

Integrated optical quantum manipulation and measurement of trapped ions

by

Karan K. Mehta

B.S., University of California, Los Angeles (2010)

S.M., Massachusetts Institute of Technology (2012)

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

Doctor of Philosophy in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2017

© Massachusetts Institute of Technology 2017. All rights reserved.

Author
Department of Electrical Engineering
December 22, 2016

Certified by.....
Rajeev J. Ram
Professor of Electrical Engineering
Thesis Supervisor

Accepted by.....
Leslie A. Kolodziejski
Professor of Electrical Engineering
Chair, Department Committee on Graduate Students

Integrated optical quantum manipulation and measurement of trapped ions

by
Karan K. Mehta

Submitted to the Department of Electrical Engineering
on December 22, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering

Abstract

Individual atomic ions confined in designed electromagnetic potentials and manipulated via lasers are strong candidates as physical bases for quantum information processing (QIP). This is in large part due to their long coherence times, indistinguishability, and strong Coulomb interactions. Much work in recent years has utilized these properties to implement increasingly precise quantum operations essential for QIP, as well as to conduct increasingly sophisticated experiments on few-ion systems. Many questions remain however regarding how to implement the significant classical apparatus required to control and measure many ions (and indeed any physical qubit under study) in a scalable way that furthermore does not compromise qubit quality.

This work draws on techniques in integrated optics to address this question. Planar-fabricated waveguides and gratings integrated with planar ion traps are demonstrated to allow optical addressing of individual $^{88}\text{Sr}^+$ ions $50\ \mu\text{m}$ above the chip surface with diffraction-limited focused beams, with advantages in stability and scalability. Motivated by the requirement for low crosstalk in qubit addressing, we show also that intuitively designed devices can generate precisely tailored intensity profiles at the ion locations, with diffraction-limited sidelobe intensities characterized to the 5×10^{-6} level in relative intensity up to $25\ \mu\text{m}$ from the focus. Such devices can be implemented alongside complex systems in complementary metal-oxide-semiconductor (CMOS) processes. We show in addition that the multiple patternable metal layers present in CMOS processes can be used to create complex planar ion traps with performance comparable to simple single-layer traps, and that CMOS silicon avalanche photodiodes may be employed for scalable quantum state readout. Finally we show initial results on integrated electro-optic modulators for visible light.

These results open possibilities for experiments with trapped ions in the short term, and indicate routes to achieving large-scale systems of thousands or more ions in the future. Though ion qubits may seem isolated from scalable solid-state technologies, it appears this apparent isolation may uniquely allow a cooperation with complex planar-fabricated optical and electronic systems without introducing additional decoherence.

Thesis Supervisor: Rajeev J. Ram
Title: Professor of Electrical Engineering

Acknowledgments

It is a pleasure to reflect on all the people I can thank for having somehow made possible the work below. I'm very grateful to have been able to work with Rajeev Ram, and for his guidance over these years; this project absolutely wouldn't have happened without him, and I think I'm aware (but probably not completely) of the patience he's had with me at various points along this course. His intuition and sense for problems and solutions has been valuable sometimes in ways that I would realize later, and without his openness to discussion and criticism I'd have learned and done far less.

I'm fortunate also to have been able to work with John Chiaverini and Jeremy Sage. Their expertise in trapped-ion experiment and enthusiasm in the project was invaluable in getting the project going, and I've continued to learn from them over the last few years. I'm thankful to them for their welcoming me into their group. Colin Bruzewicz and Robert McConnell were great people to work alongside in the lab and very helpful when questions arose.

I owe big thanks to many members of our group, past and present. I learned a lot from Jason Orcutt early in graduate school, much of which has shaped this work in significant ways. Parthiban Santhanam I had many good discussions with, technical and otherwise; I particularly remember some around the time when I was trying to define a thesis topic which were very meaningful. Amir Atabaki more recently has always been willing to discuss any problems that arise, and has helped me out many times with issues in fabrication. Peter Han, Bill Herrington, Huaiyu Meng, Luca Alloatti, Jin Xue, Zheng Li, and Neerja Aggarwal, have been great coworkers and always helpful labmates. It has been a privilege also to get to work with Gavin West and Kramnik Daniel in the last few months as they begin along some related lines.

Isaac Chuang helped us start to see the virtues of trapped ions early on and encouraged us to start along this path; I'm thankful to him for that, as well as many stimulating conversations at various points along the way since. I want to thank Amira Eltony as well, for many encouraging early discussions and work together on CMOS traps, and Michael Gutierrez for discussions and camaraderie throughout, not to mention help in several moves in Cambridge and many good dinners.

A number of people in neighboring groups also deserve mention for lending time to help at various points. Jie Sun and Ehsan Hosseini helped me get started fabricating waveguide devices, and their early pointers and willingness to spend time discussing were a big help. Faraz Najafi spent much time with me getting the cryostat used for APD measurements here working; Adam McCaughan, also from Karl Berggren's group was a great help in questions related to the cryostat as well. Min Sun and Tomas Palacios deserve thanks as well for lending equipment used for the high-voltage DC breakdown tests of the first CMOS trap.

Mark Mondol's work on MIT's e-beam lithography systems and willingness to discuss problems/solutions were invaluable for fabrication, as was James Daley's work in keeping many parts of the NSL running. I'd also like to thank the MTL staff, and in particular Bernard Alamariu, Bob Bicchieri, and Vicky Diadiuk for the constant effort that goes into keeping these facilities and tools running well.

I am fortunate to have had support from a DOE Science Graduate Fellowship for the first few years of my graduate work; an NSF grant that supported much of the work in this thesis; from the iQuISE IGERT program at MIT; as well as from the Integrated Quantum Initiative at Lincoln Labs/MIT.

Finally, I don't think I could imagine any of this without the support and encouragement I've had from my mother, father, and brother during these years and before – thank you.

Contents

Cover page	1
Abstract	3
Acknowledgments	5
Contents	7
List of Figures	11
1 Introduction: quantum manipulations of individual ions	15
1.1 Physics of trapped ion qubits	16
1.1.1 Radio-frequency electromagnetic Paul traps for ions	16
1.1.2 Carrier electronic transitions, single qubit operations	18
1.1.3 Coupling to a motional mode; the Lamb-Dicke regime	20
1.1.4 Two-qubit operations	22
1.1.5 Atomic structure of commonly used ions; hyperfine and quadrupole qubits	24
1.1.6 Fluorescence readout of qubit state	27
1.1.7 Cooling of trapped ions	28
1.2 Experimental implementation and approaches to scalable optics	28
1.2.1 Bulk trap experiments	29
1.2.2 Planar electrode trap experiments	29
1.2.3 Approaches to scalable optics in planar traps; overview of work presented here	31
1.3 Summary of scaling challenges and those addressed here	32
1.3.1 Scalable control apparatus	33
1.3.2 Heating rates	33
1.3.3 Limitations to two-qubit gate fidelities, and addressing	34
1.3.4 Speed limitations on motional-mode coupling in multi-ion systems	35
1.3.5 General ion movement, interconnections between distant ions	35
1.4 Concluding comments and thesis overview	36
2 Overview of the proposed approach	37
2.1 Routing, shaping, beam-forming, and detecting light within a trap chip	38
2.2 Waveguide-based routing and ion addressing	39

2.2.1	Focusing couplers for ion addressing	41
2.3	Advantages and limitations of waveguide-based routing and focusing .	42
2.3.1	Laser power	42
2.3.2	Raman gate operation, detuning and spontaneous scattering .	43
2.3.3	Pointing instability	44
2.3.4	Phase stability	47
2.3.5	Individual addressing	47
2.3.6	Wavelength ranges and polarizations	49
2.3.7	Power handling and self-phase modulation effects	49
2.4	Parallel modulation	51
2.5	CMOS integration and on-chip photodetection	52
2.6	Summary and overview of chapters ahead	55
3	Design, fabrication, characterization of focusing waveguide grating couplers	57
3.1	Overview and previous work	57
3.2	Grating design	58
3.2.1	Dimensional considerations for ion addressing	58
3.2.2	Longitudinal design	61
3.2.3	Transverse design	64
3.2.4	Predicted focus distance and waist for a given design	66
3.3	Layout, fabrication, and material characterization	67
3.3.1	Silicon nitride deposition and loss	69
3.4	Optical characterization	70
3.4.1	Detailed characterization at focus	72
3.4.2	Polarization purity	76
3.5	Discussion	78
3.5.1	Possible improvements to sidelobe suppression in present devices	78
3.6	Conclusion and future work	80
4	Ion addressing with waveguide optics	81
4.1	Overview	81
4.2	Integrated waveguide trap design and fabrication	82
4.2.1	Choice of substrate	82
4.2.2	PECVD and LPCVD SiN deposition	83
4.2.3	Chip layout and design	84
4.2.4	Grating design	86
4.2.5	Device fabrication	86
4.3	Grating optical properties	88
4.3.1	Input grating couplers	89
4.4	Individual ion experiments	90
4.4.1	Coherent manipulations	93
4.4.2	Individual addressing	95
4.4.3	Crosstalk quantification and comparison to existing experiments	95
4.4.4	Optical losses	98

4.4.5	Polarization purity	99
4.4.6	Photo-induced charging and stability	101
4.5	Conclusion and future work	101
4.5.1	Direct fiber coupling	102
4.5.2	Waveguide losses	102
5	CMOS integration and silicon avalanche photodiodes	103
5.1	CMOS-fabricated ion trap	103
5.1.1	Design and implementation	103
5.1.2	CMOS trap characterization	109
5.2	Integrated silicon avalanche photodiodes	111
5.2.1	Trap-integrated device layout and design	112
5.2.2	CMOS APD designs and room temperature DC characteristics	113
5.2.3	STI-GR device characterization and low-temperature measure- ments	117
5.2.4	Full device characterization	125
5.3	Conclusion and summary	126
6	Visible electro-optic modulation	127
6.1	Approaches to modulation	127
6.2	Silicon nitride/Lithium niobate hybrid waveguides	128
6.2.1	Simulated waveguide properties	128
6.2.2	Approximate expected electro-optic properties and modulator designs	132
6.3	Fabrication	133
6.3.1	Electron-beam proximity effect on LiNbO ₃	133
6.4	Sample 1: uncontacted waveguide devices	136
6.4.1	Designs included and brief background on ring resonators . . .	136
6.4.2	Characterization of IR ring resonators and MZIs	139
6.4.3	Visible waveguides	143
6.5	Sample 2: contacted MZ and ring devices	143
6.5.1	Fabrication	144
6.5.2	674 nm Mach-Zehnder modulator	145
6.5.3	EO tunable infrared rings	147
6.6	Conclusion	148
7	Conclusion, future work, and outlook	149
7.1	WG devices for blue and UV wavelengths	149
7.1.1	Alumina as a core material for blue/UV wavelengths	150
7.1.2	Materials for electro-optic modulators	151
7.1.3	Characterization of waveguide nonlinearities and damage thresh- olds in the visible	151
7.2	Direct fiber attaching, and silicon substrates	152
7.3	Short-term ion experiments	153
7.3.1	High-fidelity, fast two-qubit gates	153

7.3.2	Transport gates in a 2D geometry	153
7.3.3	Optical addressing of microwave operations	154
7.4	APDs and devices for ion movement	154
7.5	Errors and scale	155
7.5.1	Noise/infidelity sources and magnitudes in this approach . . .	155
7.5.2	Concluding remark, integration with isolated qubits	158
A	Motivation for constraint on curvature radius in grating design	159
B	E-beam write order and focusing grating sidelobe profiles	161
	Bibliography	165

List of Figures

1-1	Electrode configurations and pseudopotential profiles for bulk and planar linear paul traps	17
1-2	Schematics of transition in two-qubit gates	23
1-3	Level structures of commonly used ions	25
1-4	Illustration of readout process and errors	28
1-5	Overview of some example scalable optics approaches	30
1-6	Schematic of approach envisioned here	32
2-1	Second detail schematic	38
2-2	Example mode profiles of dielectric waveguides	40
2-3	Bend-loss vs. bend radius for representative silicon nitride waveguides at $\lambda = 674$ nm and 1092 nm.	41
2-4	Pointing instability benefits with waveguide coupling	46
2-5	2 ion spacings for typical axial trap frequencies	48
2-6	Phase shifts per unit length due to self-phase modulation vs. power	50
2-7	Detector collection efficiency as a function of size	53
3-1	Schematic and overview of focusing grating design	59
3-2	Dimension considerations for integrated grating design	60
3-3	Coordinate definition for calculation of desired field in the grating plane	62
3-4	Example calculation of grating parameters from 2D cross-section simulations	63
3-5	Focusing grating design parameters	65
3-6	Cosine-profile waveguide mode field and maximally overlapping Gaussian profile	67
3-7	Mask layout for grating device, with inset showing discretization on a 5 nm rectangular grid.	68
3-8	SEM image of slightly underexposed grating feature, showing lines lifting off (obtained using a 0.28 μ s dwell time with a 2 nA beam current, equivalent to 0.11 μ s with 5 nA for comparison to the numbers given in the text)	70
3-9	Prism coupling film characterization	71
3-10	SEM image of input fiber coupler	72
3-11	Measured beam profiles for fabricated grating device	73

3-12	Experimental characterization of low-intensity sidelobes in radiated beam profile at $z = 50 \mu\text{m}$	75
3-13	Polarization purity characterization at focus	77
3-14	Expected improvement to sidelobe profile using larger grating	79
4-1	Schematic of integrated waveguide and ion trap device	85
4-2	Schematic and optical characterization of focusing grating devices	87
4-3	Micrograph of input couplers used for input coupling to free-space beams, with inset showing SEM image of curved grating lines near center.	89
4-4	Relevant level structure of $^{88}\text{Sr}^+$	91
4-5	Experimentally observed spectra of the S-D manifold	92
4-6	Coherent manipulations and ion addressing with the grating-generated beam	94
4-7	EMCCD images of 5 ions, with the center ion addressed and undergoing quantum jumps	95
4-8	Crosstalk quantification via intensity measurements and via ion excitation probability using the grating-generated beam	96
4-9	Waveguide loss measurements before and after metal deposition	99
4-10	SEM image of fabricated waveguide cross-section under metal electrodes, on quartz.	100
5-1	“Back-end” stackup of the 90-nm CMOS process used for the foundry trap	104
5-2	CMOS trap chip micrograph and perspective	105
5-3	Images of mask layout for CMOS trap	107
5-4	Micrograph of a diced and wirebonded CMOS trap, with inset showing the image of a pair of ions trapped.	108
5-5	Leakage current in DC HV measurement on CMOS trap chip	109
5-6	Heating reate measurement in CMOS foundry trap	110
5-7	AFM measurement of CMOS trap surface roughness	111
5-8	APD trap chip micrograph and mask layout of STI GR device	112
5-9	Layers involved in main APD designs attempted	114
5-10	Dark I-V Characteristics of main APD designs attempted	115
5-11	Dark and illuminated I-V characteristics of STI GR devices	116
5-12	Schematic of passive quenching circuit used	118
5-13	Pulse traces observed in Si APD with current-mode quenching	119
5-14	Dark count rates at room temperature for STI GR devices of various diameters	120
5-15	Dark count rates with $10 \mu\text{m}$ -diameter STI GR APD at various temperatures	120
5-16	DC photocurrent measurement as a function of temperature at 674 nm	121
5-17	Observed afterpulse bunches and dependence on quench resistance	122
5-18	Dark counts and response to low light 405 nm illumination at 90 K	124

6-1	Effective indices of TE modes for thin LiNbO ₃ slabs at $\lambda = 674$ nm and 1550 nm	129
6-2	Simulated optical mode profile at $\lambda = 674$ nm for a SiN/LiNbO ₃ hybrid waveguide, together with DC field profile from adjacent contacts on x -cut LiNbO ₃	130
6-3	Simulated bend losses for hybrid waveguides at $\lambda = 674$ nm	131
6-4	Electron range plots for Si and LiNbO ₃ substrates	134
6-5	SEMs of written features showing artifacts from proximity effect	135
6-6	Mode profiles of waveguides as fabricated	136
6-7	Optical micrograph of fabricated devices on LiNbO ₃	137
6-8	Ring resonator transmission spectra after anneals	140
6-9	Measured spectra of ring resonators at IR wavelengths, with theoretical fits	141
6-10	Transmission spectra of a high- Q ring as at various applied biases	142
6-11	Waveguide loss measurement at 674 nm	143
6-12	Alignment mark patterned in SiN layer for e-beam registration	144
6-13	674 nm Mach-Zehnder modulator in SiN-on-LiNbO ₃	146
6-14	LiNbO ₃ -based ring resonator tuning	147
7-1	Losses in single-mode Si ₃ N ₄ waveguides at various wavelengths	150
7-2	Schematic of multi-layer trap device (only ground plane shown for simplicity), together with waveguides on a silicon substrate.	152
A-1	Simulated and measured emission profiles in waveguide plane, showing mode distortion in grating region	160
B-1	Illustration of write orders employed for focusing grating devices of Ch. 3	162
B-2	Comparison of sidelobe profiles obtained with different e-beam lithography write orders	163

Chapter 1

Introduction: quantum manipulations of individual ions

The ability to precisely engineer the quantum states of many-particle systems would offer unprecedented abilities to study the behavior of large entangled and quantum coherent systems and the space between quantum and classical, simulate other quantum systems in ways that are intractable with classical computation, and solve certain problems with quantum computers for which no efficient classical algorithms are known [NC10]. The experimental realization of such systems gained much traction in the mid 1990s and around the time of Shor’s discovery of the factoring algorithm [Sho97] and has been an ongoing effort with increasing momentum for roughly 15-20 years now. Individual ions confined in electromagnetic potentials and controlled with laser light were among the first studied physical systems for implementation [MMK⁺95], and a number of groups have made much progress honing the basic operations required for computation, reducing error rates to levels workable with error correcting codes [BHL⁺16, GTL⁺16], and in experimental contexts which, with some imagination, can be envisioned to scale to larger systems [KMW02, MK13]. Many authors have indicated that no fundamental obstacles to building a large system for QIP are known – and the breadth of technical advances in the basic operations makes it look increasingly as though the practical situation may be tenable as well. But one of the major outstanding challenges, which has perhaps received less attention so far in communities focused on any physical qubit system, is reflected in the question of how exactly to implement the significant classical control apparatus, both electrical and optical, required to interact with the multi-qubit quantum system.

In this thesis we present approaches to creating the optical and electronic systems required for quantum state control of trapped ions, as well as readout via fluorescence detection; such approaches could be implemented in CMOS foundry processes, which would allow leveraging of a significant established infrastructure for integrating complex electrical/optical systems. The approach here relies on new optical structures for implementing known physical processes in the ions in a scalable fashion, and aims to combine the scalable nature of planar fabricated waveguide devices and electronics with the uniformity and isolation from environment, essential to precise quantum manipulation, of individual atoms.

Before proceeding to discuss the approach taken in this work and the relevant devices, this chapter aims to review very generally the physics of the trapped ion system to facilitate discussion of the role of optics in constructing larger systems. The aim is primarily to establish the role of optics and the general state of the art in trapped-ion systems, and to point to the relevant literature, in a way that may be useful to readers with minimal background in trapped ions. We first give an overview of the physics of trapped-ion systems, starting with the basic ideas of RF Paul traps, and then give an overview of the theory of laser-ion interactions as used for quantum gates with trapped ions. The different ion species used in various experiments are discussed, as well as the method for quantum state readout and laser cooling. Next the techniques for optics employed in present ion experiments are described, together with prior work on optics of some flavor integrated with planar traps. Our approach is briefly summarized (Chapter 2 discusses it in far more detail), and finally we attempt to put into context the scope of the challenges to scaling trapped-ion QIP systems that may be addressed by this work.

1.1 Physics of trapped ion qubits

Quantum information processing with trapped ions relies on a variety of atomic techniques built up over decades, and comprising very generally the stable confinement of individual atoms in vacuum, and the manipulation and measurement of their electronic states and vibrations. This section aims to describe the relevant techniques for ion qubits, to establish a context for the devices developed in this work.

1.1.1 Radio-frequency electromagnetic Paul traps for ions

Unlike traps for neutral atoms and particles that rely on electric or magnetic dipole forces, strong forces are exerted on ions via the Coulomb interaction, which offers a strong confinement mechanism and indeed presents a challenge to confinement via optical dipole forces owing to inevitable stray fields in practice [SEHS10]. A stable potential extremum for a charged particle in vacuum and a static electric field is prohibited by Gauss' law ($\nabla \cdot E = 0$) and Earnshaw's theorem in particular, but stable orbits exist in both static electric/magnetic fields (Penning traps), or in quadrupole electric field configurations oscillating at radio frequencies. The later approach forms the basis for Paul traps, as are most commonly used in quantum information experiments; such traps offer very deep potentials (trap depths on the order of 100 K), and often lifetimes of days or longer, and with very small field-induced level shifts on the trapped species. All of these features are beneficial both for demonstrated ion trap atomic clocks and early quantum processors. The most salient features of ion traps for quantum information experiments are very briefly reviewed here.

The dynamics of charged particles moving in oscillating quadrupole fields are described by Matthieu equations [Gho95], which can be solved exactly, and whose solutions along each dimension include "secular" motion at a frequency governed by the steepness of the potential and (typically) small-amplitude "micromotion" at the




Figure 1-1: (a) Schematic electrode configuration of RF and DC control electrodes for a bulk linear Paul trap, and RF pseudopotential plots for a bulk linear Paul trap (b) and a planar electrode Paul trap (c); figures from [CBB+05].

frequency of the oscillating potential. When the ion mass is large enough for a given spatial potential variation and temporal oscillation frequencies, its inertia is great enough that the high-frequency micromotion can be neglected to a good approximation, and in this limit the dynamics are well described by the “pseudopotential” [CBB+05]:

$$U(\vec{r}) = \frac{q^2}{2m\omega_{\text{rf}}^2} \langle E^2(\vec{r}) \rangle, \quad (1.1)$$

where q and m are the particle’s charge and mass, respectively, and ω_{rf} denotes the angular RF oscillation frequency. For a quadrupole field arrangement this potential is harmonic; the ion’s low frequency “secular” motion in this harmonic pseudopotential governs much of the dynamics.

