, R. Shewmon, and I. L. Chuang. Demonstration of a quantum logic gate in a cryogenic surface-electrode ion trap. *Phys. Rev. A*, **81**(6), 062332, 2010.
- [WLL<sup>+</sup>11] S. X. Wang, G. H. Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, and I. L. Chuang. Laser-induced charging of microfabricated ion traps. *J. Appl. Phys.*, **110**(10), 104901, 2011.
- [WMI<sup>+</sup>98] 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. *J. Res. Natl. Inst. Stand. Technol.*, **103**, 259, 1998.
- [WPF<sup>+</sup>01] Y. S. Weinstein, M. A. Pravia, E. M. Fortunato, S. Lloyd, and D. G. Cory. Implementation of the quantum fourier transform. *Phys. Rev. Lett.*, **86**, 1889–1891, 2001.
- [XZL<sup>+</sup>11] N. Xu, J. Zhu, D. Lu, X. Zhou, X. Peng, and J. Du. Quantum factorization of 143 on a dipolar-coupling NMR system. preprint available at <http://arxiv.org/abs/1111.3726>, 2011.
- [Yan05] X. M. Yang. Hybrid exposure strategy: Combining e-beam direct writing with optical lithography for magnetic recording heads. *J. Vac. Sci. & Tech. B*, **23**(6), 2624–2630, 2005.
- [Yan09] J. K. Yang. *Superconducting Nanowire Single-Photon Detectors and Sub-10-nm Lithography*. Ph.D. thesis, Massachusetts Institute of Technology, 2009.
- [YKD<sup>+</sup>07] J. K. W. Yang, A. J. Kerman, E. A. Dauler, V. Anant, K. M. Rosfjord, and K. K. Berggren. Modeling the electrical and thermal response of superconducting nanowire single-photon detectors. *Applied Superconductivity, IEEE Transactions on*, **17**(2), 581–585, 2007.
- [zvo] Zvode: Variable-coefficient ordinary differential equation solver, with fixed-leading-coefficient implementation. <http://www.netlib.org/ode/zvode.f>.