In a typical “linear” trap, a quadrupole field is generated in two dimensions x and y , and ions are weakly bound along z only by static potential produced by DC voltages on control segments, as illustrated schematically in Fig. 1-1(a). Typical RF frequencies are of order a few 10s of MHz, and secular oscillation frequencies in the radial directions (x and y in these figures) determined by the RF voltage applied are typically a few MHz. Along the direction of the trap axis (z), the oscillation frequency is set by the DC control voltages applied. Confinement is usually weaker in this direction to allow multiple ions to arrange into a “Coulomb crystal”, where the spacing between ions is set by the balance between the confining potential and their mutual Coulomb repulsion. N ions in such a crystal then oscillate according to $3N$ normal modes, each with its own phonon frequency; importantly, the center-of-mass (COM) mode of motion is, in a purely harmonic potential, independent of N . This frequency can be widely tuned in experiments based on the DC control voltages, and is usually in the range of a few $2\pi \times 100$ kHz to 1 MHz. For many of the interactions with a single motional mode, such a selected mode of the ion or ion crystal is referred to as the “trap frequency” ω_t .

The ion spacing in a linear crystal depends of course on the ion mass, and confining potential, and is nonuniform for $N \geq 4$; for two ions, an analytic solution exists for the spacing between 2 ions $s_2 = 2^{1/3}s$ and between 3 ions $s_3 = (5/4)^{1/3}s$, where the

characteristic spacing s is [WMI+97]:

$$s = \left(\frac{q^2}{4\pi\epsilon_0 m\omega_z^2} \right)^{1/3}; \quad (1.2)$$

numerical values corresponding to this spacing are plotted later on in Fig. 2-5 for a few ion species.

The bulk linear traps have been used for many of the scientific experiments thus far, while planar electrode trap structures, with the potentials for multiple electrodes in complex geometries to manipulate large populations of ions, are recognized as promising candidates for large-scale systems. The reduced optical access, as well as difficulties posed by the proximity to surfaces, are some of the major additional challenges posed by such structures for current experiments though, to be discussed further below.

1.1.2 Carrier electronic transitions, single qubit operations

In the following we review the quantum description of the interaction of a general electronic two-level system with a classical radiation field. “Carrier” electronic transitions refers to transitions that leave the motional state of the ion unchanged, and which capture the physics of single qubit operations; the next section will include coupling to a motional mode of the trap in “sideband” transitions, essential for multi-qubit operations.

A two-level system coupled to a radiation field is described by the Hamiltonian

$$\hat{H} = \hat{H}_0 + \hat{H}_I, \quad (1.3)$$

where \hat{H}_0 gives the evolution of the system in the absence of the radiation, and \hat{H}_I represents the atom-field interaction. For excited and ground states $|e\rangle$ and $|g\rangle$, separated in energy by $\hbar\omega_a$, we can write

$$\hat{H}_0 = \hbar\omega_a |e\rangle\langle e|, \quad (1.4)$$

where the labels are chosen to correspond to qubit notation and the energy of $|g\rangle$ is 0. The interaction Hamiltonian represents coupling between the states at the Rabi frequency Ω ; the expression for this frequency as a function of the field(s) strength at the ion depend on the nature of the transition and the manner of addressing, and for now we suppose a dipole interaction, such that

$$\hat{H}_I = -\hat{d} \cdot \hat{\epsilon}_r E(t), \quad (1.5)$$

with the dipole operator $\hat{d} = e\hat{r}$ (whose matrix elements are defined as $\langle a|\hat{d}|b\rangle = \mu_{a,b}\hat{\epsilon}_d$), electric field polarization unit vector $\hat{\epsilon}_r$, and oscillating electric field $E(t) = \mathcal{E}e^{-i\omega_r t} + \mathcal{E}^*e^{i\omega_r t}$. Writing the dipole operator in terms of its matrix elements and assuming the field and atomic dipole polarization are aligned, and letting $\mu = \mu_{g,e} =$

$\mu_{e,g}^*$, the interaction Hamiltonian is

$$\hat{H}_I = -(|g\rangle\langle e|\mu + |e\rangle\langle g|\mu^*) \cdot (\mathcal{E}e^{-i\omega_r t} + \mathcal{E}^*e^{i\omega_r t}). \quad (1.6)$$

With the wavefunction $|\psi\rangle = c_0|g\rangle + c_1|e\rangle$ and the Schrödinger equation $\hat{H}|\psi\rangle = i\hbar|\dot{\psi}\rangle$, transforming the wavefunction so as to rotate at the radiation frequency so that the atomic wavefunction is described by $\tilde{c}_0 = c_0$ and $\tilde{c}_1 = c_1e^{i\omega_r t}$ (the energy spacing in the rotating frame is now given by $-\hbar\delta$ with the detuning $\delta = \omega_r - \omega_a$), and neglecting terms that oscillate as $e^{2i\omega_r t}$, results in the following two equations for the wavefunction amplitudes:

$$i\hbar\dot{\tilde{c}}_1 = -\hbar\delta\tilde{c}_1 - \mathcal{E}\mu c_0(t) \quad (1.7)$$

$$i\hbar\dot{\tilde{c}}_0 = -\mathcal{E}^*\mu^*\tilde{c}_1. \quad (1.8)$$

Defining the Rabi frequency $\Omega = \mathcal{E}\mu/\hbar$, and considering the resonant case $\delta = 0$, we have simply

$$\dot{\tilde{c}}_1 = i\Omega\tilde{c}_0, \quad (1.9)$$

$$\dot{\tilde{c}}_0 = i\Omega^*\tilde{c}_1, \quad (1.10)$$

and hence also

$$\ddot{\tilde{c}}_1 = -|\Omega|^2\tilde{c}_1 \quad (1.11)$$

$$\ddot{\tilde{c}}_0 = -|\Omega|^2\tilde{c}_0. \quad (1.12)$$

The neglecting of the fast varying term is commonly known as the “rotating wave approximation” (RWA), valid for radiation fields near resonance $\delta \ll \omega_a, \omega_r$, and weak so that $\Omega \ll \omega_r$; under these conditions, if the state initially is $|0\rangle$ and we let $\Omega = |\Omega|e^{i\phi}$, we have

$$\tilde{c}_0 = \cos(|\Omega|t) \quad (1.13)$$

$$\tilde{c}_1 = ie^{i\phi}\sin(|\Omega|t). \quad (1.14)$$

This particular case illustrates that the magnitude of the applied field sets the rate of transfer between the states, and that the phase ϕ of the applied field is imprinted onto the excited state amplitude arising from the interaction.

In this derivation for explicitness the rotating frame transformation and rotating wave approximation were made after expanding the Hamiltonian into equations for the amplitudes – both could have been applied to the Hamiltonian to begin with, and the Hamiltonian describing the system in this frame and approximation would be

$$\hat{H} = -\hbar\delta|e\rangle\langle e| - i\hbar[\Omega|e\rangle\langle g| + \Omega^*|g\rangle\langle e|], \quad (1.15)$$

which will be the starting point for consideration of motional mode coupling in the next section.

Clearly a resonant field can induce oscillations of probability amplitudes between the two electronic states; to see resonant fields of controlled phase and amplitude can map any initial state to any other state and thus carry out arbitrary single qubit operations, a more general picture based on the Bloch sphere representation is helpful. It is possible to express the movement of the state vector on the Bloch sphere as a rotation around the unit vector [HR13]

$$\hat{n} = \frac{-\delta\hat{z} - |\Omega|\sin\phi\hat{x} + |\Omega|\cos\phi\hat{y}}{\sqrt{\delta^2 + |\Omega|^2}}, \quad (1.16)$$

with a degree of rotation controlled by the time of interaction (pulse length). The surface of the Bloch sphere being a two-dimensional one, rotations around two axes are sufficient to map between any pair of states; hence, with $\delta = 0$, rotations about both x and y axes are possible, and hence arbitrary single qubit operations are possible with amplitude and phase control of resonant light.

1.1.3 Coupling to a motional mode; the Lamb-Dicke regime

In trapped ions, shared motional modes of multiple ions in a single harmonic potential are crucial to the two-qubit operations to be discussed below, and which together with the single qubit operations effected by resonant light form a complete basis for state manipulation [NC10]. This section describes the physics of the coupling between the “internal” ion state (i.e. the valence electron level) and its external state, described by the quantum state of the harmonic oscillators corresponding to its motional modes. Considering near-resonant excitation of just a single motional mode illustrates the essential physics.

Including the contribution from this motional mode, the Hamiltonian with radiation turned off is now

$$\hat{H}_0 = \hbar\omega_a|e\rangle\langle e| + \hbar\omega_t a^\dagger a, \quad (1.17)$$

where a^\dagger and a are the creation and annihilation operators the motional mode harmonic oscillator of trap frequency ω_t , and the ground and excited states of the qubit are now labeled $|g\rangle$ and $|e\rangle$ to avoid confusion with the motional eigenstates $|0\rangle, |1\rangle, |2\rangle$, etc. The energy levels of this “bare” hamiltonian form a ladder of states $|g, n\rangle$ and $|e, n\rangle$, with energies equal simply to the sum of the internal and motional energies.

A very similar Hamiltonian as in Eq. 1.15 is used to describe coupling between these eigenstates, with the simple addition of phase seen by the ion depending on the ion position. That is, if the ion oscillation is along \hat{z} and the electric field of a classical beam propagating with angle θ to \hat{z} is $E(z, t) = \mathcal{E}e^{ikz\cos\theta - i\omega_r t}$, the interaction Hamiltonian in the RWA (which corresponds to neglecting terms leading to oscillations at $2\omega_r$) is

$$\hat{H}_I = -\hbar\Omega|e\rangle\langle g|e^{ik\hat{z}\cos\theta - i\omega_r t} - \hbar\Omega^*|g\rangle\langle e|e^{-ik\hat{z}\cos\theta + i\omega_r t}, \quad (1.18)$$

where coupling between the classical electric field and the quantized ion motion is

described via the atomic position operator $\hat{z} = \sqrt{\frac{\hbar}{2m\omega_t}} (\hat{a} + \hat{a}^\dagger)$, which has taken the place of z in the otherwise classical field. Here a and a^\dagger are the annihilation and creation operators of the relevant motional harmonic oscillator mode.

Substituting the creation and annihilation operators in for \hat{z} , the exponentials are proportional to $e^{\pm i\eta(\hat{a} + \hat{a}^\dagger)}$, where

$$\eta = \sqrt{\frac{\hbar k^2}{2m\omega_t}} \cos \theta \equiv \eta_0 \cos \theta \quad (1.19)$$

is known as the Lamb-Dicke parameter. η_0 can be expressed equivalently as

$$\eta_0 = \sqrt{\frac{E_r}{\hbar\omega_t}} = \frac{2\pi a_0}{\lambda}, \quad (1.20)$$

where $E_r = \frac{\hbar^2 k^2}{2m}$ is the recoil energy imparted to the ion of mass m after scattering a photon of momentum $\hbar k$, and the ground state wavefunction extent is $a_0 = \sqrt{\frac{\hbar}{2m\omega_t}}$, typically on the order of 10 nm. These two descriptions of the maximum (with beam propagating along the trap axis) Lamb-Dicke parameter η_0 show it describes both the extent of the ground state wavefunction relative to the wavelength and the ratio of the recoil energy to the motional mode quanta; ion traps are typically quite tight and this parameter is typically around 0.01 – 0.3, depending on the particular setup. When η_0 is much less than 1, the energy delivered to the ion by a scattering event is small compared to the quantum of energy in the motional mode, and so the probability for a photon to excite a phonon is low in proportion to η (and tunable via the angle θ , which reflects the photon momentum along the direction of ion oscillation) – and the probability of multiple phonons even more so (by successive powers of η_0).

The details of the derivation are presented elsewhere (e.g. in Ch. 8 of [HR06]), but here we want to note that, in an interaction picture with respect to the bare atom/motional mode Hamiltonian, by expanding the exponential as a power series of η , we obtain a series of terms, indexed by q (any integer), each term proportional to $(i\eta)^q e^{i(\omega_a - \omega_r + q\omega_t)t}$. When the laser frequency is resonant ($\omega_r = \omega_a + q\omega_t$) such that one of these time dependencies vanishes, that single term is time-independent and the others oscillate at at least ω_t . When the dynamics are slow compared to the trap frequency ($\Omega \ll \omega_t$), a second RWA allows neglecting these ω_t time-scale oscillations. And in this case, we can consider separately the dominant term in each of the three situations, corresponding to $q = 0$, the carrier transition (when the laser is tuned to ω_a , and $q = \pm 1$, the “red” and “blue” sideband transitions, in which the laser is tuned to a sideband of the carrier transition so as to excite or de-excite one motional quantum).

The time-independent terms corresponding to these three resonance conditions,

to first order in η (valid for $\eta \ll 1$), are:

$$H_c \approx -\hbar\Omega(|e\rangle\langle g| + |g\rangle\langle e|) \quad (1.21)$$

$$H_r \approx -\hbar\Omega\eta(a|e\rangle\langle g| - a^\dagger|g\rangle\langle e|) \quad (1.22)$$

$$H_b \approx -\hbar\Omega\eta(a^\dagger|e\rangle\langle g| - a|g\rangle\langle e|) \quad (1.23)$$

where H_c , H_r , and H_b represent the dominant terms when the laser frequency addresses transitions on the carrier, and red and blue sidebands.

When any of these terms operates on states $|g, n\rangle$, representing the atom in its internal ground state and the motional mode in a Fock state with n motional excitations, they drive transitions that excite the internal state to $|e\rangle$ and preserve n (carrier transition), reduce n by 1 (red sideband) or increase n by 1 (blue sideband) at effective Rabi frequencies $\Omega_{n,n}$, $\Omega_{n,n-1}$, and $\Omega_{n,n+1}$ that evaluate to (again to lowest order in η):

$$\Omega_{n,n} \approx (1 - \eta^2 n)\Omega \quad (1.24)$$

$$\Omega_{n,n-1} \approx \eta\sqrt{n}\Omega \quad (1.25)$$

$$\Omega_{n,n+1} \approx \eta\sqrt{n+1}\Omega, \quad (1.26)$$

indicating the dependence of the transition strengths both on η and the motional mode occupancies.

1.1.4 Two-qubit operations

The coupling to collective motional modes described in the previous section offers a powerful mechanism for coherent interaction and is the basis for two-qubit gates in trapped ions. The first proposed such gate is the Cirac-Zoller gate [CZ95], which in fact sparked work on experimental quantum computation generally as the first practical two-qubit gate proposal. The process is illustrated in the figure reproduced in the top panel of Fig. 1-2. Here, two ions are cooled to the ground state, and both are generally in a superposition of the $|0, g\rangle$ and $|0, e\rangle$ states, where the 0 indicates the number of motional quanta. A red sideband π -pulse (i.e. a pulse with time and amplitude chosen so that $\int \Omega_{n,n-1}(t)tdt = \pi/2$, and the population is wholly transferred on the resonant transition) applied to ion 1 transfers its population in $|0, e\rangle_1$ to $|1, g\rangle_1$; this motional quantum excited is of a mode shared by the two ions, and this now allows an operation on ion 2 that can be conditioned on the motional mode being excited, namely in another red sideband coupling from the ground state to some auxiliary level. A 2π pulse applied to this transition does not alter population, but induces a π phase shift on the $|1, g\rangle_2$ state. On the other hand the $|0, g\rangle_2$ state would have been coupled to no red sideband transition, and would not have acquired this phase. Finally the ion 1 is returned to its original state by another π -pulse on




Figure 1-2: (Top) electronic/motional transitions involved in a two-qubit Cirac-Zoller gate, and (bottom) in a Molmer-Sorensen gate. Images from [HRB08].

the red sideband, completing the conditional phase shift summarized in the mapping:

$$|g\rangle_1|g\rangle_2 \rightarrow |g\rangle_1|g\rangle_2 \quad (1.27)$$

$$|e\rangle_1|g\rangle_2 \rightarrow -|e\rangle_1|g\rangle_2 \quad (1.28)$$

$$|g\rangle_1|e\rangle_2 \rightarrow |g\rangle_1|e\rangle_2 \quad (1.29)$$

$$|e\rangle_1|e\rangle_2 \rightarrow |e\rangle_1|e\rangle_2 \quad (1.30)$$

This technique relies on the ions being precisely in the ground state of motion – any excitation would allow the red sidebands to be addressed on the second ion to some extent regardless of ion 1’s state. This places considerable demands on ions being in the absolute ground state of motion prior to beginning the gate, and is a challenge in achieving high fidelities with this gate technique [SKHR+03].

Another kind of interaction, used in what are referred to as Molmer-Sorensen type gates, was later proposed which is less sensitive to ions being in the ground state [SM99, SM00]. Here also the interaction is due to a shared motional mode, but both ions are illuminated by two frequencies simultaneously, at frequencies $\omega_a + (\omega_t + \delta)$ and $\omega_a - (\omega_t + \delta)$, where ω_t is the motional mode frequency and δ a detuning. The two frequencies sum to $2\omega_a$, corresponding to the energy spacing between $|g\rangle_1|g\rangle_2$ and $|e\rangle_1|e\rangle_2$, and were shown to be able to generate the transformation $|g\rangle_1|g\rangle_2 \rightarrow \frac{1}{\sqrt{2}}(|g\rangle_1|g\rangle_2 - i|e\rangle_1|e\rangle_2)$ in a time

$$\tau = \frac{\pi}{\eta\Omega} \sqrt{K}, \quad (1.31)$$

for an appropriate detuning such that $\frac{\eta\Omega}{\omega_t - \delta} = \frac{1}{2\sqrt{K}}$, with K any integer and Ω the single-qubit Rabi frequency [SM00]. In addition to relaxing the need for ground-state cooling, this gate does not require individual addressing of each of the two so long as those qubits involved in an operation can be addressed together. The highest-fidelity multi-qubit gates have employed this interaction, or the closely-related geometric phase gate [SKK+00, LDM+03, BKRB08b, Ba114], which have various relative technical advantages and disadvantages [LBD+05]. Of course, though implementation of this gate on two ions alone does not require individual addressing of either ion, to implement this gate between any given two ions of a larger system would of course require some form of addressing, whether purely optical, or together with appropriate motion in and out of addressed sites.

1.1.5 Atomic structure of commonly used ions; hyperfine and quadrupole qubits

Having summarized the quantum coherent operations fundamental to state manipulation, we now consider the particular features of real ions and the possibilities for implementing such interactions. Though our experiments will focus on the $^{88}\text{Sr}^+$ ion, many of the concepts in this thesis are more generally applicable to different ion species and qubit types, e.g. the $^{40}\text{Ca}^+$ or $^{43}\text{Ca}^+$ ions used by the Innsbruck and Oxford groups and the $^9\text{Be}^+$ ion used by the Colorado group. We briefly discuss the two major implementations of qubits – those based on narrow linewidth quadrupole transitions and those based on hyperfine ground state levels, each having different advantages and disadvantages.

Fig 1-3 shows the valence electron level structures for the major experimentally investigated ions. Hyperfine qubits are encoded between particular Zeeman levels in the s-orbital manifold of the $^9\text{Be}^+$ and $^{171}\text{Yb}^+$ ions by the NIST groups and Maryland/-Duke groups, respectively, with transition frequencies given by the hyperfine structure splitting of 1.25 GHz and 12.6 GHz. Optical qubits, on long-lived dipole-forbidden quadrupole transitions, are used in the $^{88}\text{Sr}^+$ and $^{40}\text{Ca}^+$ by the MIT/Lincoln Labs groups and the Innsbruck group; both of these species have 0 nuclear spin and thus no hyperfine structure. The $^{43}\text{Ca}^+$ qubit offers an accessible optical qubit in addition to having a nuclear spin, and has a complex level structure that would allow the combination of both qubits, for manipulations and memory separately; a major difficulty it poses is in pumping the ground-state electron into just a single state in the complex ground state manifold. In all cases, S-P transitions at wavelengths spanning 310 - 420 nm are used for Doppler cooling and readout.

Hyperfine and optical qubits have different advantages. Due to the $\propto \omega^3$ scaling of the photon spontaneous emission rate, the GHz transitions in hyperfine qubits have very long intrinsic lifetimes as compared to the spontaneous emission-limited optical qubits (lifetimes of 390 ms and 1.2 s for Sr and Ca), and are instead limited in coherence primarily by magnetic field noise (aside from infidelities in gate operations). The different mechanisms for addressing and manipulation however pose different compromises. For the hyperfine qubits, coupling between the qubit states,

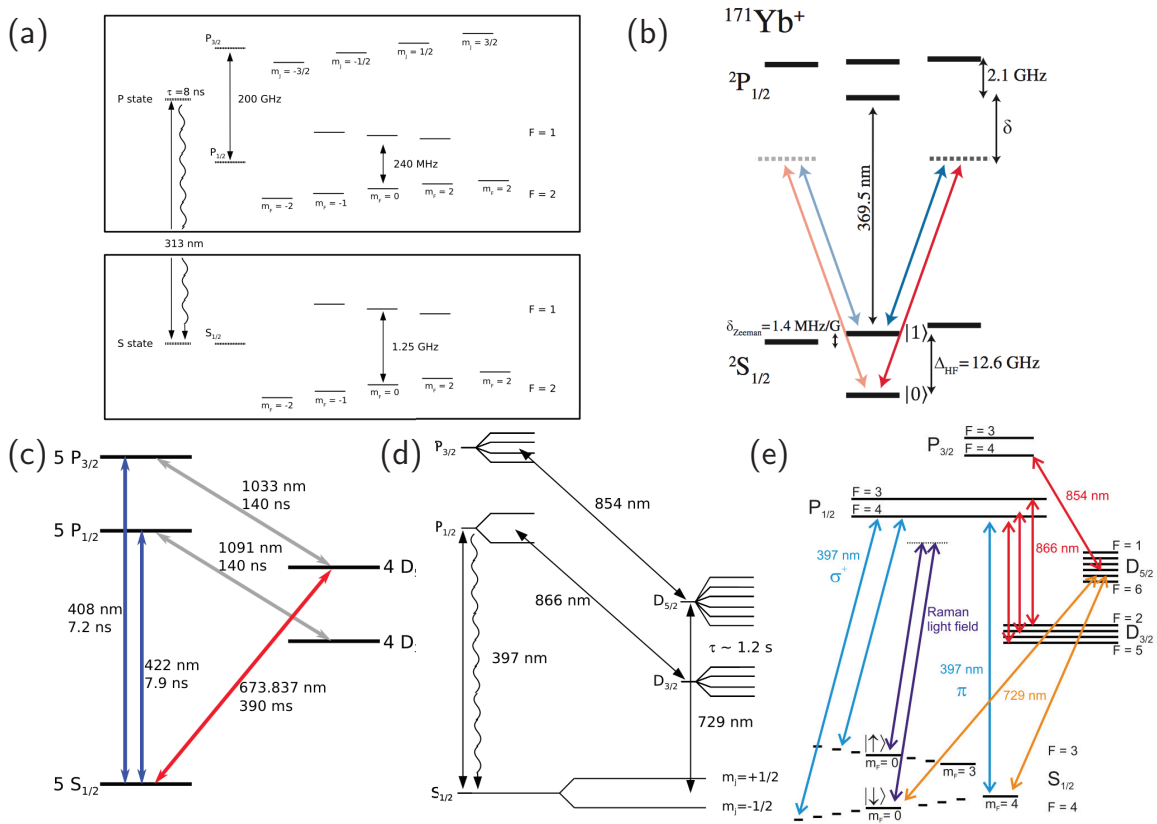


Figure 1-3: Example level structures of experimentally investigated ions: (a) ${}^9\text{Be}^+$ used by the NIST group, figure from [HRB08]; (b) ${}^{171}\text{Yb}^+$ used by the Maryland and Duke groups [MBB⁺13]; (c) ${}^{88}\text{Sr}^+$ by the MIT and Lincoln Labs groups [Wan12]; (d) ${}^{40}\text{Ca}^+$ [HRB08] and (e) ${}^{43}\text{Ca}^+$ [BKRB08a], used and considered by the Innsbruck, Oxford, and Berkeley groups.

as well as to motional modes, are carried out with stimulated Raman transitions by pairs of laser beams, each detuned from the resonant transition and with difference frequency equal to the hyperfine splitting. Experimental complexity is not obviously much greater on one side or the other, because very narrow (Hz-level) linewidth is required for the optical qubits to preserve the spontaneous-emission limited coherence time; whereas although the Raman approach requires two different frequencies, the frequency difference is the most important parameter in the technical noise, which can be low without great difficulty when both frequencies are derived from a single laser for the ${}^9\text{Be}^+$ ion; for the ${}^{171}\text{Yb}^+$ ion the larger hyperfine-structure splitting makes this somewhat difficult (owing to the finite bandwidths of AOMs, although electro-optic modulators can address this [OYM⁺07]), though high fidelity gates have been reported with this qubit as well [BKG⁺16]. Optical frequency combs may allow convenient addressing of this difference frequency [CMQ⁺10, HMM⁺10] and with the possibility of faster gate times than those based on modulated CW lasers, but high fidelity operations have yet to be demonstrated with such an approach.

Spontaneous emission from virtually excited states in the stimulated Raman process, however, can be a limitation to the fidelity of such transitions, although the probability of excitation of such states can be decreased by increasing detuning of the lasers (and also the intensity, to maintain a given operation speed) [WMI⁺97, OIB⁺07]. The spontaneous emission probability of the D-levels in the optical qubits over a gate time can more easily be low (10^{-4} - 10^{-5}) for reasonable gate times and lower intensities, making them appealing from this perspective for both optically-addressed single and two-qubit gates.

The long coherence time of the hyperfine states, especially those in “clock” states, makes them attractive for memory [LOJ⁺05]. And they offer the possibility of very high-fidelity microwave-driven single qubit operations [HAB⁺14]. However, the long wavelength of microwaves implies that, due to the Lambe-Dicke parameter-dependence ($1/\lambda$) of the motional-mode coupling strength, motional mode coupling to free-space microwaves is very slow for reasonable powers. Although near-field subwavelength electrodes within a planar trap can offer such coupling [OLA⁺08, OWC⁺11], such an approach is likely to be highly dissipative, and still struggles to achieve comparable two-qubit gate times as optical approaches; furthermore, whether for free-space single-qubit operations or for near-field-induced motional mode couplings, limiting the spatial extent of the microwave to address a single ion is significantly harder than with optical beams. Potential approaches to using magnetic field gradients to differentially shift Zeeman splittings for different ions to allow individual addressing pose their own challenges [WOC⁺13].

Clearly a rather broad set of both technical experimental and fundamental characteristics comes to bear on which qubit type and addressing method is ideal for what purpose, and the various trade-offs motivate consideration of more complex schemes in which, for example, an optical qubit is used for entangling operations and mapped to a hyperfine one for memory [BKR⁺08a], or in which hyperfine states are used exclusively and addressed optically via Raman transitions for two-qubit gates and via microwaves for single-qubit operations. Here we simply presented an overview of the various approaches, and to establish the general wavelength ranges involved for the

most important operations in the $^{88}\text{Sr}^+$ used in this work but also ion species under study elsewhere; more will be said in the next chapter about the particular role of optics, but it is clear that fluorescence is the only readout mechanism used, and optical addressing has significant technical advantages for individual addressing of all operations (although fundamental disadvantages compared to microwaves for single-qubit operations), and both fundamental and practical advantages for two-qubit operations.

1.1.6 Fluorescence readout of qubit state

When one qubit state can be coupled on an optical frequency cycling transition (one in which spontaneous decay is primarily to the initial state) to an excited state, readout is achieved simply by detecting scattering of light resonant with this transition; light is scattered only if the qubit was in one of its two states. As long as the expected number of photons collected per measurement time after all losses are considered is large compared to the dark count rate, such readout offers a conceptually simple route to high fidelity measurement [WMI⁺97]. In Fig. 1-4, for example, Poisson-distributed count probabilities for some measurement time are illustrated for expected value of dark counts $\langle N \rangle_{dk} = 4$, simulating a hypothetical distribution of counts in the “dark” state (owing either to detector dark counts or scattered light driving the transition), and the expectation number of counts for a bright ion $\langle N \rangle_b = 30$, simulating the ion fluorescing ($\langle N \rangle_b = \langle N_s \rangle \eta$, where $\langle N_s \rangle$ is the expectation number of photons scattered in the measurement time and η is the collection efficiency). If in a given measurement the number of counts is above the optimum threshold, the ion is judged to be bright, and vice versa; clearly a large separation between the dark and bright distributions will correspond to a minimal error probability, and also plotted is the probability of a wrong inference as a function of the expected “bright” counts, for a few values of expected “dark” counts; in the case of 0 dark counts, this is simply the probability that 0 photons are collected from a fluorescing ion, equal to $e^{-\langle N \rangle_b}$. We note that this is illustrative only – while the distribution of counts from the fluorescing ion is Poissonian to a good approximation, the dark counts may not be.

When using a qubit encoded on an optical transition, such a cycling transition for readout is provided by an S-P transition, for example in $^{88}\text{Sr}^+$ the $5S_{1/2} \rightarrow 5P_{1/2}$ transition at 422 nm; an ion will scatter light on this transition only if it is measured in to be in the $5S_{1/2}$ state. When a hyperfine qubit is used the situation is somewhat more complex as particular sublevels of the S and P-state manifolds have to be coupled such that when exciting from the bright qubit state, decay to the dark qubit state is minimized, via selection by polarization and frequency of a particular excited sublevel; however in practice due to off-resonant couplings and imperfect polarizations there is invariably some leakage, and as a result of the greater isolation between qubit states when an optical qubit is employed, these have achieved higher readout fidelities [MSW⁺08]. An exception is when population from a hyperfine qubit is transferred to a metastable D-state prior to cycling [HAB⁺14].

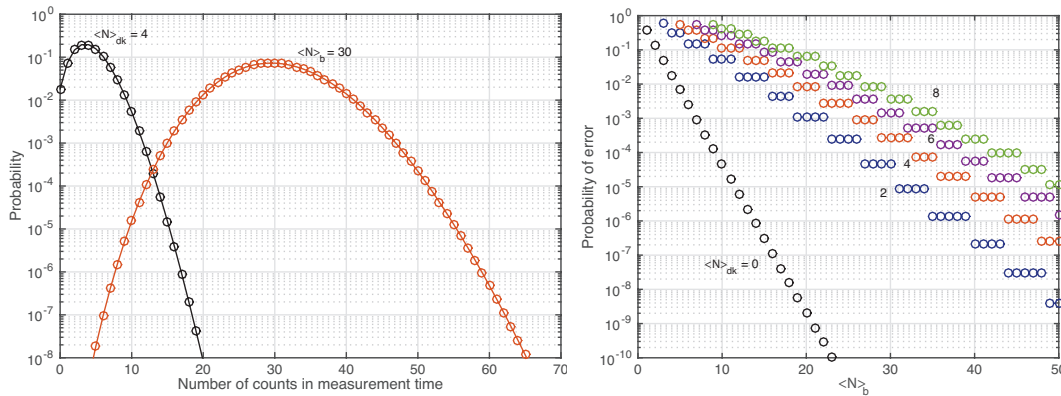


Figure 1-4: Illustration of readout process and errors. (left) Example histograms of Poisson-distributed count probabilities in measurement time, with $\langle N \rangle_{dk} = 4$ expected counts for the dark state and $\langle N \rangle_b = 30$ for the fluorescent state. (right) Readout error probabilities as a function of the expectation number of counts from the fluorescent state, for various expected dark count levels $\langle N \rangle_{dk} = 2, 4, 6, 8$.

1.1.7 Cooling of trapped ions

It is important that ions be cooled to near the ground state of motion; for the Cirac-Zoller gate this is essential to the principle of operation, but in general, this is required for high fidelity operations, so that for example the ion position is not varying with respect to the applied fields and hence that the relevant Rabi frequency and pulse time required is known precisely. Doppler cooling is performed on a cycling transition, often the same as used for readout, using a laser red-detuned from the transition, for cooling to on the order of $100 \mu\text{K}$; for ions, in the Lamb-Dicke regime as usual, resolved sideband cooling, using the red-sideband of the qubit transition, is then used to walk the ion down the ladder of internal/motional states and cool to near the ground state [WMI⁺97], with average motional quanta on the order of 0.05 achieved routinely.

1.2 Experimental implementation and approaches to scalable optics

Current experiments can be divided into those based on bulk traps such as the four rod design pictured in Fig. 1-1a,b:, and those based on planar ion trap structures. Common between them are the ion loading processes, which typically involve evaporation off a hot source and photoionization near the region of the trap. The principles of the experiments are the same, but the planar ion experiments reduce the available solid angle for optical addressing, and introduce added challenges of avoiding scatter off the chip surface, and typically also suffer from larger heating rates, and have hence so far not been used for experiments at the same scale as the bulk trap experiments.

1.2.1 Bulk trap experiments

Most experiments on quantum operations in single or few ions so far have been carried out in “bulk,” few mm-scale traps of the type illustrated in fig. 1-1(a). These are capable of large trap depths of typically order ~ 100 meV, favorable noise properties due to the large distance between ions and the metal surfaces, and convenient optical access from almost all directions. These experiments range from the original demonstrations of two-qubit gates, to the more recent experiments with multiple ions in a single linear ion crystal [DLF⁺16, MMS⁺16]. The Maryland and the Innsbruck group have pursued most of the multi-ion bulk trap experiments so far, achieving for example entanglement of 14 $^{40}\text{Ca}^+$ qubits in a Green-Horne-Zeilinger state by collectively addressing all qubits with the same fields [MSB⁺11], so far the largest entangled state created with any qubit, executing more general algorithms on 3-5 qubits using a focused laser beam in free space directed to different ion sites with an electro-optic deflector [SNM⁺13].

It is notable that only two groups in the world, the Innsbruck and Maryland groups, has at the time of this writing demonstrated high-fidelity operations on a >2 -ion system with individual addressing, as required for general operations [SNM⁺13, MNM⁺15, DLF⁺16]. The Innsbruck system scans a single tightly-focused beam serially across the ion chain using electro-optic deflectors, and to avoid the necessity of phase stability between this single-ion addressed beam and the other beams that globally address all ions, the ion addressing beam is detuned from resonance and serves only to create a stark-shift on the addressed ion.

While measurement of one or two ions is often done using a PMT, the need to spatially-resolve ion fluorescence in multi-ion experiments motivates imaging onto an EMCCD for such experiments [Hem14].

The highest-fidelity two-qubit gates carried out so far have been in bulk traps. Two-qubit phase gates, using optically addressed Raman transitions in $^{43}\text{Ca}^+$ or $^9\text{Be}^+$, have achieved maximum fidelities of 99.9% [BHL⁺16, GTL⁺16], as compared to the previously demonstrated 99.3% fidelity Molmer-Sorensen gates carried out with the optical qubit in $^{40}\text{Ca}^+$ [BKRB08b].

1.2.2 Planar electrode trap experiments

While bulk traps have allowed a number of interesting experiments of gradually increasing size and precision, architectures based on the complex electrodes fabricable in planar structures bring significant advantages in the long run for large arrays [CBB⁺05], and hence the approaches to scalable optics discussed below focus on integration with such traps. The conventional picture of light-ion interactions in such traps, and most current experiments, rely on laser beams propagating parallel to the surface of the trap (which is entirely coated with electrodes, to avoid any regions of floating potential except in the $\approx 5\mu\text{m}$ gaps between electrodes, that can destabilize the trap).

The ion-electrode distance is typically in the range of 30-300 μm , smaller than most bulk trap experiments. Like the bulk trap experiments, the traps can be operated at

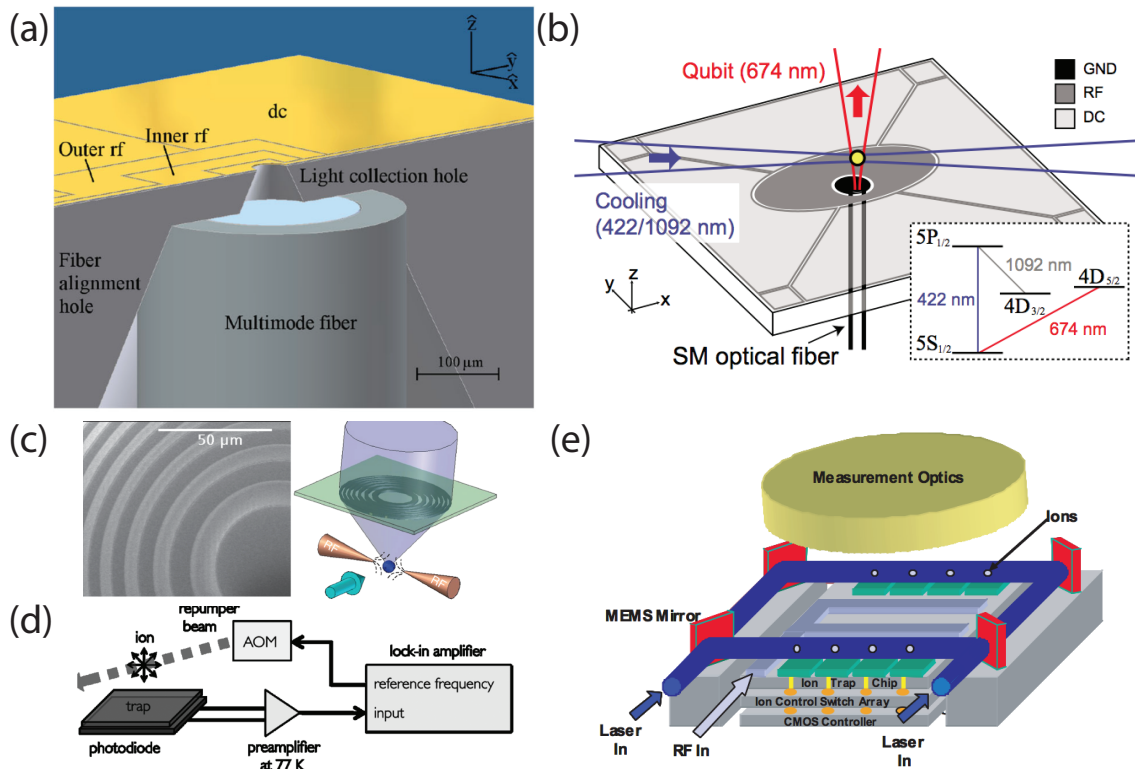


Figure 1-5: Various steps towards scalable optics integration with planar ion traps. (a) An integrated multimode fiber for fluorescence collection [VCA⁺10], (b) Single mode fiber for light delivery to ions [KHC11], (c) A microfabricated fresnel lens for imaging [SNJ⁺11], (d) A transparent trap with an integrated photodetector [EWA⁺13], and (e) A MEMS-optics based approach to integrated optical control [KK09, MK13]

room temperature, but the heating rates observed in traps with electrodes near the ions increase strongly with reduced distance, and is orders of magnitude higher than heating expected from Johnson noise in the electrodes. The mechanism for this so-called “anomalous heating” is not understood, but cryogenic cooling has been shown to allow significant reduction of the heating rate [LGA⁺08, CS14], to levels of around 5 motional quanta/s in traps with $d = 50 \mu\text{m}$ ion-electrode distance. The distance dependence may very roughly scale with $1/d^n$ with n between 2 and 4, but the scaling is very difficult to ascertain from the various experiments conducted so far [BKRB14].

It is worth noting that the highest fidelity single qubit gates performed in ions (or in any qubit) to date were carried out in a planar trap using microwave pulses to couple hyperfine “atomic clock” states of a $^{43}\text{Ca}^+$ ion [HAB⁺14], reaching average single-qubit gate fidelities of 99.9999%.

1.2.3 Approaches to scalable optics in planar traps; overview of work presented here

The need for integrated optics approaches to trapped ion quantum information processing has been recognized for some time [KK09], and various groups have been working on initial steps towards such a goal.

A handful of experiments have pointed to particular devices or elements that may form a part of a scalable approach in planar ion traps; examples include multi-mode fiber optic readout of fluorescence for relatively efficient and perhaps multiplexible readout [VCA⁺10], delivery of light to ions via single mode fibers [KHC11], fluorescence imaging through lithographically patterned fresnel lens [SNJ⁺11], and a large area photodetector integrated with a transparent trap for efficient collection [EWA⁺13], with some illustrative figures included in Fig. 1-5. While many of these experiments demonstrate performances superior to the traditional free-space optics approaches (fluorescence collection of 2%, e.g. in the multi-mode fiber collection experiment), they in general have relied on rather specialized fabrication techniques that do not clearly point to practical scalable implementations.

One of the most thorough existing proposals with an eye to large-scale implementations is based on MEMS optics; an early proposal considered beams propagating parallel to the chip surface and routed by switchable MEMS mirrors [KK09]. The architecture proposal is appealing in simply allowing for the same polarizations (both circular and linear) as are used in bulk experiments across a wide wavelength range while still manipulated with microfabricated optics; however, due to concerns about beam diffraction and scattering off the chip surface, the beams propagating parallel to the surface pose significant constraints on the degree of focusing and the chip size; for example, 30 μm beam waists limit the chip size to about 1 cm^2 , and tighter beams can come only at the expense of smaller overall chip size. MEMS mirrors are also limited to similar μs -scale switching times as free-space electro-optic deflectors, and the integration of MEMS mirrors with planar traps entails significant fabrication challenges as well; experimental demonstrations along these lines have so far used MEMS optics external to the vacuum chamber and imaged onto the ions through a sizable set of external optics [CMBK14, MK13], which is not as clearly scalable to multiple ions. Although the through-substrate trenches used in these studies allows μm -scale focuses of beams propagating normal to the trap surface, in contrast to the surface-parallel beams in the original proposal, the deep etch and the requirement for distant MEMS mirrors and imaging optics poses problems for stable alignment, high-yield fabrication and integration of electronics. Despite the clear challenges, the MEMS optics approach is one of the most plausible approaches to really scalable optics integration already under study.

This thesis presents an approach that instead utilizes waveguiding optics, fully integrated within an ion trap chip, to route and direct light to ions. Focusing grating couplers are designed to couple the light guided in 100 nm-scale high-index-contrast dielectric waveguides to beams propagating in vacuum to a focus co-located with ions trapped above the chip, as illustrated schematically in Fig. 1-6. Integrated electro-optic modulators together with passive waveguides and couplers are envisioned to




Figure 1-6: Overview of approach for control; waveguides patterned beneath the trap electrodes split and route light; grating couplers direct the guided light to points in free-space to address the ions. Modulators within the chip are required to encode control pulses to implement particular operations.

allow for parallel manipulation and delivery of light for the various tasks to ions at large numbers of sites. While the focus of this thesis is on the establishment of the basic building blocks for such an approach, our argument is that these basic pieces will be more practical to scale to larger systems, and will both greatly simplify experiments at the levels of 10s of qubits, and make it possible to scale the optical control and measurement systems to 1000s. A crucial point in this is the ability to fabricate devices of the kind imagined and studied here within full CMOS foundries [MEB⁺14, SGO⁺15], which presents a path to fabrication of large-scale systems with complex waveguide circuits and multi-layer trap structures designed together. The next chapter will discuss in some more detail the potential advantages and challenges of the approach presented and studied in this work.

1.3 Summary of scaling challenges and those addressed here

Many challenges face the practical implementation of large-scale trapped ion quantum information processing; this thesis aims to address a few of them, but the scope of the challenge is large enough that we necessarily focus on only a subset of the outstanding issues. But it may still be worthwhile to attempt to keep in view a more full set of major apparent challenges as of this writing, to help to define the scope of this work (which is also elaborated more in the next chapter).

1.3.1 Scalable control apparatus

The efforts towards integration of optical components into planar ion traps discussed above are part of an attempt to address the broader problem of implementing the devices and systems required for trapped ion control in a manner that can be practically applied to many ions [MK13]. Integrating such components at a large scale in a way that avoids an excessive experimental overhead, and furthermore does not add sources of noise or instability to the quantum system, is the challenge here; this challenge applies to the optics as well to the electrical systems for trap control and pulse shaping.

With a view to an eventual system, at early stages in the development of quantum processors when experiments are generally limited to a few qubits [BKM⁺14, MNM⁺15], it would be easy to underestimate the significance of this problem relative to the properties and limits of the more basic quantum operations. However, these problems may prove decisive in future years, and are likely to strongly affect what basic techniques for qubit implementations or gate operations prove useful in the long run. As an example, the planar superconducting qubits that have shown long coherence times (relative to such qubits historically) of 10-100 μs in recent years, e.g. [BKM⁺13] have been fabricated using shadow evaporation techniques standard to Josephson junction device engineering, and samples fabricated with direct etching have shown coherence times 1-2 orders of magnitude lower. The interfaces, defects, and loss channels introduced by additional dielectrics and metal layers are known to present various additional decoherence sources that reduce lifetimes from the high values obtained in simple few-qubit samples; but given the necessity of more complex structures simply for wiring up a multi-qubit system, how well these issues can be addressed will be crucial for the practical applicability of superconducting qubit techniques to larger systems [BPW⁺15, BMR⁺16]. Ion qubits face the same general problem, though with very different technical features, and the practical implementation of this control apparatus is the main problem addressed in this work.

1.3.2 Heating rates

As discussed above ions are initialized in motional states near the ground state of the mode used for two-qubit interactions; heating of this motional mode can contribute to errors in two-qubit operations, and the heating rate is a crucial metric for planar traps. Given the strong scaling with ion distance from surfaces [BKRB14], this has limited distances to trap electrodes so far to roughly $>30 \mu\text{m}$. Much effort has been dedicated to understanding and reducing this loss; though ion traps are usually operated at room temperature (with only the ion cooled to μK levels), cryogenic cooling of the trap electrodes has been shown to allow reduction of heating rates [LGA⁺08], for example from a few hundred quanta/s at room temperature to approximately 5 quanta/s at 4K, using a 1.3 MHz trap frequency with $^{88}\text{Sr}^+$ ions at 50 μm trap height [CS14]. Alternatively, surface cleaning by Argon ion bombardment has also been shown to allow significant reduction at room temperature, by approximately a factor of ~ 100 [HCW⁺12].

Though heating rates at a few quanta per second are in fact comparable to those in the bulk trap used for the highest fidelity two-qubit gates demonstrated so far [Bal14] (at room temperature, typical for bulk trap experiments), where errors were not dominated by motional heating, it will be desirable to further reduce these heating rates, as they are a prominent limit to further reduction of ion trap dimensions. Our work does not contribute to this, except so far as the additional devices we introduce to the trap vicinity should not dramatically increase this heating rate.

1.3.3 Limitations to two-qubit gate fidelities, and addressing

Resource overheads required for quantum error correction scale strongly with the infidelities of physical gates [FSG09], and hence reducing gate infidelities (and reducing them well beyond any fault-tolerant “threshold”, at which the resource overheads are usually extreme) is an important task for achieving an interesting system.

In trapped ions, single-qubit operations have been performed with quite low infidelities, of order 10^{-5} or 10^{-6} using either optical stimulated Raman [GTL⁺16] or microwave [HAB⁺14] transitions between hyperfine states. Two-qubit gates have generally posed a larger problem, owing to the greater complexity involved in coupling to the motional modes, but recently a few groups have achieved infidelities of order 10^{-3} [BHL⁺15, GTL⁺16], though at present it appears these gates will be the limiting factors in most algorithms.

A number of factors contribute to the errors in these gates, discussed in detail in the references cited, many of which have to do with fluctuations in trap parameters or laser noise, which would be independent of the sort of structures considered here. However, we note that in both of the highest fidelity entangling gates presented so far, spontaneous emission during the stimulated Raman transitions used was the largest contribution to error, and hence we expect that the ability to much more tightly focus the excitation light, without introducing additional beam pointing instability (effective intensity noise), may allow larger detunings and lower errors for a given power and gate time. Alternatively, the use of higher intensities together with appropriate pulse shaping could also allow for faster gates, which would suffer less from the errors associated with motional mode heating during the gate. For optical qubit single-qubit gates, similarly [BKR08b], the same arguments about focusing without introducing pointing instability would apply.

A concern separate from the fidelities achievable in single or two-qubit experiments are unintended couplings to neighboring ions when operating on some subset, or crosstalk. Avoiding crosstalk when addressing operations in a system of closely spaced ions is an important challenge, one where the optical approaches considered here can play a significant role, as discussed in more detail later in the thesis.

To sum up, many of the limitations to fidelities currently would remain challenges independently of the way the optics and electronics around the ion system are implemented, but it appears plausible that there would be at least a few significant ways to take advantage of the approaches presented here, primarily for purposes of scaling up operations to larger numbers, to also reduce these errors.

1.3.4 Speed limitations on motional-mode coupling in multi-ion systems

Increasing fidelities on systems of just two ions is one challenge, and implementing similar fidelity operations involving motional modes in larger chains of ions presents additional problems associated to the larger number of closely spaced transitions that can be off-resonantly excited by the pulses involved. Though some pulse-shaping protocols have been developed for attempting to minimize these excitations in few-ion chains without excessively slowing down operations, fidelities so far have been substantially lower in such systems than in two-ion experiments [DLF⁺16]. These problems appear to be largely independent of the approach discussed in this work.

Additionally, even aside from issues arising in multi-ion chains, excitation of unwanted transitions becomes an issue for implementing high-fidelity fast gates (with timescales of order or lower than the trap frequency, typically \sim MHz even for just two ions in a well), and though it appears theoretically possible to circumvent these excitations by appropriate pulse-shaping [PMGL⁺16], it remains to be done in practice.

1.3.5 General ion movement, interconnections between distant ions

Most architectures require some sort of ion movement, to either bring ions on a 2D lattice [KBB11] near each other when enacting a gate between neighboring ions [KHL⁺16] (and adopting such an architecture can help avoid some of the problems associated with multiple ions in the same well discussed above), or to move ions between different registers based on few-ion linear chains [KMW02], requiring also separation of ions in single trapping potentials [HS04]. Implementation of these operations generally requires a series of complex RF and DC potentials, with amplitudes depending on the trap geometry anywhere from many volts and tens of volts. Particularly for geometries requiring high voltages (including the planar trap geometries pursued here, relatively to multi-layer or 3D electrode configurations), this can entail considerable complexity in circuitry. The use of CMOS foundries may address this in part, though for higher voltages it may be necessary to incorporate power transistors, not typically integrated with CMOS environments.

Aside from motion of the ions themselves, distant ions may interact optically, e.g. in schemes where coincident photon detection heralds entanglement of ions at a distance [MMO⁺07, MRR⁺14]. This brings a number of challenges which are not largely addressed here, except speculatively as far as this work may lead to planar dielectric mirrors for cavity integration within trap chips, or perhaps some metallo-dielectric waveguides/resonators which could enter the near field of the ions while still having surface potentials controlled to prevent destabilizing the trap. Such steps would assist in implementing the high- Q cavities that seem likely necessary for sufficiently high success rates with such entanglement schemes, but which have been very difficult to implement with trapped ions so far.

1.4 Concluding comments and thesis overview

The purity of individual atomic ion qubits and the strength of Coulomb-based interactions makes them attractive candidates for quantum information processing. However the fact that these individual ions are doomed to spend their existences floating alone in the void often leads many to expect that they are unlikely to ever really be coaxed operate meaningfully in concert at a large scale. Only in the last few years have individual atoms started to interface with integrated optical devices [TTdL+14, GHY+14], a convergence that may prove very fruitful. By comparison, solid state approaches, primarily those based on superconducting qubits, but also with quantum dots or diamond nitrogen-vacancy centers, for example, have long appeared far more amenable to the same ideas that are behind the success of integrated circuit technology.

The broader theme in this work is that atomic ions' apparent isolation from solid state technologies is only superficial, and that in fact a certain amount of distance between the pristine quantum system and the the tumult of the chip below may even allow for a more favorable combination between the two. That is, the ability to combine scalable optics and classical control in planar-fabricated systems with a quantum system, without sacrificing quality and purity of the qubit to do so, may prove particularly powerful.

The next chapter will present the details of the general approach pursued here, its various features and challenges. Chapters 3 and 4 discuss the waveguide and grating devices used for ion addressing, first describing the optical design and characterization, and next experiments on their integration with planar ion traps. Chapter 5 then discusses CMOS foundry-fabricated ion traps, which are a step towards implementing the previously discussed optics in a readily scalable platform, together with a variety of other components in the same chip. This chapter for example also discusses integrated APDs implemented in the same process and ongoing experiments on these devices. Chapter 6 then discusses an approach to electro-optic modulator devices at the relevant wavelengths, and our initial work towards this, and Chapter 7 summarizes and discusses avenues for future work.

Chapter 2

Overview of the proposed approach

Light is implicated in various steps in quantum information processing with trapped ions – these include ionization, cooling, optical pumping, coherent state manipulation, and fluorescence readout. We want to implement these in a way that we may expect to practically scale beyond few-qubit experiments, to ion numbers where new computational capacity and the advantage over classical simulation or computation becomes clearly manifest. This scale would be as low as around 50 qubits [BIS⁺16], from the perspective of maximally large quantum state whose evolution could be modeled on a classical supercomputer (and hence beyond which scale a quantum simulator may allow new insight), and as high as many millions of qubits, for an error-correcting factoring machine factoring large numbers with reasonable per-gate errors [FMMC12].

In this chapter we discuss the approach to control and measurement pursued in the proposed thesis; the requirements, advantages, and particular challenges are discussed in relation to specific aspects of current experiments as well as other proposals for integration in planar traps. Though an architecture study can become a major undertaking in itself, this chapter aims to still point to the general features of an architecture based on such components as discussed in this proposed thesis and their benefits and drawbacks. A variety of architectures are possible using planar ion trap technology, broadly divided at present into those based on linear chains of some length connected either through ion shuttling between different registers or optical interconnections [KMW02, MRR⁺14], and those based on 2D arrays of ions with short range interactions [CLJ08]. The bulk of this chapter discusses features of this approach which are largely decoupled from this choice of geometry for the ion interactions, and which are generally compatible with any planar ion trap architecture. Our experiments in the subsequent chapters will center on systems using the $^{88}\text{Sr}^+$ ion and an optical qubit transition and many of our particular examples here will concern this species; however, the approach here is of course in principle applicable more generally, and throughout we discuss challenges and opportunities with other species and qubits.

In what follows we first give a brief description of the kind of overall system envisioned; then we describe the properties of the high-index-contrast waveguide structures involved, together briefly with the grating couplers that would direct light from such waveguides to propagating beams focused at the ion locations. The next section

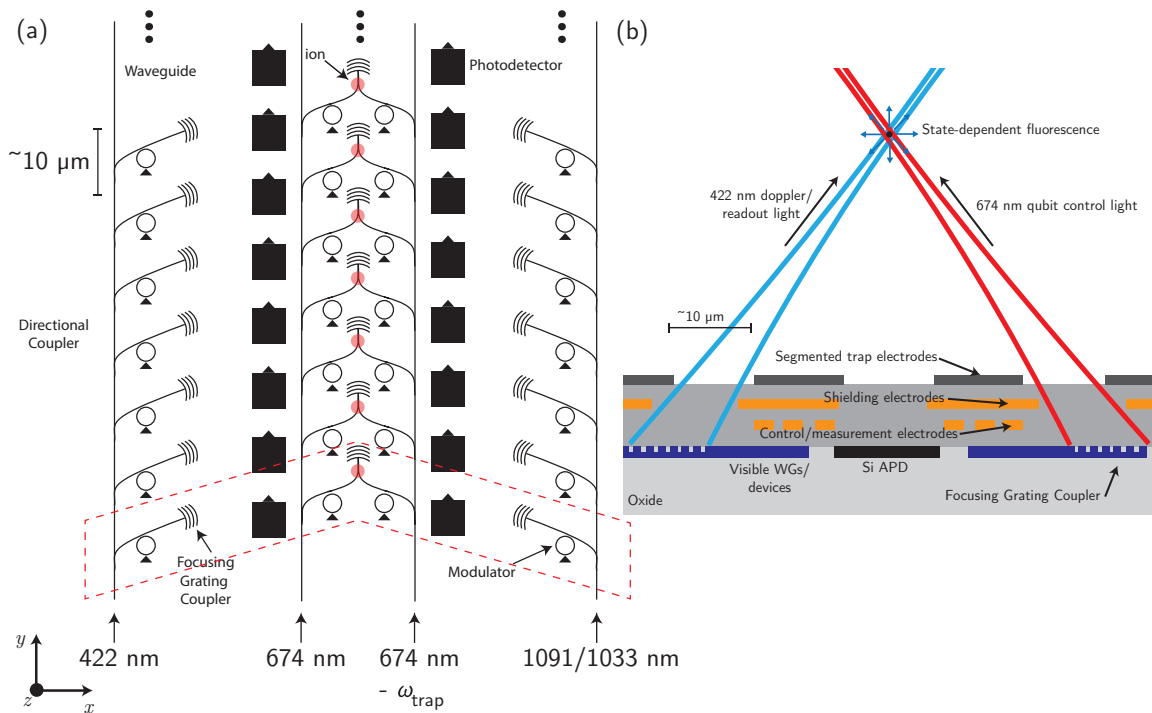


Figure 2-1: (right) Cross-section of a possible chip incorporating the required devices – focused beams are directed to the ions by grating couplers below the trap electrodes, and silicon APDs are incorporated nearby for fluorescence collection. Multiple metal layers (wiring for CMOS circuitry) are employed for shielding between the trap electrodes and devices. (left) Carton illustration of an imagined layout for carrying out operations on a chain of $^{88}\text{Sr}^+$ ions.

discusses in more detail the various advantages and challenges associated with this kind of routing and addressing. We then discuss the possibility and motivations for parallel modulation in such a platform, followed by integrated photodetection and implementation of such ideas in CMOS processes, which we expect will play a significant role in making large-scale systems based on such ideas a possibility.

2.1 Routing, shaping, beam-forming, and detecting light within a trap chip

We envision a system such as was shown above in Fig. 1-6; here power at the relevant frequencies (674 nm and sideband for qubit manipulations and 422 nm for cooling/readout in $^{88}\text{Sr}^+$ drawn, omitting pumping beams) is coupled in through a single bus and split off from the main bus using a variety of splitters, followed by modulators and grating couplers to couple the waveguided light to free-space beams propagating to the ions through openings designed into the the trap electrodes. A slightly more detailed illustration of a cross-section of a chip that may incorporate these functions is shown in Fig. 2-1; both 1D chains and 2D arrays and other trap

geometries could be accommodated in related ways, drawing on low-loss waveguide crossings and μm -scale bend radii achievable in waveguides. The promise of such routing for scalable light distribution and modulation to multiple ion locations in parallel in a way relatively flexible to particular ion trap designs motivates much of this work in general. The next section focuses more specifically on the particular features associated with ion addressing with focusing grating couplers located near the ions as proposed here.

2.2 Waveguide-based routing and ion addressing

Dielectric films can be patterned in diverse ways to create a rich variety of optical devices, including waveguides, splitters, filters, resonators and various couplers. In the field of silicon photonics, a variety of devices have been studied for various applications in classical interconnects [BJO⁺08], and large systems with hundreds of optical elements cooperating nontrivially have been made [STY⁺13].

As compared to single mode fiber optical waveguides, where the optical mode's energy is distributed over an area of order $50 \mu\text{m}^2$ and where the waveguide core has an index only on the order of 0.01 higher than the cladding, nanophotonic waveguides confine the mode's energy to areas of order $0.1 \mu\text{m}^2$, as a result of higher index contrast between the core and cladding (e.g. SiN/SiO₂ – 2.0/1.45, Si/SiO₂ 3.5/1.45). The higher index contrast allows for much tighter low-loss bends than in fiber optics, and denser packing of devices due to the smaller mode area. In addition, many of the more sophisticated devices fabricable in the context of nanophotonics, like photonic crystals, resonators, and grating couplers benefit from the higher index contrast; fundamentally this is because of the stronger scattering at perturbations to a waveguide of a higher index. While a boon in terms of optical devices that can be designed with a small per-device area and densely packed, this stronger scattering makes loss due to waveguide imperfections much stronger. So while optical fibers achieve propagation losses of 0.1 dB/km, nanophotonic waveguide losses are strongly influenced by scattering by high index discontinuities. Si waveguides in current technology are limited typically to around 1 dB/cm, though SiN waveguides can be significantly less 0.01-0.1 dB/cm, though this is strongly dependent on wavelength. In any case, certainly for systems on the scale of a few cm and possibly larger depending on exactly how power is distributed, such a platform allows for dense integration of complex functions.

Example mode profiles for $\lambda = 674 \text{ nm}$ light in waveguides of Si₃N₄/SiO₂ are shown in Fig. 2-2, for light in the fundamental quasi-TE mode. The field is polarized predominantly along x as labeled in the figure, but as is typical for high index contrast, subwavelength-scale waveguides (as well as for wavelength-scale focused beams in free space), the wave is not purely transverse, and appreciable longitudinal components (along z) are present for highly-confined modes. Waveguides significantly wider than the wavelength have much reduced nontransverse field components, as shown, important for the polarization purity generated by such waveguides. Depending on the aspect ratio of the guide, quasi-TM modes (with the field predominantly along y) can be coupled to as well.




Figure 2-2: Example E -field mode profiles of rectangular Si_3N_4 waveguides ($n \approx 2.0$) surrounded by SiO_2 ($n \approx 1.45$), for waveguides of 120 nm height and 540 nm width (top row, $n_{\text{eff}} = 1.59$) and 3 μm width (bottom row, $n_{\text{eff}} = 1.66$), for the lowest-order quasi-TE mode (field predominantly along x) at $\lambda = 674$ nm.

Single-mode waveguides allow light to be routed over large distances and complex paths without phase or amplitude distortions; thus, diffraction is naturally controlled over even complex paths, avoiding a major problem in free-space or MEMS-optics approaches [KK09]. Bend radii give a sense for general footprints of routing components; Fig. 6-3 shows simulations of optical loss per 90° bend for example rectangular waveguides (with widths chosen so that the guides remain single-moded for the quasi-TE modes) at $\lambda = 674$ nm and $\lambda = 1092$ nm. As is evident in both cases bends of a few 10 μm s can provide low losses, and though at longer wavelengths the lower confinement in a film of given thickness restricts bends to larger values, with reasonable bend losses and single-mode operation can be achieved over similar film thicknesses across a wide wavelength range.

A significant challenge in integrating the various wavelengths required to implement all the control, cooling/readout, and pumping required for ion experiments will be the choice of waveguide material that allows operation from the blue (or for some species, as reviewed in Chap. 1, UV) up to the near IR. Table 2.1 lists a few of the possible waveguide core materials, their approximate refractive indices in the visible, short wavelength absorption limits, and maximum electro-optic coefficient. These are approximate numbers and short-wavelength operation can be difficult even above the absorption edge, due to stronger scattering from roughness at shorter wavelengths, as well as to optical damage, e.g. UV solarization in silica and photodamage in LiNbO_3 . Prospects for this will be discussed in the section of wavelength ranges and polarization below.




Figure 2-3: Simulated bend-loss vs. bend radius for representative rectangular silicon nitride waveguides (dimensions labeled) at $\lambda = 674$ nm and 1092 nm. The Si_3N_4 core index here is wavelength dependent but approximately 2.0, and that of the SiO_2 cladding is approximately 1.46.

2.2.1 Focusing couplers for ion addressing

Optical devices near the ions allow for the equivalent of high NA optics to be employed in directing light the ions; whereas planar ion trap experiments using beams propagating parallel to the chip surface require beams of approximately $30 \mu\text{m}$ diameter (so as to keep scattering from beam clipping on a 1 cm^2 chip) [KK09], couplers emitting light from the trap chip itself could focus to spots on the order of $1 \mu\text{m}$.

We note that more unusual trap structures with a narrow raised pedestal forming the optical access region [Mau16] with beams propagating still parallel to the chip surface, or structures in which the trap chip is etched entirely through the substrate and the mount is arranged such that the beam can pass through normal to the surface [CMBK14] can bypass the beam-clipping concern mentioned above and allow μm -scale focusing without integrated optics. Both certainly enable tightly focused addressing of ions in a single linear trap. However, the first case does not allow such addressing of an array of linear traps, essential for eventual scalability. The latter requires etching through the substrate and chip mount to avoid unwanted scatter of the surface-normally propagating beam, which would significantly limit the density of addressable zones, as well as the possibility of integrating other functionality into the

Material	n at 632 nm	Absorption edge (λ)	Max EO coefficient (pm/V)
Si ₃ N ₄	2.0	290 nm [Pal98]	0
LiNbO ₃	2.3	302 nm [Won02]	33
GaN	2.4	365 nm (bandgap)	3
AlN	2.1	206 nm (wurtzite-phase bandgap)	1
Al ₂ O ₃	1.8	< 250 nm [AWB ⁺ 10]	0

Table 2.1: Table of possible waveguide materials; values of absorption edge, as estimates for pure material with references, are to be taken with a grain of salt, as the level absorption near the edge will depend on material quality in practice, along with possible optical damage by higher energy (UV) photons. Ease of fabrication and etching with fine feature size and acceptable roughness is a further concern.

substrate. Thus both are applicable for experiments in 1D array of ions (and indeed the traps used in such experiments are designed for this), but not clearly applicable to larger systems requiring 2D geometries.

Smaller focuses implemented in scalable geometries have significant implications for quantum operations, which stem basically from the lower total power required for the same intensity at the focus (3 orders of magnitude, comparing 1 and 30 μm diameter waists), and the tighter spatial localization of the focused light. The next chapter will detail the design and demonstrate these devices; here we aim only to establish the relevance of capability to trapped ion QIP.

2.3 Advantages and limitations of waveguide-based routing and focusing

2.3.1 Laser power

Laser power to manipulate a multi-qubit system in parallel becomes appreciable. Table 2.2 lists the various wavelengths required for manipulation of Strontium qubits and the powers necessary per qubit, and for convenience the corresponding power in a 50 μm -diameter beam as would typically be used in a planar ion trap experiment with free-space optics. So for the 674 nm light, in a 5 μm^2 beam, area, approximately 25 μW would be required per ion, to enact ~ 1 MHz single qubit operations and two-qubit gates of interaction times of a few 10 μs . For a given laser power (current narrow linewidth lasers and tapered amplifiers can output of order 1 W of Hz-level linewidth

Function	λ (nm)	Order of magnitude intensity (mW/cm ²)	Switching speed
Qubit operations, sideband cooling	674	10,000-100,000	1-10 MHz
Doppler cooling and readout	422	100-1000	1-10 MHz
Repumping	1091, 2033	10-100	Slow
Two-photon photoionization	405, 460	1000-10,000	Slow

Table 2.2: Required wavelengths, approximate intensities and switching times for trapping and quantum state manipulation in $^{88}\text{Sr}^+$. The power orders of magnitude very approximately follow the values given in [Wan12].

radiation), the number of ions individually addressable in parallel is drastically higher ($500 \times$) with focusing as compared to with $2500 \mu\text{m}^2$ beams.

2.3.2 Raman gate operation, detuning and spontaneous scattering

Although in our work we use an optical transition on the quadrupole transition in $^{88}\text{Sr}^+$, qubits encoded in hyperfine ground states of ions with nuclear spin would allow for longer coherence times not limited by the spontaneous emission of the quadrupole transition (of order 1 s). Here, transitions between the two levels can be enacted by pairs of lasers with difference in frequency equal to the spacing between the qubit levels ω_0 , and each with a detuning Δ from one of the P levels (chosen to be with respect to the $P_{1/2}$ level by convention – positive Δ corresponds to a photons blue detuned with respect to this level), which effect two-photon stimulated Raman transitions between the levels. Here, however, for a given Δ a fundamental error source is present here which results from spontaneous scattering during a Raman transition from the excited state. By increasing the laser detunings from the excited state (p - orbital), the probability of this decay can be reduced, at the cost of needing higher optical power to maintain the operation speed.

Early considerations of this effect assumed the qubit (encoded in hyperfine states in s-orbital manifold) would decay upon any spontaneous scattering event from the excited state. A π -pulse time for the stimulated Rabi transition is given by $t_\pi = \pi/2\Omega_R$, with the Raman Rabi frequency scaling as [WBB⁺03]

$$\Omega_R \propto \frac{\omega_f}{\Delta(\Delta - \omega_f)}, \quad (2.1)$$

where ω_f is the spacing between the $P_{1/2}$ and $P_{3/2}$ states. Thus t_π goes as Δ^2 for large Δ , but since the elastic Rayleigh scattering rate from the p levels goes as Δ^{-2} , the probability of a scattering event in a π -pulse time stops decaying for large Δ , and no benefit in gate fidelity would be achieved. Thus it seemed that there was an optimal detuning, with Δ limited on the order of ω_f , and no benefit could be had in gate fidelity by increasing detuning (and intensity) to larger values.

However, it was later realized that Brillouin scattering events, in which the ion's electron ends up in the same state it was in to begin with, carry away no information about the qubit state as long as $\Delta > \omega_f$, and hence elastic Rayleigh scattering does not contribute to decoherence [OIB⁺07] in this limit. Inelastic scattering into the other ground states was shown to scale as approximately Δ^4 for large detunings, and as a result, owing to the fact that only such inelastic scattering contributes errors, large detunings can in fact result in arbitrarily low operation errors from spontaneous scattering back into the ground state manifold, and for species with low-lying D levels the probability of scattering into the D-level becomes the limiting factor for infinitely high detunings.

Hence, higher intensity becomes an asset in reducing Raman gate errors from inelastic spontaneous scattering, as a constant π -time can be maintained while increasing detuning while increasing intensity proportional to Δ^2 . This brings about a reduction in error that goes as Δ^{-2} , i.e. for a given gate time, for both single and two-qubit gates the gate error decreases linearly with available optical intensity at the ion.

Even at the highest-fidelity gates demonstrated to date (with Raman-addressed hyperfine transitions), this power limitation is a dominant contributor to loss. In the experiment of [Bal14], Raman photon scattering was the largest known source of gate infidelity even in the slowest gate ($t_g = 100 \mu\text{s}$), using 5 mW per beam in spot sizes of $w = 27 \mu\text{m}$ and a detuning of -3.0 THz. Faster gates obtained by decreasing the detuning showed higher infidelity due to photon scattering. A robust method of focusing a given amount of power more tightly to μm -scale spots, can reduce this scattering contribution to error by two orders of magnitude, essentially removing it as a significant source of error. The clear need for high powers in the laser beams addressing these transitions, especially for multiple ions in parallel, has provided incentive for development of new laser sources capable of high power and stability [BLWW15]; so long as other losses are managed well, tight focusing to each ion site, by allowing efficient use of whatever power available, would allow an optimization of error and parallelization of such operations. While this advantage is clear, the challenge will be in achieving low-loss waveguides at the short wavelengths used for Raman-addressed gates; prospects for this will be discussed below.

2.3.3 Pointing instability

Beams in ion trap experiments typically propagate in free space over many meters from their source before interacting with the ions. Small positional or angular fluctuations in the various optics in the beamline can translate into significant displacements at the ion, resulting effectively in intensity noise. With much care, this can be reduced

to a level lower than other noise sources, but this generally poses a significant challenge; and furthermore since a given displacement corresponds to a larger intensity variation for tightly focused beams it is a larger challenge the more tightly the beam in question is focused. Light launched from the same chip as the ion would of course eliminate this problem entirely at the ion, though it would translate it into errors resulting from fluctuations in where the light was coupled into the chip. However this allows us to decouple the amount of focusing at the ion location from the pointing instability-induced fluctuations, and as we outline below can be expected to be much more robust than when light is directly incident on the ion.

If we consider a coupling a beam propagating in free-space into a grating coupler on-chip designed to couple the free-space Gaussian to a waveguide mode, the input power coupling efficiency would be given by

$$\eta = \left| \int E_g^* E_i dA \right|^2, \quad (2.2)$$

where E_g is a normalized ($\int |E_g|^2 dA = 1$) field profile corresponding to that emitted by the grating and E_i is that for the incident beam. To evaluate the effect of a given beam displacement on the coupling, we consider the ideal case of Gaussians for both profiles, with matched beam radii w_0 and some displacement d . Along one dimension, $E_g = w_0^{-1/2} \pi^{-1/4} \exp(-x^2/2w_0^2)$ and similarly for E_i , so the coupling would be:

$$\eta = \left| \frac{1}{w_0 \sqrt{\pi}} \int \exp\left(-\frac{x^2}{2w_0^2}\right) \exp\left(-\frac{(x-d)^2}{2w_0^2}\right) dx \right|^2 = \exp\left(-\frac{d^2}{2w_0^2}\right), \quad (2.3)$$

which should be compared to the relative intensity for a beam displacement d if the same beam was directly incident on the ion, simply proportional to $|E_i|^2 = \exp\left(-\frac{x^2}{w_0^2}\right)$. Hence for small displacements, (small arguments in the exponentials and using $\exp(-z) \approx 1 - z$ for small z) the power deviation at the ion for a given displacement of the beam would vary $\sim 2\times$ less if the power is coupled into a grating than if a beam of the same radius w_0 were directly incident on the ion. This modest advantage in coupling is made more significant by the fact that the input coupler and beam can be chosen to have a dimensions much larger than that of the beam addressing the ion; in this case, the fractional deviations due to displacement d at the coupler are $\frac{d^2}{2w_{0,g}^2}$ and for the same displacement at the ion $\frac{d^2}{w_{0,ion}^2}$, and the deviation when waveguide-coupled is lower by a factor of $\frac{w_{0,ion}^2}{2w_{0,g}^2}$, which can be very significant.

This would be best realized in practice when directly fiber coupling to a chip, and in this case the coupling deviation would occur at the collimator addressing the fiber which could have a large beam waist for example of many millimeters. These points are illustrated in Fig. 2-4.




Figure 2-4: Deviations in intensity relative to peak vs. displacement in units of w_i , the waist of a beam addressing an ion. Red line shows the case of direct incidence on the ion; black line when a beam of the same waist is incident on a matched grating coupler; and black dotted line when a beam with waist $10w_i$ is in-coupled through a larger grating, with greater tolerance to absolute displacement.

2.3.4 Phase stability

Also in the context of quantum information, waveguide systems have been investigated also for linear quantum optics systems [PCR⁺08, MPSO09, SBO⁺14], where a noted advantage is the phase stability arising from the fact that path lengths are fixed (in waveguide geometry) and not susceptible to air circulation or air temperature variations which would affect experiments with free-space optics, where path lengths are in addition typically significantly longer. The same advantage of waveguide approaches would apply to trapped ion QIP in multi-ion systems, when phase-stable light delivery to multiple ions would be required for quantum gates across the system.

Already in small-scale processors (3-5 ions), the need for phase-stability between beams if each was involved in coherent population transfer motivated an approach in which one was significantly detuned and used only for individually addressed Stark shifts, and the other effected population transfer globally [SNM⁺13]; this was chosen since the Stark-shifting beam need not be phase-locked to the other, reducing a significant difficulty (and error source) for the free-space optical system employed. In large systems, an approach relying on a single global addressing beam would be infeasible, and an approach that allows phase-stable distribution of light routed along complex paths to multiple separate locations will be required; waveguide-based approaches, in which a single beam can be input at one location and split a few hundreds or more ways [STY⁺13] should be capable of providing this essential capability.

2.3.5 Individual addressing

Implementation of single or multi-qubit gates on particular elements of an array, required for any general quantum information processing, will rely on individual addressing of different ions' qubit transitions. The two-ion spacing, for two-ions in a linear trap with axial frequency $\omega_z/2\pi$, is given by [WMI⁺97]:

$$s_2 = \left(\frac{2q^2}{4\pi\epsilon_0 m\omega_z^2} \right)^{\frac{1}{3}}, \quad (2.4)$$

where the weak dependence on m , the ion mass, implies that ions with very different masses have spacings on the same order, as shown in Figure 2-5 for $^{88}\text{Sr}^+$, $^{43}\text{Ca}^+$, and $^9\text{Be}^+$. Since gates typically target axial frequencies of ~ 1 MHz to not suffer in speed and/or error due to off-resonant excitation at a given speed [WMI⁺97, SM00], and since with more ions added to the chain the spacings decrease from these values somewhat, to address ions in a chain clearly requires spatial addressing of only a few μms .

A straightforward way to do this is to focus a laser beam down to a spot size small compared to the inter-ion spacing, and is indeed the way this has been done so far in many experiments in bulk traps [NLR⁺99, JLH⁺14]. In planar ion traps, with beams propagating parallel to the surface, diffraction limits the waist that can be achieved for a given chip size before the beam will clip on the edges of the trap, to




Figure 2-5: 2 ion spacings s_2 for typical axial trap frequencies.

approximately 20-30 μm for a 1- cm^2 chip, and optical individual addressing has not been demonstrated in standard planar traps, but only in a chip in which the substrate was fully etched through, opening a trench to allow light to pass through the chip perpendicular to the surface [MBB⁺13, CMBK14]; however, the complete etch is a complex procedure that would limit further integration within the trap chip and it may be difficult to scale such an approach requiring an etch through the full substrate (wide enough to allow the diffracting beam to propagate out over the thickness of the substrate) to parallel operations. Focusing to μm -scale spots from within a standard single metal layer trap chip, as demonstrated in the subsequent chapters, may provide a comparatively simple solution.

Other approaches to individual addressing in planar traps include the use of magnetic field gradients to shift transition energies across space [WLG⁺09], or with purely microwave approaches in which microwave field gradients generated on chip provide the same selectivity [WOC⁺13]. However, both present particular challenges to scaling up; nevertheless, given the promise of high-fidelity single-qubit microwave gates for hyperfine qubits with microwave fields from on-chip electrodes [HAB⁺14], an interesting possibility, for optimal and addressed single-qubit gates, may be to provide individual addressing capability to such an approach by using focused beams to locally Stark-shift hyperfine levels.

2.3.6 Wavelength ranges and polarizations

Although for $^{88}\text{Sr}^+$ the shortest wavelength required is 405 nm, as alluded to above, to accommodate the variety of ions used by experimental groups, waveguides extending into the UV would be required. $^{43}\text{Ca}^+$, used by a number of groups now, requires a similar range, down only to 397 nm; however, to accommodate the $^9\text{Be}^+$ and $^{24}\text{Mg}^+$ ions used by the Colorado group, wavelengths down to 313 nm and 297 nm, respectively, would be needed. Silicon nitride, used in our first devices, has been shown to allow losses < 1 dB/cm at 470 nm [GJG⁺08].

While materials do exist that have been shown to show thin-film guided modes with losses of a few dB/cm down to even lower wavelengths than any needed for the ions mentioned above, such as Al_2O_3 (down to below 250 nm) [AWB⁺10], and though good etching of rectangular waveguides in the infrared has been demonstrated in the same material [BAWP07], patterned single-mode waveguides below 400 nm have not been demonstrated with low losses to the best of our knowledge. Due to the scaling of scattering loss from sidewall roughness with decreasing wavelength, this could prove a significant challenge, and prospects for operation in the UV will require significant study.

While all functions in $^{88}\text{Sr}^+$ are enacted by linear polarizations, some functions in other ions (particularly optical pumping and readout in species with complex ground state manifolds) require circular polarization, and furthermore with high purity when polarization (as opposed to frequency selection, as for example in pumping into one of the two ground state sublevels in $^{88}\text{Sr}^+$) is used as the means of transition selection. While our approach can eliminate the impurity arising from nonzero birefringence in thick vacuum windows, achieving high-purity circular polarization from waveguide devices may not be trivial. While our approach is thus most immediately applicable to species (or operations) which can rely on linearly polarized beams (e.g. most coherent operations on both quadrupole and hyperfine transitions), prospects for circular polarization-generating components which would increase the generality of the approach will be discussed in the last chapter of the thesis.

2.3.7 Power handling and self-phase modulation effects

For systems directing light in parallel to large numbers of ions, light at each frequency would be input at one or a few points and split on chip to the various sites. Depending on material used, the total power at the input could reach limits where nonlinear loss mechanisms such as two-photon absorption (TPA) would set in, or else where whatever small linear absorption is present results in sufficient heating to damage the waveguide.

Power-handling capabilities of low-loss SiN waveguides however can be quite high; due to its high energy absorption edge, 2 W of 1550 nm light have been launched without damage in a $2.8 \times 0.1 \mu\text{m}^2$ waveguide as pump for nonlinear frequency conversion [KAS⁺15]. If 25 μW is required per beam for ~ 1 MHz single-qubit gates at 674 nm, and supposing an average 6 dB total loss between the input waveguide and each ion location, 1 W would suffice for 10,000 ions; one bus can thus be expected to




Figure 2-6: Phase shifts per unit length due to self-phase modulation vs. power, for a SiN waveguide with $A_{\text{eff}} = 0.1 \mu\text{m}^2$, supposing $n_2 = 2.4 \times 10^{-15} \text{ cm}^2/\text{W}$ and $\lambda = 674 \text{ nm}$.

handle light for a large number of ions. However, two-photon absorption may occur at shorter wavelengths, particularly in the blue; this is expected to occur only for wavelengths less than around 470 nm [TISF10], although it does not appear to have been well characterized in silicon nitride.

A problem likely to arise earlier on is self-phase modulation over long optical path lengths due to the intensity-dependent refractive index, which could modulate the phases of the beams seen at the various ions. This is usually parametrized by an intensity dependent refractive index $n = n_0 + n_2 I$, where n_2 has been measured to be approximately $2.4 \times 10^{-15} \text{ cm}^2/\text{W}$ at $\lambda = 1550 \text{ nm}$ in silicon nitride [ISAF08]. Using this number to get a rough feel (these coefficients have not been characterized for SiN in the visible), and noting that the phase shift due to this effect over a length l is

$$\phi_I = n_2 \frac{P}{A_{\text{eff}}} k_0 l, \quad (2.5)$$

where P is the guided optical power, A_{eff} is an effective area that accounts for the waveguide mode profile, group velocity, and confinement, and k_0 is the free-space wavevector, the phase shift per unit length owing to this effect for $\lambda = 674 \text{ nm}$ light in a waveguide with a typical $A_{\text{eff}} = 0.1 \mu\text{m}^2$ is plotted as a function of P in Fig. 2-6.

This rough analysis indicates that for mW level powers, phase shifts over 1 cm are at the $\pi/10^{-4}$ level; at the 1 W level (i.e. in the input waveguide carrying light to multiple ions), this shift can reach $\pi/10$ over 1 cm. Thus for light that has been split, this should be a manageable effect; and indeed prior to splitting, this self-phase modulation would equally affect light that will be addressed to each ion simultaneously, and hence would entail only approximately a global phase-shift. Hence it does not appear to be a large concern, but this is a coarse analysis that only

gives a rough idea for this effect.

These nonlinear loss and Kerr effects have not been well characterized in SiN and particularly in the visible, and measurement of these effects in materials suitable for visible waveguiding would be essential (and interesting generally for applications for visible integrated optics) for a more precise quantification of the problems they pose.

2.4 Parallel modulation

The light directed to each ion has to be switched and modulated in time to implement meaningful operations, and eventually a method to do so in parallel for each of the signals going to each ion will be required. Some form of parallel single or two-qubit gates are involved in many quantum algorithms, but as one illustrative example, we can consider an error correcting system, in which errors would be diagnosed by a series of gates applied to each logical qubit followed by particular measurements [FMMC12]. Error diagnoses on each logical qubit are independent, so an inability to perform these diagnoses in parallel would entail a cost in time (and hence more stringent requirements on errors) proportional to the number of logical qubits. A second example is provided by transversal two-qubit gates, in which a two-qubit gate between two logical qubits is enacted by carrying out in parallel a two-qubit gate between each of the multiple pairs of physical qubits of the logical qubits; such operations could make clear use of parallel operations and would be sped up by a factor equal to the number of physical qubits per logical qubit.

This parallelism could be achieved a few ways; reliance on optical modulation can in the first place be tempered through the use of motional gates, where operations are encoded through ion movement through stationary beams [LKOW07]. The free-space AOMs presently used could be replaced by fiber-coupled AOM modules and multiplexed. But perhaps the most desirable, though also the most ambitious, would be multiplexed planar-fabricated modulators, which as in classical integrated photonic systems [XSPL05] could in principle be integrated with the waveguide and grating coupler devices proposed for addressing.

Such waveguide modulators are typically based on some material with a χ_2 nonlinearity, i.e. with an electro-optic effect. Candidate materials include LiNbO₃, which is widely used for optical modulation in telecommunications and has in the last few years been employed for thin-film photonic devices [PHSG12]; and AlN [XPS⁺12], which has a lower EO coefficient but can be sputtered and is more likely to be compatible with CMOS approaches. These are two among many possibilities (Table 2.1).

For a LiNbO₃ waveguide, given the electro-optic coefficient of $r_{33} \approx 31$ pm/V, a π -phase shift could be achieved with a 5 V amplitude pulse across a distance of 2.5 μ m in a waveguide length of approximately 1 mm (see Ch. 6 below), and hence a phase modulator or Mach Zehnder intensity modulator could have this length; resonant modulators could be significantly smaller. Specific designs and the challenges for modulators particular to trapped ion QIP will be discussed in Chapter 5, but we note here that if such devices could be reliably achieved according to the requirements of trapped ion QIP, these dimensions would allow packing multiple phase and intensity

modulators onto a single chip, integrated with the addressing optics.

2.5 CMOS integration and on-chip photodetection

A few considerations motivate implementation of ion trap systems in CMOS settings. Eventual systems will require tens to hundreds of thousands of electrodes [Ste04], and reliable, scalable electrode fabrication, utilizing multiple layers of metals to connect particular and possibly distant electrodes, is hence an obvious requirement. CMOS fabrication would present a clear path to creating such large systems, and together with the waveguide devices and systems described above [OMS⁺12, SGO⁺15, MOTZ⁺14].

In addition, trapped ion quantum computers will likely require significant classical computing resources for pulse shaping of qubit control signals [KPM⁺05], and trap electrode voltage control for ion motion [BGL⁺12, HOS⁺06]; and error correcting quantum computers in any qubit platform will make heavy use of classical computation for error syndrome diagnoses and feedback [Ste04]. To minimize delays from these classical operations, integration of electronics within the trap chip will be desirable, and CMOS approaches enable this in a very robust way.

These considerations alone are strong justification to consider CMOS approaches, which even in the short term will be desirable for research groups so as to avoid the significant labor involved in multi-layer fabrication even for relatively small chips of 10s of electrodes. In addition however, the availability of avalanche photodiodes (APDs) in standard CMOS processes offers a route to on-chip detection. For example, in a 130 nm process node, APDs with lower than 100 Hz dark count rate at room temperature and detection efficiencies of over 20% at 425 nm have been achieved [FLC⁺10], with yet higher performance in devices implemented in CMOS imaging processes [WGH12]. Such devices would offer a few advantages for trapped ion QIP. Most straightforwardly is again the feedback requirement for error correction – the data from measurements on ancilla qubits must be processed to determine operations on the other physical qubits, and detectors together with electronics on chip can significantly reduce the delay associated with this step.

But, given an ion height, the efficiency with which these photons can be collected, and what size of detector is required, and the confidence with which photons from distinct ions can be distinguished are crucial metrics, and we discuss these here.

To get a rough feel for collection efficiencies, if the detector collects light with a cone of half-angle θ , the collected fraction is calculated as

$$\eta_c = \int_0^{2\pi} d\phi \int_0^{\theta_r} d\theta \sin(\theta) |T(\theta)|^2, \quad (2.6)$$

where for a trap height h and detector radius r , $\theta_r = \arctan(r/h)$ is the maximum angle accepted by the detector area, and T is the angle-dependent transmission into the device from vacuum. Neglecting this reflection, this reduces to the solid angle fraction $\eta_c = \sin^2(\theta/2)$. We can plot this efficiency, assuming for convenience a




Figure 2-7: Detector collection efficiency for isotropic emission (as on the $5S_{1/2} - 5P_{1/2}$ transition in $^{88}\text{Sr}^+$) η_c , neglecting detector quantum efficiency, as a function of detector radius (expressed as a fraction of trap height). Dashed and dot-dotted blue and red curves show the same accounting for angle-dependent reflection for s- and p- polarized light at the vacuum-oxide interface above a detector. The thick horizontal dashed line marks the $\sim 4\%$ collection efficiencies achievable with bulk optics with $\text{NA}=0.4$, with the vertical dashed line marking the line $r/h = 0.5$ which is the minimum ratio required to achieve comparable efficiencies.

circular detector, as a function of the detector radius r as a fraction of trap height h , as shown in Fig. 2-7, including also the curves accounting for non-zero reflection of s- and p-polarized light for an oxide interface ($n \approx 1.47$ at 422 nm); the emission from the ion will be some combination of the two depending on emission angle, but the solid angle fraction is clearly the dominant factor. For comparison, an imaging optics in present setups to collect emitted fluorescence have NAs of approximately 0.4 [SKC12], which corresponds to $\eta_c = 0.042$, and implying that a detector with $r \approx 0.4h$ is required to be comparable in collection efficiency. Already this is an appreciable area to be taken up beneath the ion, and for a 50 μm trap height would significantly limit area available for couplers and electronics, posing additionally a challenge for trap design leaving so much area uncovered by electrodes (unless transparent electrodes can be applied at scale).

Thus having detectors close to the chip does not itself guarantee higher effective NA than is possible with bulk collection optics, though it can reduce total system size drastically. And in the future, if heating can be alleviated and ions brought much closer to the surface, e.g. with $h = 5 - 10 \mu\text{m}$ (or, if the height can be dynamically reduced immediately before a readout step, so that ions are trapped at a larger height for computation steps and brought closer to the chip only for the readout steps), detector areas can be much more manageable compared to ion spacings.

Furthermore, such detectors would not be able to discriminate ions separated by less than the trap height, as can be done by imaging fluorescence onto an EMCCD. This could be dealt with by reading out ions sequentially, illuminating only one at a time with the readout light, but at the expense of parallel readout, and given the length of readout times required (of order 100 μs , to scatter $\approx 10^4$ photons in Sr and collect $\approx 10^2$) this would entail a significant cost. While it may be possible to employ arrays of detectors to allow some kind of computational discrimination of ion locations even for closely-spaced ions, this may prove quite challenging and resource intensive; nevertheless it may offer another route to spacial selectivity without sacrificing parallelism.

Even if logical operations are carried out with ions at a 50 μm height, both concerns above, regarding NA and ion discrimination, could be addressed in principle with traps designed with variable heights such that ions could be brought near the surface before readout (the readout operation is insensitive to heating of the ion motional mode), allowing detectors with $r/h < 0.1$ to still collect a few percent or more. This would be a significant design challenge in its own right, but quite possible in principle. Successful on-chip detection, even if at first low efficiency, would give serious motivation to develop such drastically variable-height trap structures for a large-scale system, or for example transparent electrodes that would allow large detector areas without overly perturbing the trap potentials.

The dark count rates of APDs can be low, especially at low temperature. Error rates of $< 10^{-5}$ are achievable [MSW+08] with dark counts per measurement time 10 times below the signal counts with the ion in the bright state; hence, if 5% of the photons emitted near saturation at about 100 MHz are collected with a 20% QE giving signal counts at 1 MHz, we should have below at most 10-100 kHz dark counts, readily achievable with CMOS detectors demonstrated to date. A number

of potential challenges can be anticipated, however; in particular, APDs operating at low temperatures (close to 4K) as appears advantageous for ion experiments, can suffer from afterpulsing phenomena which can translate even extremely low rates of thermal carrier generation in to excessive noise [RLA⁺07]. Additionally, if not sufficiently shielded, the high RF voltages applied to the trap may interfere with the detector operation, or the Volt-level pulses experienced by the detector may couple to the potential experienced by the ion, and either coupling may prohibit cooperation of these devices.

2.6 Summary and overview of chapters ahead

Ion traps designed together with integrated optics for both ion addressing and read-out, and furthermore in a platform that allows close integration of control electronics, would help ion trap quantum systems scale beyond few-qubit demonstrations, while also bringing a number of advantages in performance and efficiency. The following chapters will describe our steps towards this goal, first in the design, fabrication and characterization of the focusing waveguide devices that couple the ions and waveguides (Chapter 3), the first experiment demonstrating ion addressing with such optics (Chapter 4), our work on CMOS-fabricated traps and traps with integrated APDs (Chapter 5), and on integrated electro-optic modulators for the visible (Chapter 6), followed with a concluding chapter discussing outlook and further work motivated by that here.

Chapter 3

Design, fabrication, characterization of focusing waveguide grating couplers

As discussed in the previous chapter, waveguiding optics integrated together with planar ion traps would offer a promising route to scalable optical manipulation of trapped ion quantum states. This chapter aims to address questions related to the experimental feasibility of such an approach, starting with the basic element of a coupler to direct light in the waveguide to the ion above the chip. This is done in such a way as to focus the light to a sufficiently small spot to both couple to a single ion of a group, and to minimize power requirements.

This chapter discusses the design, fabrication and optical characterization of these gratings, with much of the material based on that presented in [MR16]; experiments regarding their integration with planar traps and ion addressing are presented in the next.

3.1 Overview and previous work

The focusing grating devices presented here couple visible-wavelength light from single-mode, high index-contrast dielectric waveguides to free-space beams forming micron-scale diffraction-limited spots a designed distance and angle from the grating.

With a view to application in spatially-selective optical addressing, and in contrast to previous work on similar devices, deviations from the main Gaussian lobe up to 25 microns from the focus and down to the 5×10^{-6} level in relative intensity are characterized as well; we show that along one dimension the intensity of these weak sidelobes approaches the limit imposed by diffraction from the finite field extent in the grating region. Additionally, we characterize the polarization purity in the focal region, observing at the center of the focus a low impurity $< 3 \times 10^{-4}$ in relative intensity. Our approach allows quick, intuitive design of devices with such performance, which may be applied in trapped-ion quantum information processing and generally in any systems requiring optical routing to or from objects 10s–100s

of microns from a chip surface, but benefitting from the parallelism and density of planar-fabricated dielectric integrated optics.

A number of systems may employ integrated waveguiding optics, formed in a planar dielectric layer, that also require directing light to objects external to the chip. In atomic physics these may include atom chips [KAZ⁺16], broadly speaking, in which trapped atoms are manipulated in close proximity (typically 1-100 microns) to a chip which defines a trapping potential, or in planar ion trap devices [MBM⁺16], for scalable implementations of experiments relying on quantum control of individual trapped ion qubits [SNM⁺13, DLF⁺16]. In such experiments, highly precise control over the beam profile is often necessary, a challenge especially when combined with the requirement for scalability. Other areas may include structures to create and efficiently illuminate large arrays of focused spots for certain microscopy techniques [WCZ⁺10, OC12], waveguide-coupled arrays optical trapping potentials [DG98], components for optically-assisted data storage [MIH06, CPI⁺09], or targeted delivery of light to multiple sites for biological experiments requiring optical inputs [PRH13].

The designs presented here can be generated with simple numerical calculations and two-dimensional electromagnetic simulations of uniform periodic structures; hence designs can be drawn relatively rapidly, and this approach may serve as an efficient starting point for further numerical optimization. In contrast to previous work on waveguide devices generating focused beams [USNK86, HO81, SLY⁺97, SGG99, KNI⁺04], these devices are fed by single-mode (SM) waveguides on-chip which allows precise tailoring of the transverse field profile and control over both low-intensity sidelobes and polarization purity of the beams generated, which we characterize here as well.

3.2 Grating design

In designing the devices, amplitude and phase shaping of the output is considered separately for the dimensions along and transverse to the propagation in the waveguide layer. The approach used along both dimensions is described below, after a brief discussion of the considerations that motivated the choice of device dimensions.

3.2.1 Dimensional considerations for ion addressing

Ion traps currently are constrained by heating rates to operate with ions a few 10s of microns from the chip surface, and we designed couplers for use with traps with a 50 μm ion height. Diffraction imposes a straightforward tradeoff between the spot size at the focus and the required size of the grating required to generate this spot, which can be considered precisely using the standard equation describing the evolution of a gaussian beam waist as a function of distance d along its propagation direction (given by θ) from the waist:

$$w(d) = w_0 \sqrt{1 + (d/z_R)^2}, \quad (3.1)$$

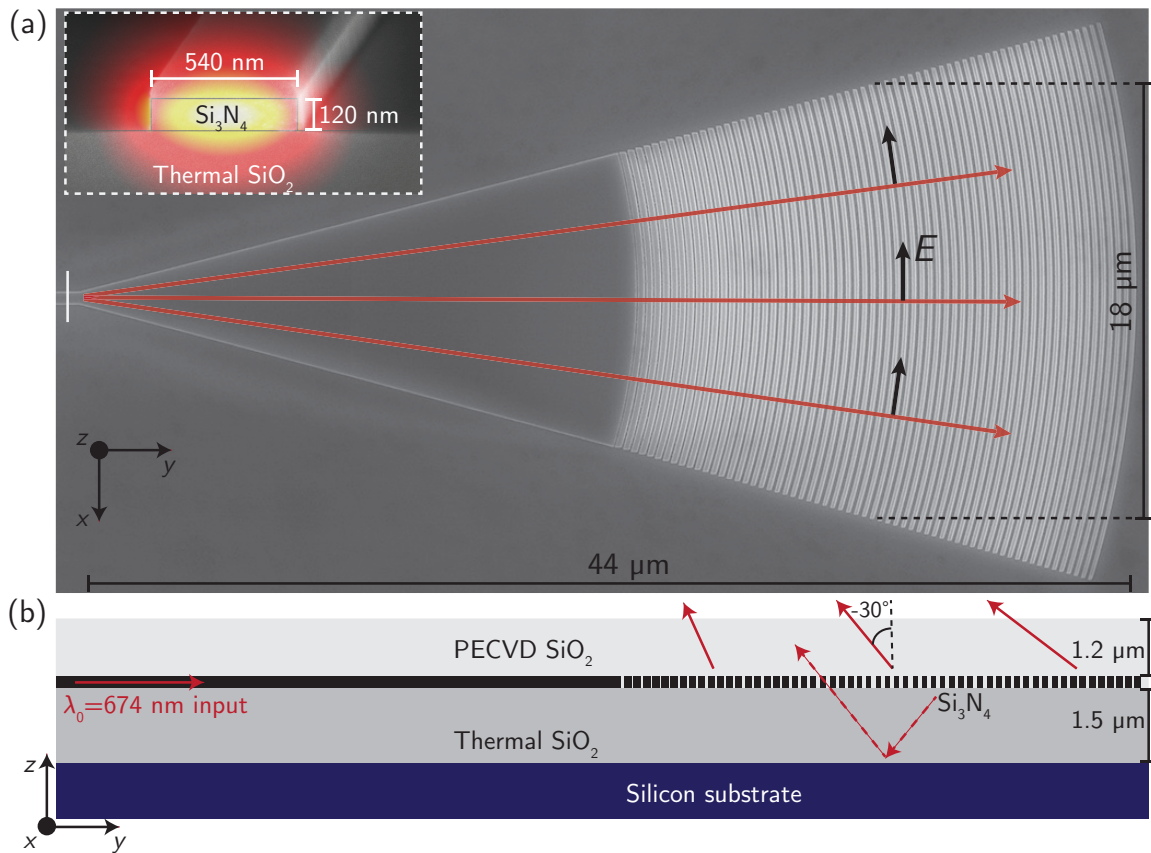


Figure 3-1: Device overview. (a) Scanning electron microscope view of the device, showing the taper from the SM waveguide and the grating region. The inset shows the single-mode quasi-TE waveguide feeding the device (white line in main image), overlaid with a simulated field profile (field points predominantly horizontally). (b) Schematic cross-section of device, showing varying angle of emission along grating length and the two interfering emission paths.




Figure 3-2: Plot of beam waist in the chip plane for a given focused beam waist at height $z = d \cos(\theta) = 50 \mu\text{m}$ above the chip, for different beam angles θ as defined in Fig. 3-1b.

where w_0 is the waist at the focus and the Rayleigh range is

$$z_R = \frac{\pi w_0^2}{\lambda}. \quad (3.2)$$

The resulting beam waists as a function of the focused waist for a few different propagation angles is plotted in Fig. 3-2. As shown below, along the transverse direction, the waveguide width required in the grating section for a given beam waist w is equal to $\sim 2.844 \times w$. Our designs in both this section and the next targeted a focused waist of about 2.0 microns, to avoid requiring an overly large grating area. Future designs incorporating arrays of such gratings will of course have to account for the same tradeoff, and the size of the gratings will determine how closely they can be packed to address neighboring ions. For the present devices though, operating at $\theta \sim 30^\circ$ and $z = d \cos(\theta) = 50 \mu\text{m}$, the choice of $w_0 = 2.0 \mu\text{m}$ led to a waist in the waveguide plane of $6.4 \mu\text{m}$ and hence a grating dimension of $18 \mu\text{m}$.

The Si_3N_4 film thickness of 120 nm was chosen as it gives a reasonably well-confined mode (slab guided mode effective index is $n_{\text{eff}} = 1.66$ when surrounded oxide), and reasonably strong grating strengths when fully etched. Thicker films would correspond to a higher n_{eff} and lower periodicity for a given emission angle (since $k_{y,\text{em}} \sim \beta - \frac{2\pi}{\Lambda}$) which can be more difficult to fabricate, and past a point increasing thickness no longer increases grating strength, and the film thickness here was reasonably close to the maximum.

3.2.2 Longitudinal design

Along the direction of waveguide propagation (y as labeled in Fig. 3-1), the emitted field profile is tailored via the local grating period (Λ) and duty cycle (DC), which together set the local angle of emission θ and grating strength α (defined such that along the length of a uniform grating the electric field magnitude would decay as $e^{-\alpha y}$). We approximate the local θ and α as equal to those of a uniformly periodic grating with the same Λ and DC, accurate for gratings in which these parameters vary sufficiently slowly over length. To determine these parameters in the designs presented here, we use the standard paraxial-limit equations for Gaussian beam propagation to calculate the field in the waveguide plane, $E(y, z = 0) = |E(y)| e^{i\phi(y)}$, that would propagate to a focus with $w_0 = 2.0 \mu\text{m}$, $z = 50 \mu\text{m}$ above the waveguide plane and an angle $\theta = -30^\circ$. The corresponding wavenumber along y is simply

$$k_y = \frac{d\phi(y)}{dy}, \quad (3.3)$$

from which the local emission angle is calculated as

$$\theta(y) = \sin^{-1}(k_y/k_0). \quad (3.4)$$

Similarly the amplitude profile $|E(y)|$ is used to calculate the necessary $\alpha(y)$, via:

$$2\alpha(y) = K |E(y)|^2 \left(1 - \eta \int_0^y |E(y')|^2 dy' \right)^{-1}, \quad (3.5)$$

where η is the fraction of power outcoupled by the end of the grating length and K is a normalization factor that enforces the relationship between grating strength and total power emitted:

$$1 - \eta = \int_0^L \exp[-2\alpha(y)] dy, \quad (3.6)$$

with L the length of the grating.

In general, the electric field in the waveguide plane from a focus at some height h above the chip can be written as (see Fig. 3-3 for illustration) [Hau84, Ver89]:

$$E(y', z') = E_0 \frac{w_0}{w(z')} \exp \left[-\frac{y'^2}{w^2(z')} \right] \exp \left[-i \left(kz' + k \frac{y'^2}{2R(z')} - \psi(z'), \right) \right] \quad (3.7)$$

and where z' is the length along the propagation axis (y' and z' represent the coordinates in the beam's coordinate system, with origin at the focus as shown in red in Fig. 3-3); w_0 represents the beam waist at the focus (at height $z = h$); $w(z')$ is the waist along the propagation axis, as given above; and $R(z')$ and $\psi(z')$ are the evolving




Figure 3-3: Coordinate definition for calculation of desired field in grating plane; the bold line represents the grating extent (with grating length L) and the red dot represents the origin of the beam's coordinate system, at the focus.

radius of curvature and Guoy phase, given by

$$R(z') = z' \left[1 + \left(\frac{z_R}{z'} \right)^2 \right] \quad (3.8)$$

$$\psi(z') = \arctan \left(\frac{z'}{z_R} \right) \quad (3.9)$$

The field at the grating plane can be obtained by substituting, for points along the bold line in Fig. 3-3 representing the grating plane with $z = 0$:

$$y'(z = 0) = y \cos \theta \quad (3.10)$$

and

$$z'(z = 0) = \frac{h}{\cos \theta} + y \sin \theta. \quad (3.11)$$

Together these allow calculation of the necessary field $E(y, z = 0)$, from which the grating can be assembled as described below.

In the weak grating limit, the emitted angle corresponding to the phase of $E(y, z = 0)$ can be calculated simply by observing that the first order diffracted beam has wavevector along y

$$k_y = \beta - \frac{2\pi}{\Lambda}, \quad (3.12)$$

where $\beta = n_{\text{eff}} k_0$ is the guided mode's propagation constant and Λ is the periodicity of the grating. In this limit the necessary period could be calculated independently of the required grating strength, and while this intuition is helpful in designing devices,

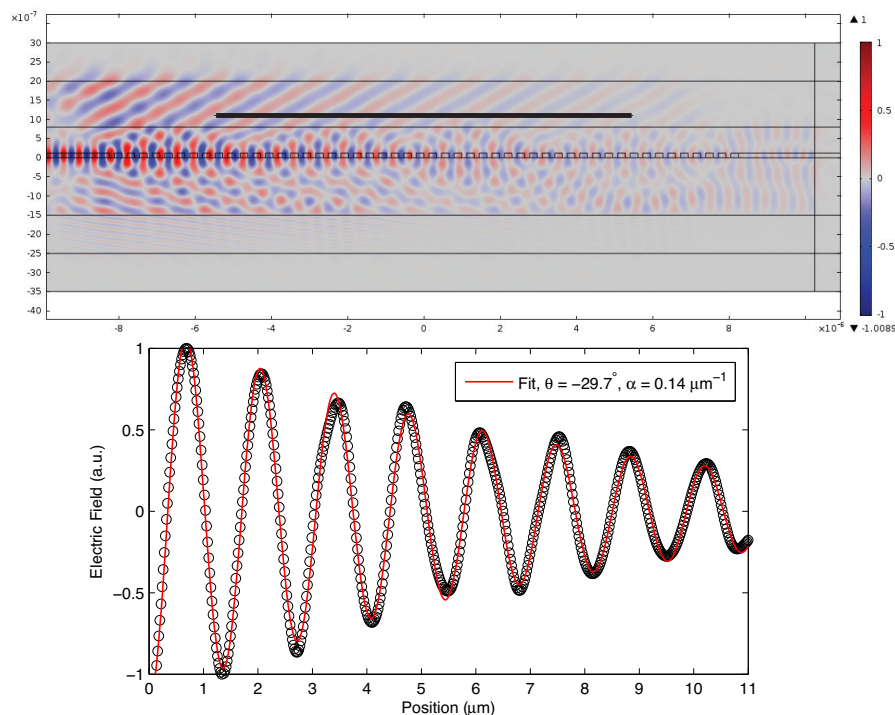


Figure 3-4: (top) example output of COMSOL simulation for grating 2D cross section; and (bottom) calculated electric field along the dark black line, showing the fit from which the emission angle and grating strength are inferred.

the weak grating approximation does not accurately describe the $\text{Si}_3\text{N}_4/\text{SiO}_2$ gratings considered here, and our designs were performed accounting for both Λ and DCs influence on both θ and α .

To relate the required $\alpha(y)$ and $\theta(y)$ to the physical grating parameters $\Lambda(y)$ and $\text{DC}(y)$ (which we define here as the fraction of a grating period where the Si_3N_4 is etched away and occupied by the low-index SiO_2), frequency domain, finite-element-method simulations of uniform periodic 2D grating cross sections were carried out in COMSOL, from which the decay lengths (giving α) and emission angles are calculated as a function of Λ and DC . In these simulations, the input is in the form of the calculated mode profile of the uniform waveguide, and the other boundaries are surrounded by perfectly matched layers, which function as absorbing layers. Similar simulations could also be carried out via frequency difference time domain FDTD methods, for example in MEEP [ORI⁺10], but since we were interested in single frequency response in this case and the 2D geometry allows a simulation with small enough domain, the frequency domain calculation can be faster. An example simulation domain and field profile, together with the fit used to extract α and θ for this particular period and duty cycle, is shown in Fig. 3-4.

The results of such calculations over a range of grating parameters are shown in Fig. 3-5, together with the physical grating parameters assembled from such calculations to result in focusing along both x and y to an approximately $2 \mu\text{m}$ spot $50 \mu\text{m}$ above the chip, and at an average angle in the yz plane of -30° . We choose $\theta < 0$,

corresponding to emission with direction along y opposite that of the guided mode, so as to ensure no second diffraction order; this reverse emission also turns out to be essential for focusing given the the method used to define the grating arc radii, as discussed below. The simulated efficiency of these devices (calculated as the upwards-radiated power divided by the incident) is 80%, taking advantage of the Si substrate as a reflector of downwards-radiated light (Fig. 3-1a) and using angles where, given the bottom oxide thickness here, constructive interference maximizes the grating strength (Fig. 3-5a and b).

The design represented in Fig. 3-5 is for a coupler that focuses along y ; the same procedure could have been used to assemble, from the data in Fig. 3-5a and b, a coupler emitting a collimated beam along y by simply choosing the corresponding $E(y, z = 0)$ (this was done for the trap-integrated devices shown in the next chapter).

We note that although a few previous designs have employed holographic methods to find the grating line spacings, these implicitly assume a low index contrast; the method we have adopted here, particularly for the longitudinal design parameters (and in a fashion related to work on silicon photonic grating couplers to SM fibers[TBB04, NPW+16]), is directly applicable to high index-contrast structures.

3.2.3 Transverse design

Transverse focusing (along x as labeled in Fig. 3-1) is controlled by the curvature of the grating arcs, which relates to the degree to which the phase of the radiated beam at the outer edges of the grating in the transverse direction (at $\pm x$) is advanced relative to at center, and hence the strength of the focusing action.

To minimize distortion of the field profile as it propagates through the grating region, the gratings presented here are designed such that the radius of curvature of each grating arc is equal to the distance from the start of the taper; since the guided field expands through the taper such that the radius of curvature as a function of distance from the taper start is approximately equal to that distance, this condition approximately ensures that each grating arc is parallel to the phase front incident on it (or perpendicular to the effective rays propagating through the structure as illustrated in Fig. 3-1c). Though this is not required for focusing action in general, as shown below, choosing the curvature this way reduces the strength of the low-intensity sidelobes away from the focus as compared to devices in which this constraint on the radii was not imposed [MBM+16]. The observations on various previous devices that led to the imposition of this constraint are described in Appendix B.

We approximately predict the height of the focus based on the radius of the arc at the center of the grating longitudinally (which we call R_g), where the emission amplitude is engineered to be maximum (Fig. 3-5); the radius of curvature of the phase fronts emitted, along the dimension of propagation, are expected to be roughly $R_i = R_g / \sin(\theta)$. This, together with the standard equation for evolution of radius of curvature for a Gaussian beam as a function of distance from waist along the propagation direction d

$$R(d) = d \left[1 + \left(\frac{z_R}{d} \right)^2 \right], \quad (3.13)$$

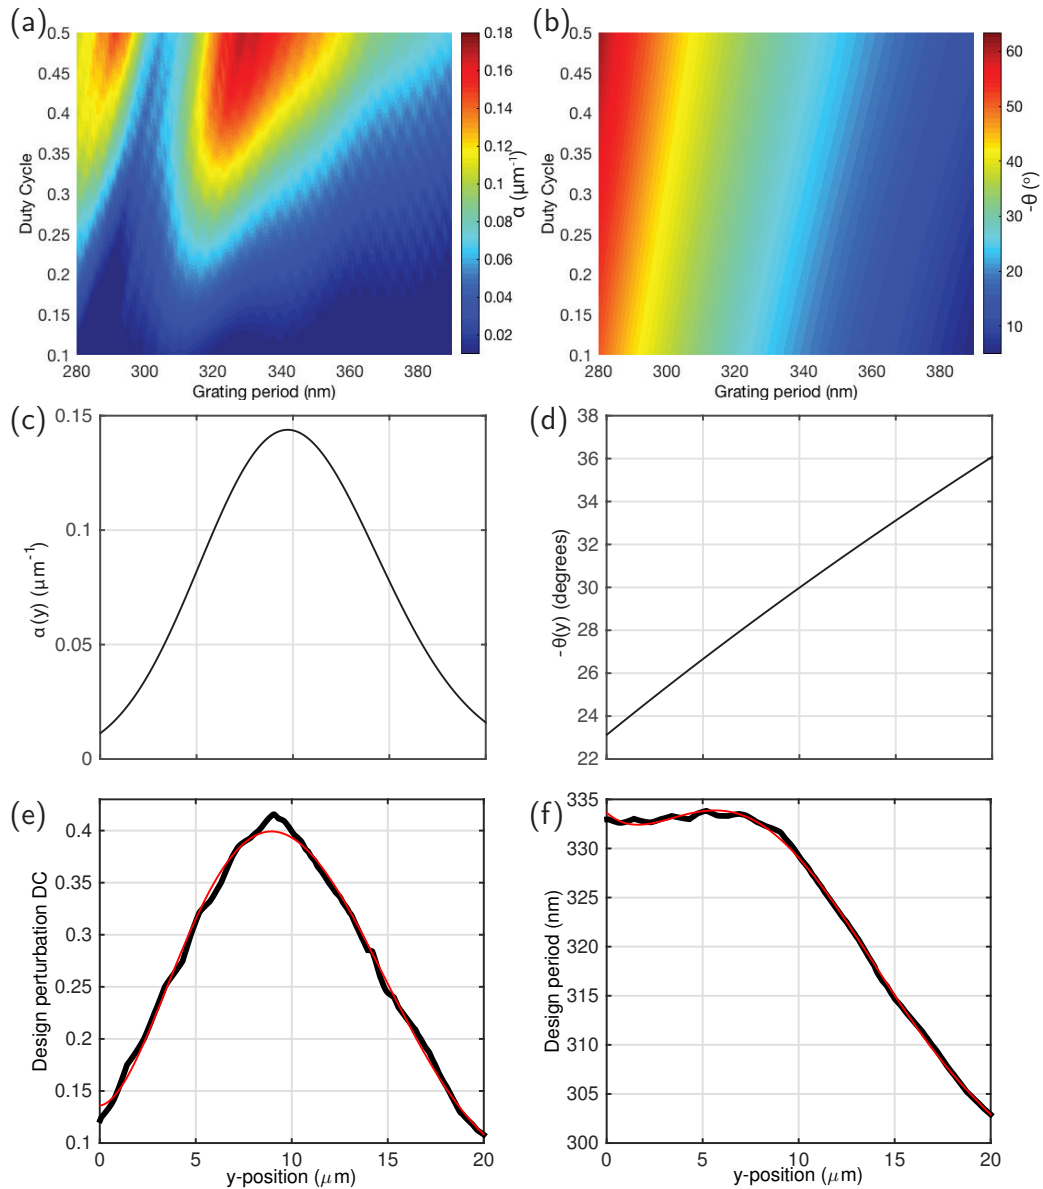


Figure 3-5: Grating design parameters. (a) Simulated α and (b) θ as a function of grating period and DC; (c) Desired local α and (d) local θ to produce the intended focus for a grating $20 \mu\text{m}$ -long along y ; (e) and (f) inferred physical DC and Λ profiles to approximate the desired grating (black lines), together with polynomial fits used to specify the design.

gives the necessary radius of curvature of the grating lines to satisfy this condition.

In the transverse direction (along x), the amplitude is given by the mode profile of the wide waveguide mode, a cosine function to a good approximation:

$$E_c(x) = \sqrt{\frac{2}{w}} \cos(x\pi/w), \quad (3.14)$$

where normalization is chosen such that $\int_{-w/2}^{w/2} E^2 dx = 1$.

Ideally, this would be a Gaussian, with desired initial beam waist w_i (i as opposed to the focused beam waist w_0):

$$E_g(x) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \frac{1}{\sqrt{w_i}} \exp\left(-\frac{x^2}{w_i^2}\right), \quad (3.15)$$

with the same normalization.

From the focused beam waist and distance to waveguide plane we know what the desired w_i is; the question is just how to choose the physical w to maximize the overlap

$$\kappa_{c,g} = \sqrt{2} \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \int_{-w/2}^{w/2} \frac{1}{\sqrt{ww_i}} \cos(x\pi/w) \exp\left(-\frac{x^2}{w_i^2}\right) dx, \quad (3.16)$$

and this gives $w \approx 2.844w_i$, at which value the integral evaluates to $\kappa_{c,g} \approx 0.9947$ implying that the cosine would couple to a Gaussian of width w_i with $|\kappa_{c,g}|^2 \approx 98.9\%$ power efficiency. We therefore expect this Gaussian approximation to the waveguide mode to predict the focusing properties of the main lobe of the beam, although the deviations from an ideal Gaussian (at the $\sim 1\%$ power level, and as illustrated in Fig. 3-6) will lead to propagation of higher-order Hermite-Gaussian modes as well, with implications for crosstalk away from the main lobe which will be discussed later.

3.2.4 Predicted focus distance and waist for a given design

For reference, and since these equations are helpful in design, we briefly note that if the initial curvature R_i and beam waist (from waveguide width) w_i are known, the expected focused beam waist is

$$w_0 = \frac{R_i w_i \lambda}{\sqrt{\pi^2 w_i^4 + R_i^2 \lambda^2}} \quad (3.17)$$

and the distance to the focus is

$$d_f = \frac{\pi^2 R_i w_i^4}{\pi^2 w_i^4 + R_i^2 \lambda^2}. \quad (3.18)$$

These follow from a simultaneous solution of the equations for the Gaussian beam phase curvature radius $R(d)$ and waist $w(d)$ written above.




Figure 3-6: Cosine-profile waveguide mode field (solid line) and maximally overlapping Gaussian profile (dotted line); position is in units of waveguide width, and the Gaussian has $1/e^2$ width of $1/2.844$.

3.3 Layout, fabrication, and material characterization

The grating and waveguide patterns are all generated by scripts that encode the designed geometry, in SKILL code within Cadence Virtuoso, for convenience using simple approximations to the curved shapes in the form of rectangles of 5 nm width – an example of such a discretization is shown in Figure 3-7, and is the same general approach used for CMOS photonics by our group previously [OR10].

Devices are fabricated starting with silicon wafers coated with 1.5 μm of thermal oxide, followed by 120 nm of stoichiometric, LPCVD Si_3N_4 (details on deposition in the section below). Electron-beam lithography is performed with a system operating with electrons accelerated to 125 keV energies (Elionix ELS-F125) and using HSQ resist developed with a mixture of NaCl and NaOH [YB07]. The develop step was done in this solution at room temperature for 4 minutes, by simply immersing the piece in a beaker containing the developer solution. Reactive ion etching is performed with CHF_3 and O_2 gases [HBP⁺06] using the RIE chamber in the nanostructures laboratory (NSL) at MIT. The CHF_3 and O_2 gases are flowed in at rates of 4.0 and 15.0 sccm, respectively, at a chamber pressure of 10 mTorr. This etch is followed by PECVD cladding deposition of SiO_2 using TEOS precursor with the Oxford-100 PECVD tool in ICL, using the recipe for “0-stress” TEOS CVD.

Using the Virtuoso software a gds file containing the desired pattern was generated, after which LayoutBeamer was used to translate into a format readable by the ebeam tool, being namely a series of points over which to scan the beam. The e-beam current used for exposure was typically 5 nA (chosen to be high enough to write full chip patterns in a few hours without being too large as to sacrifice resolution), and




Figure 3-7: Mask layout for grating device, with inset showing discretization on a 5 nm rectangular grid.

written patterns were discretized (in LayoutBeamer) to squares 2.5 nm on a side. For the beam current I_{eb} , dose time τ , and pitch $a = 2.5$ nm in this case, the delivered dose in Coulombs/cm² is calculated as $I_{eb}\tau/a^2$.

The range of the 125 keV electrons in silicon is approximately 50 μm , which is to say this is the depth down to which the injected electrons may propagate. Some of the electrons scattered from within the substrate are reflected back to the resist, where they may expose resist within a radius on the order of the range from the point at which the beam is pointed, and this exposure from backscattered electrons is responsible for the ‘‘proximity effect’’ in electron beam lithography. As a result the substrate, as well as the characteristic dimensions and pattern density of the features being written, alters the optimal exposure dose, and exposures for different scale features must be independently checked. In this case, we used a single dose over the focusing grating area, a single dose for the fiber couplers, and a single dose for the SM waveguides. With HSQ these dosages can shift somewhat over the shelf life of the resist, but here the dwell times per 2.5×2.5 nm² dot were, respectively, 0.14, 0.22, and 0.20 μs , corresponding to dosages of 11.2, 17.6, and 16 mC/cm², respectively. Underexposed features were observed to result in grating lines that would lift off from the substrate and appear wavy or curved when imaged (for example in Fig. 3-8, and overexposed grating features translated into gaps that were partially or completely filled. Dose calibrations were performed by writing a version of a desired device using a series of different dwell times, incremented by 0.02 μs (this was due to constraints on the dose time, requiring that it be a multiple of 0.02 μs , in the system configuration when these writes were performed – later experiments, presented in Ch. 6, used more fine dose stepping with a constant multiplication factor between dose tests). For single-mode waveguides, approximately optimal dosage can be estimated by checking at which dosages the waveguide width realized is approximately as designed (underexposure typically results in eating away of the sidewalls, and vice versa – often also additional roughness in the resist), and grating dosages can be estimated by avoiding the evident underexposure as visible in Fig. 3-8 and avoiding gap filling in, evident under SEM. Highly overexposed gratings can also be identified by optical microscopy, where typically the grating region appeared significantly darker than the wide taper region, but this contrast is not seen when the lines are significantly filled in. The optimal doses for HSQ could change noticeably (on the 0.1-0.2 μs dwell times at least) over months, so these calibrations had to be repeated every 2-3 months, roughly, though how often would depend on exactly what features are being written.

We note additionally that the particular order in which shapes in the design were exposed by the e-beam tool could significantly affect the intensities in the weak sidelobes; the results shown below were from devices where care was taken to ensure the pattern was exposed symmetrically, and a comparison with observations from other devices is presented Appendix B.

3.3.1 Silicon nitride deposition and loss

Many of the gratings were fabricated in LPCVD-grown stoichiometric Si₃N₄ deposited at 720°C on silicon wafers with a 1.5 μm thermal oxide; the wafers with thermal oxide




Figure 3-8: SEM image of slightly underexposed grating feature, showing lines lifting off (obtained using a $0.28 \mu\text{s}$ dwell time with a 2 nA beam current, equivalent to $0.11 \mu\text{s}$ with 5 nA for comparison to the numbers given in the text)

were supplied by Rogue Valley Microdevices, and the LPCVD material was deposited in the tool in ICL at MTL. The guided mode in the film can be measured before waveguide patterning to assess material quality using a prism coupling system made by Metricon; here a HeNe laser beam (632 nm) is shone in through a prism whose surface is in contact (to within $\lambda/2\pi$) of the nitride. The angle of the input is swept through a range where the beam totally internally reflects within the prism, and at a particular angle when the light within the prism has a wavevector parallel to the substrate surface equal to the guided mode's, the light is coupled in and a dip in the transmission is observed, as in Fig. 3-10. The loss can be measured in this system with a fiber detector that scans along the path of the launched light and measures scattering out of the mode, fitting the profile to a decaying exponential.

The unpatterned film material on silicon substrates had low loss ($< 0.4 \text{ dB/cm}$ as measured with the prism coupler at 632 nm) and was ideal for tests of designs.

Silicon substrates, especially doped but also undoped [NLK⁺14], present difficulties for the RF fields used in ion traps, so the experiments towards integration with the trap structure had to make use of either a substrate transfer or a different substrate with low RF loss, such as crystalline quartz, sapphire, or fused silica.

3.4 Optical characterization

The grating emission is characterized by imaging the emission in a microscope using $50\times$ objective with a 0.95 NA (Olympus MPLAPON50x, in an Edmund Infinitube Proximity Series microscope tube assembly). This NA implies an acceptance cone half-angle of 72° , large enough to ensure the emission of the couplers is collected. A





Figure 3-9: Image of prism-coupling into guided film mode on a 6 inch silicon wafer, and trace of detected (reflected) intensity as a function of angle, plotted against the effective index corresponding to the physical angle.




Figure 3-10: Fiber coupler used for input to waveguides; the width of the grating region is $6 \mu\text{m}$.

series of images is taken scanning the focal plane of the imaging system up from the waveguide layer, and the resulting stacks of images are integrated along x or y to yield intensity profiles along y and x , respectively, similar in principle to a “knife-edge” measurement at each height. The resulting profiles are shown in Fig. 3-11, showing focusing behavior along both dimensions, and an average emission angle of $\theta \approx -27^\circ$. RF noise from an attenuated noise source (Toptica Laser Coherence Controller) was fed to the modulation input of the laser the light input to the gratings for these measurements, with the purpose of reducing the coherence length to order ~ 1 cm, to eliminate interference artifacts that otherwise arose in the imaging system (as are visible in the next chapter in Fig. 4-2).

By collecting the emitted beam on a photodiode and comparing to the input power, and normalizing for the loss of the input coupler and waveguide feeding the focuser, we estimate the physically realized efficiency of radiation into the focused beam to be $70 \pm 15\%$ (with uncertainty due to variation in total waveguide transmission on this sample), in reasonable agreement with simulation. Here the input grating coupler (to the SMF) loss was measured to be 7.3 dB (reproducible from device to device within 0.1 dB), the average total loss from the 1.8 cm-long waveguide was 13.6 dB, and the average loss from the focusing gratings was 1.6 dB as measured on the free-space power head.

3.4.1 Detailed characterization at focus

The spot was characterized in detail at the designed height of $z = 50 \mu\text{m}$. The inset in Fig. 3-12 shows the measured intensity profile here, together with intensity profiles along x and y . The Gaussian fits to the main lobes (shown in grey dotted lines) indicate a waist of $w = 2.0 \mu\text{m}$ along x ; along y the Gaussian fit has a $2.3 \mu\text{m}$ $1/e^2$ half-width, which corresponds also to $w = 2.0 \mu\text{m}$ after accounting for the propagation




Figure 3-11: Measured “knife-edge”-like beam profiles (a) along y and (b) x showing focusing behavior along both dimensions.

along this direction. These fits indicate the device focuses approximately as designed along both dimensions; although the minimum averaged waists are realized actually at about $54 \mu\text{m}$ along y and $40 \mu\text{m}$ along x , these offsets are within a Rayleigh range of $50 \mu\text{m}$ and the difference in beam waist with respect to that at $50 \mu\text{m}$ is small.

The intensity profiles plotted in Fig. 3-12 result from a series of images, taken with a QColor 5 scientific CCD sensor (only the red color channel data was used), with exposure times varying by a factor of 400, and with dark frames subtracted (at each exposure length, 10 frames were taken with the grating illuminated and averaged, from which 10 frames with the laser blocked were averaged and subtracted), to allow sufficient dynamic range to resolve the intensity up to $\pm 25 \mu\text{m}$ from the center. Along the transverse direction (x , along which focusing is controlled by the grating line curvature), we plot this data together with the result of a 1D diffraction integral calculation showing the expected profile at this height accounting for the effect of the finite “aperture” corresponding to the finite grating width. Since a wide waveguide’s fundamental mode profile approximates a cosine profile in the core, we calculate the diffraction from a cosine profile with zeros at $\pm w = 9 \mu\text{m}$, corresponding to the diffraction from the center of the grating region where the emitted intensity is designed to be maximum. The resulting profile, calculated from the 1D diffraction integral

$$E(x, z = 50\mu\text{m}) = \int_{-w}^w dx' E(x', z = 0) \frac{e^{ik_0 d}}{\sqrt{d}}, \quad (3.19)$$

with $d = \sqrt{(x - x')^2 + \left(\frac{h}{\cos\theta}\right)^2}$, and where $E(x', z = 0)$ is the cosine-profile field amplitude along x in the waveguide plane with appropriate radius of phase curvature, is plotted in the red line in Fig. 3-12(a), and the close correspondence of this envelope with the measured points indicates that, along x , the profile even in the low-intensity sidelobes is very nearly diffraction-limited.

This is a significant improvement in sidelobe suppression over the performance of the device previously presented [MBM+16], which is due to the condition imposed here on the radius of curvature as described above, which minimizes distortions of the transverse profile of the guided field propagating through the grating region. Along the longitudinal direction, the emitted field profile is controlled by the period and duty cycle of the grating and the low-intensity sidelobes are not as well suppressed, but we still observe values below 10^{-3} beyond $10 \mu\text{m}$ from the focus. Further optimization of these designs may allow improvement beyond the mode purity achieved here, or minimizing intensities at particular distances from the center. However, we expect these designs may already be applicable with advantages in performance, as for typical ion experiments a high degree of control over the sidelobes is necessary only along one dimension (the trap axis), and along x the profile here is already a significant improvement over what has been achieved in ion experiments [SNM+13, DLF+16]. That a straightforward, intuitive design method achieves this performance along x may be a significant aid to practical design of experiments.

We note again that the sidelobe-profile presented here was from devices in which the e-beam exposure occurred in a symmetric fashion; the sidelobe profiles of devices written in a less ideal order are shown in the Appendix.

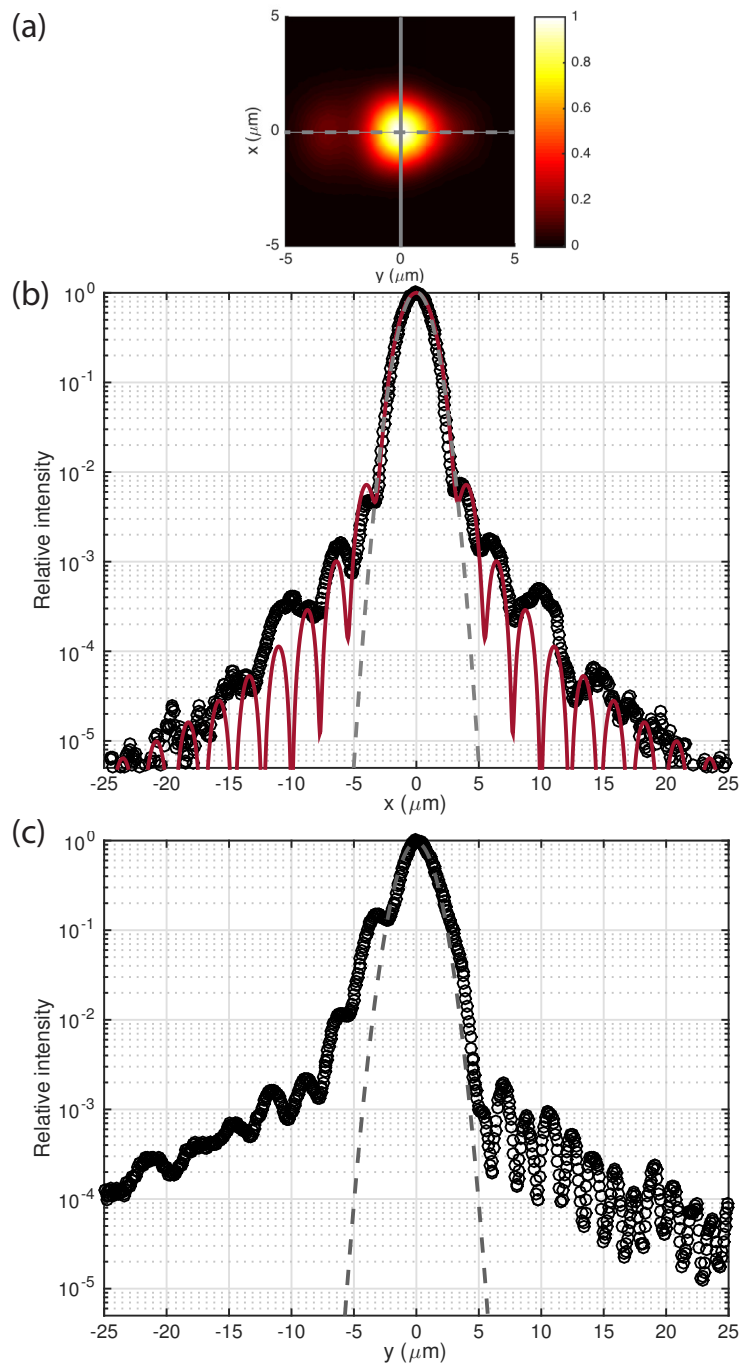


Figure 3-12: Intensity profiles along x (a) and y (b) imaged at a height of $z = 50 \mu\text{m}$. The inset shows the recorded intensity profile at this height, with the solid and dotted lines corresponding to the horizontal axes for (a) and (b) respectively. In each case the measured data points (black circles) are taken from a set of images with exposure times varying by a factor of 400 to allow sufficient dynamic range. Ideal Gaussian fits with $1/e^2$ half-widths of $2.0 \mu\text{m}$ (a) and $2.3 \mu\text{m}$ (b) are shown in the dotted gray lines, as well as in (a) the result of a 1D diffraction integral calculation (solid red line) for the intensity profile resulting from the cosine-shape field profile expected along x in the grating region.

Designs with higher effective NA, achieved by either reducing the focus height or increasing the emitting area, should result in tighter focuses; the present devices are not yet at a limit set by total internal reflection at the oxide-air interface, which would allow w_0 well below $1 \mu\text{m}$. For tighter focuses requiring smaller R_g , for a given grating waveguide width (proportional to emitting aperture diameter) the constraint on curvature radius here may not be practically applicable (i.e. it may result in tapers expanding at a greater angle than the divergence angle corresponding to the the SM waveguide mode), and in these cases the desired focusing behavior may be achieved at a trade-off with sidelobe suppression. We note also that the constraint on the radius here produces a focusing beam only when $\theta < 0$; in the opposite case (perhaps more easily fabricated in some cases since forward emission corresponds to a larger grating period) the constant phase surfaces of the profile expanding through the taper coinciding with the grating arcs would correspond to a diverging radiated beam. In that case the arc radii can modified to still produce focusing for $\theta > 0$, but also with a tradeoff in sidelobe suppression unless otherwise compensated.

3.4.2 Polarization purity

Finally, we characterize the polarization purity in the focal region. Owing to the dominant polarization of the mode feeding the taper, the radiated field is expected to be polarized predominantly along x ; furthermore, the SM waveguide mode has a dominant x component that is even about the yz -plane, with smaller y - and z -directed fields which are odd about this plane. Owing to the symmetry of the taper and grating about the yz -plane this symmetry is preserved as the field propagates through the structure (effective rays propagating through the structure illustrated in Fig. 3-1(c), with accompanying E-field, showing the odd symmetry in the y -directed components), and hence at the center of the radiated beam in the yz plane the components other than along x should be zero.

A rotating polarizer (Thorlabs LPVISE100-A in a rotating mount) inserted in to the microscope allows us to image only the light with polarization along x , or that along the other orthogonal component also transverse to the propagation direction (primarily along y). Images obtained at $z = 50 \mu\text{m}$ with the polarizer oriented along x and y are shown in Figs. 3-13(a) and (b), with a trace along the $y = 0$ axis in Fig. 3-13(c). The x -polarized profile closely reproduces the patterns obtained with no polarizer inserted, and the null in the y -polarized light at $x = 0$, as well as the side-lobes near $\pm 2 \mu\text{m}$ owing to the weak y -directed field components in the grating region, are consistent with the argument above. We measure a minimum of $< 3 \times 10^{-4}$ in relative intensity, likely limited by the extinction of the polarizer used here ($\sim 1 \times 10^{-4}$), and the birefringence of the microscope objective, not a low-stress objective optimized for polarization microscopy.

These observations indicate that these devices can produce beams with a high degree of polarization purity at the center of the focus. We note that we have imaged in the far-field the intensity in the two components transverse to the propagation direction, and our measurement is not sensitive to the longitudinal components that generally arise locally in the focal region when beams are tightly focused [TTZ⁺13];




Figure 3-13: Measured intensity profiles when imaging (a) only the dominant x -oriented polarization and (b) the orthogonal transverse polarization (with the polarizer oriented along y); color bars are scaled differently for each plot but correspond to the same scale. (c) Cross section along $y = 0$, showing intensity (relative to the peak of the x -polarized intensity) in each component along x ; black circles are points measured with the polarizer oriented along x , and red those with the polarizer oriented along y .

measurement of relative excitation rates on transitions involving different sublevels in an atom or an ion moved through the focal region could allow precise probing of the polarization profile in all three dimensions.

3.5 Discussion

Previous work has shown that photolithography and, more specifically, full CMOS processes can be leveraged to produce photonic structures like those presented here [DMR⁺04, OMS⁺12], often benefitting from optical proximity correction techniques for fine features [MOTZ⁺14]; the dimensions in the devices here should be achievable with the photolithography used for current 14-nm processes. Hence, in a slightly customized process with a patternable layer suitable for visible-wavelength waveguides (like the Si₃N₄ used here), it should be possible to integrate such devices on silicon substrates with multi-layer CMOS ion traps [MEB⁺14] for large-scale QIP systems based on such devices, or perhaps with CMOS photodiodes for wide-field microscopy.

The precision with which the transverse profile is formed here is comparable to that demonstrated with assemblies based on digital micromirror device arrays for optical lattice experiments [ZPM⁺16], and should be generally useful for highly precise definition of static optical potentials from compact and scalable devices, and without the need for additional high-NA bulk optics. Further extensions may include generating circular polarizations using either two separate couplers or ideas similar to those used in polarization-splitting couplers [MGM⁺11], as well as more complex optical profiles; for example, Hermite-Gaussian beams could be obtained along either dimension by feeding the taper and grating with a higher-order waveguide mode, or shaping the longitudinal grating profile correspondingly. In general these results demonstrate the possibility for high index-contrast waveguide devices to produce precisely tailored and tightly focused beams near a chip surface, using an intuitive and relatively quick design approach, and in a fashion that should be scalable to complex geometries.

3.5.1 Possible improvements to sidelobe suppression in present devices

A notable feature of the devices presented here is the nearly diffraction-limited sidelobe intensities achieved along x ; this indicates the possibility that such devices can play a role in systems requiring very low crosstalk to areas of space at some distance from the main focus. One such example is of course ion trap QIP, and such devices' performance, and comparison to existing experiments relying on such addressing in trapped-ion settings, is discussed in the next chapter.

Here we briefly mention possible routes to further improving the sidelobe suppression in these devices. The grating area can be increased, using the same constraints on radius of curvature, as long as the taper angle does not exceed the Gaussian beam




Figure 3-14: Straightforward possible improvement to transverse focal width and sidelobe intensity, by increasing waveguide width from 18 μm to 30 μm .

divergence angle corresponding to the waveguide mode feeding the taper

$$\theta_d \approx \frac{\lambda}{n_{\text{eff}} \pi w_g}. \quad (3.20)$$

where w_g represents the effective waist of the guided mode. This is approximately 300 nm for the waveguides here, corresponding to $\theta_g \approx 25^\circ$, in excess of the current taper angle of 15° . This suggests that for the same focal height as was used here, a larger taper angle could be employed and hence a larger average emitting area in the transverse direction by a factor of $\tan(25^\circ)/\tan(15^\circ) \approx 1.7$; the impact on sidelobe intensities, expected from the 1D diffraction integral calculation described above, is shown in the diffraction calculation in Fig. 3-14.

Depending on the application the tradeoff in device area may not be favorable, but the fractional reduction in sidelobe intensity at a given distance is in excess of the increase in dimensions.

Finally, though we have not pursued this in our present work, which has focused more on physically motivated and intuitive designs that allow fairly straightforward extension to other wavelengths or material environments, for some applications it may be interesting to modify the profile of the diffracted light so as to suppress it at specific distances and enhance it at others; in ion addressing such capability would be desirable, since one would know in a designed system at which positions other ions would be. This may be possible by tailoring the strength of the grating in the

transverse dimension (by varying the DC as a function of position along x) so as to shape the transverse mode profile of the grating emission. A number of numerical and inverse design optimization procedures may play a role in such work.

Nevertheless, the devices as demonstrated above already exceed the performance of existing trapped-ion experiments, as is discussed in more detail in the next chapter when crosstalk errors are experimentally characterized. Especially given the potential for improvement through tradeoffs with device area, or optimization at particular locations, we expect such devices should allow significant advantages in this regard.

3.6 Conclusion and future work

The focusing grating devices here have many similar features to grating devices designed in the context of silicon photonics (e.g. [MGM⁺11]), but represent the first time the side-lobes and polarization purity have been managed and characterized to the degree shown here, and indicate for the first time the potential for integrated optics to generate precisely tailored beams propagating in free space.

The devices discussed here are similar to those that were employed for scalable trapped-ion qubit addressing [MBM⁺16]. Such devices' integration with planar ion traps and their use in single ion addressing is presented in the following chapter.

Chapter 4

Ion addressing with waveguide optics

In this chapter we describe our experiments utilizing devices such as those discussed in the previous chapter, integrated within planar ion trap chips, for optical ion addressing with tightly focused beams. Much of the material here is based on the material in [MBM⁺16]. We first briefly summarize the motivations for the use of such optics in ion trap systems (discussed in more detail above in Ch. 1); describe then the design of the ion-trap integrated waveguide device; and then the optical and single-ion experiments.

4.1 Overview

We first briefly summarize the motivations for this work, as discussed in more detail primarily in Ch. 2. As compared to previous work concerning integrating optics, the gratings used here are compact compared to the optical fibers and Fresnel lenses previously integrated with planar traps for addressing [KHC11] and fluorescence collection [VCA⁺10, SNJ⁺11] (cross-sections $\geq 100 \mu\text{m}$ in diameter), and most importantly the planar fabrication used here to define the optics for both routing and addressing lends itself to intimate integration with the planar trap electrodes. Furthermore, such waveguide systems have been demonstrated to be scalable to complex geometries of thousands of devices or more [STY⁺13]. Though micro-electro-mechanical systems (MEMS) mirrors integrated with traps have been proposed as well [KK09], experiments so far have utilized MEMS components external to the vacuum chamber and separate from the chip [CMBK14], leaving full integration an essential outstanding challenge.

Integrated waveguide devices bring several advantages for ion addressing in planar traps. The ability to fabricate, in the same lithographically defined waveguide layer, multiple splitters, waveguide crossings and bends with radii less than $10 \mu\text{m}$, would enable the realization of a variety of trapped ion architectures, with flexibility as to arrangement of qubits [KMW02, CLJ08], and with light delivered in parallel to each site. This parallelism will be essential in large-scale systems in which speed is at a

premium due to finite coherence times. Additionally, grating couplers near the ions can focus light to μm -scale spots, allowing quantum logic gates of a given interaction time using 2-3 orders of magnitude less power when compared to geometries with beams propagating parallel to the chip surface, in which the beam waists are typically limited by diffraction and beam-clipping concerns to 30 – 50 μm diameters [KK09]. This focusing is crucial also for general individual addressing in an ensemble of closely-spaced ions [SNM⁺13]. In addition, the phase stability of waveguide approaches even for complex optical paths [PCR⁺08] will benefit qubit operations, which are generally phase-sensitive. Furthermore, definition of optics within the trap chip essentially eliminates beam pointing instabilities at the ion location as a noise source [BWC⁺11, SNM⁺13]. Beyond trapped ion QIP, integrated parallel distribution and focusing of light near a chip surface may find further application in atomic physics, such as in ion clocks or neutral atom dipole trap arrays, and more broadly in the various applications of nanophotonic systems.

4.2 Integrated waveguide trap design and fabrication

4.2.1 Choice of substrate

So that the chip temperature remains low despite dissipation of the RF fields applied to the trap electrodes, a material with high cryogenic thermal conductivity is ideal for the substrate. Silicon is a good candidate, but we aimed for maximum simplicity in fabrication, and silicon traps are difficult to operate with a single metal layer (i.e. without a ground plane to isolate the RF from substrate; see chapter below on CMOS traps) due both to the RF properties of silicon, and the possibility for photogenerated charge carriers in silicon to result in destabilizing stray fields.

The two prime candidates other than Si are sapphire, which the Lincoln group has used for planar traps, and crystalline quartz. Crystalline quartz is preferable due to its lower refractive index, allowing waveguides to be fabricated directly on the substrate without an intermediate thick buffer oxide, which would have to be deposited with PECVD and then chemical-mechanical-polished (CMP) to have sufficient thickness and smoothness to make waveguides on sapphire (Quartz glass or fused silica, as opposed to the crystalline form used here, has an index orders of magnitude lower at 4 K [lak]).

A consequence of this substrate choice is that the reflection off of the silicon cannot be used to enhance grating efficiency, as was done in the previous chapter. Hence these devices had lower total efficiencies, both since they emitted approximately equally along the $\pm z$ directions, and because the lack of constructive interference from the two paths results in a lower grating strength (and less total light out-coupled over the course of the grating).

4.2.2 PECVD and LPCVD SiN deposition

LPCVD deposition on crystalline quartz failed in the first attempt because the temperature ramp to 720° was done too rapidly (over 3 hours), which resulted in all the wafers shattering, perhaps due to the $\alpha - \beta$ inversion in crystal quartz, at 573° C (although Quartz is in general brittle and overly fast temperature ramps even up to 400° were observed also to result in cracks). This issue has been encountered in other high-temperature processes using crystalline quartz [PLM⁺09]. A subsequent attempt with a significantly slower ramp avoided this problem, but we nevertheless observed high loss (8 dB/cm, via prism-coupling measurements at 633 nm) in films formed as such. The reason for this relatively high loss is unclear, but may be related to the high tensile stress that LPCVD nitride grows with, which together with the crystalline Quartz substrate may result in defects not seen when the nitride grows on amorphous SiO₂ (as in the case of thermal oxide on silicon).

In any case, the reason for this behavior is not understood as of this writing, and as it happened the LPCVD chamber developed a leak soon after the observation of the high loss on crystalline quartz, preventing investigation of this for some time. Instead, we opted to use plasma-enhanced chemical vapor deposition (PECVD) to form the nitride films, which occurs at lower temperatures (typically between 300 and 400 °C), and can be more quickly performed on numerous separate samples.

The Oxford-100 PECVD tool in ICL was used to deposit Si₃N₄ films on Quartz, and no problem with shattering was encountered in deposition at 300° . For the devices presented below, the standard SiN recipe on the tool was used, which involves both low frequency and high frequency RF deposition, aimed at minimizing stress in the nitride. Such films on quartz were observed to have losses of approximately 6 dB/cm. However, in later experiments after the device presented below was fabricated, we found that using low frequency deposition alone could reduce the film loss by a few dB/cm, consistent with observations in previous studies [GJG⁺08]. The standard recipe had a rate gas flow ratio of 1 : 2 (NH₂ : SiH₄), and furthermore, reducing the relative flow rate of the silane resulted in a slightly lower index film, but lower loss, as summarized in the measurements in Table 4.1; here, the Metricon measurements on Si/SiO₂ samples refers to samples on silicon with the 1.5 μ m thermal oxide, as compared to directly on crystal quartz (right-most column). In all cases, the loss on crystal quartz was observed to be higher than on thermal oxide. Nevertheless, this loss is tunable via the deposition conditions. We note though that although LF deposition was reproducibly lower loss than HF/LF depositions, and that crystalline quartz substrates did typically result in higher losses than Si/thermal oxide, losses obtained on different crystal quartz wafers could vary by as much as a few dB/cm, for reasons that remain unclear.

Though routes to lower loss PECVD films are hence possible, all subsequent discussion in the present chapter is on the device that was made before the above measurements were carried out, with the standard PECVD SiN (using both LF/HF deposition and with a film loss of approximately 6 dB/cm – the same gas flow ratio as the middle column of Table 4.1) directly on crystalline quartz.

NH ₃ : SiH ₄ (sccm)	Reflectometry (monitor)	Metricon (Si/SiO ₂)	Metricon (quartz)
20 : 20	178 nm thickness, n(674) = 1.943	$n_{\text{eff}} = 1.679$, < 0.5 dB/cm	$n_{\text{eff}} = 1.664$, ~1.8 dB/cm
13 : 26	186 nm thickness, n(674) = 1.988	$n_{\text{eff}} = 1.727$, ~1.9 dB/cm	$n_{\text{eff}} = 1.694$, ~5.8 dB/cm
40 : 20	166 nm thickness, n(674) = 1.900	$n_{\text{eff}} = 1.637$, ~1.0 dB/cm	$n_{\text{eff}} = 1.634$, ~2.6 dB/cm

Table 4.1: Observations on film and waveguide properties of PECVD SiN deposited under various conditions on both Si/SiO₂ samples and crystal quartz. These depositions were done using only LF deposition.

4.2.3 Chip layout and design

The trap electrodes and waveguide patterns in the device presented here are visible in the optical micrograph of Fig. 4-1(a). Waveguides were fabricated on a crystalline quartz substrate in a silicon nitride (SiN) film (with refractive index $n \approx 2.0$), with cross-sectional areas of approximately 120 nm \times 540 nm, single-moded for the quasi-TE polarization at $\lambda_0 = 674$ nm. These waveguides route light on chip without phase-front distortions or diffraction. Approximately 1 μm of SiO₂ forms the top cladding, above which sit niobium (Nb) trap electrodes; the resulting cross section together with a simulated guided mode E-field profile is illustrated in Fig. 4-1(b). This guided mode is coupled to a free-space beam via a focusing grating coupler, which consists first of a taper to expand the mode to that of an 18 μm -wide waveguide, and then a series of curved grating lines with period, duty cycle, and radius of curvature chosen to couple the light to a beam focused near the ion location and polarized in the x direction, illustrated schematically in Fig. 4-2(a), (b).

Light is input to the chip via separate grating couplers designed to couple to a 30 μm -diameter beam; light is focused onto these couplers by a 15 cm focal length lens and at an angle -37° from normal to the chip surface. To reduce possible scatter from the input couplers at the ion location, these input couplers are approximately 6.5 mm from the trap center; thus the in-coupled light is routed on chip in a SM waveguide over about 8.5 mm and through two adiabatic 50/50 power splitters to three focusing couplers at the trap site, and to one output waveguide intended to produce an output beam for optimization of the input coupling. The three focusing couplers are offset by different distances from the trap axis to account for possible misalignments between trap sites and beams, as shown in Fig. 4-1; two sets of these three were included on the chip (each excited by one input coupler), and trapped ion measurements were taken with those labeled Set a in the micrograph. The emitting region of the coupler has an area of 18 \times 18 μm^2 , and design is summarized in the

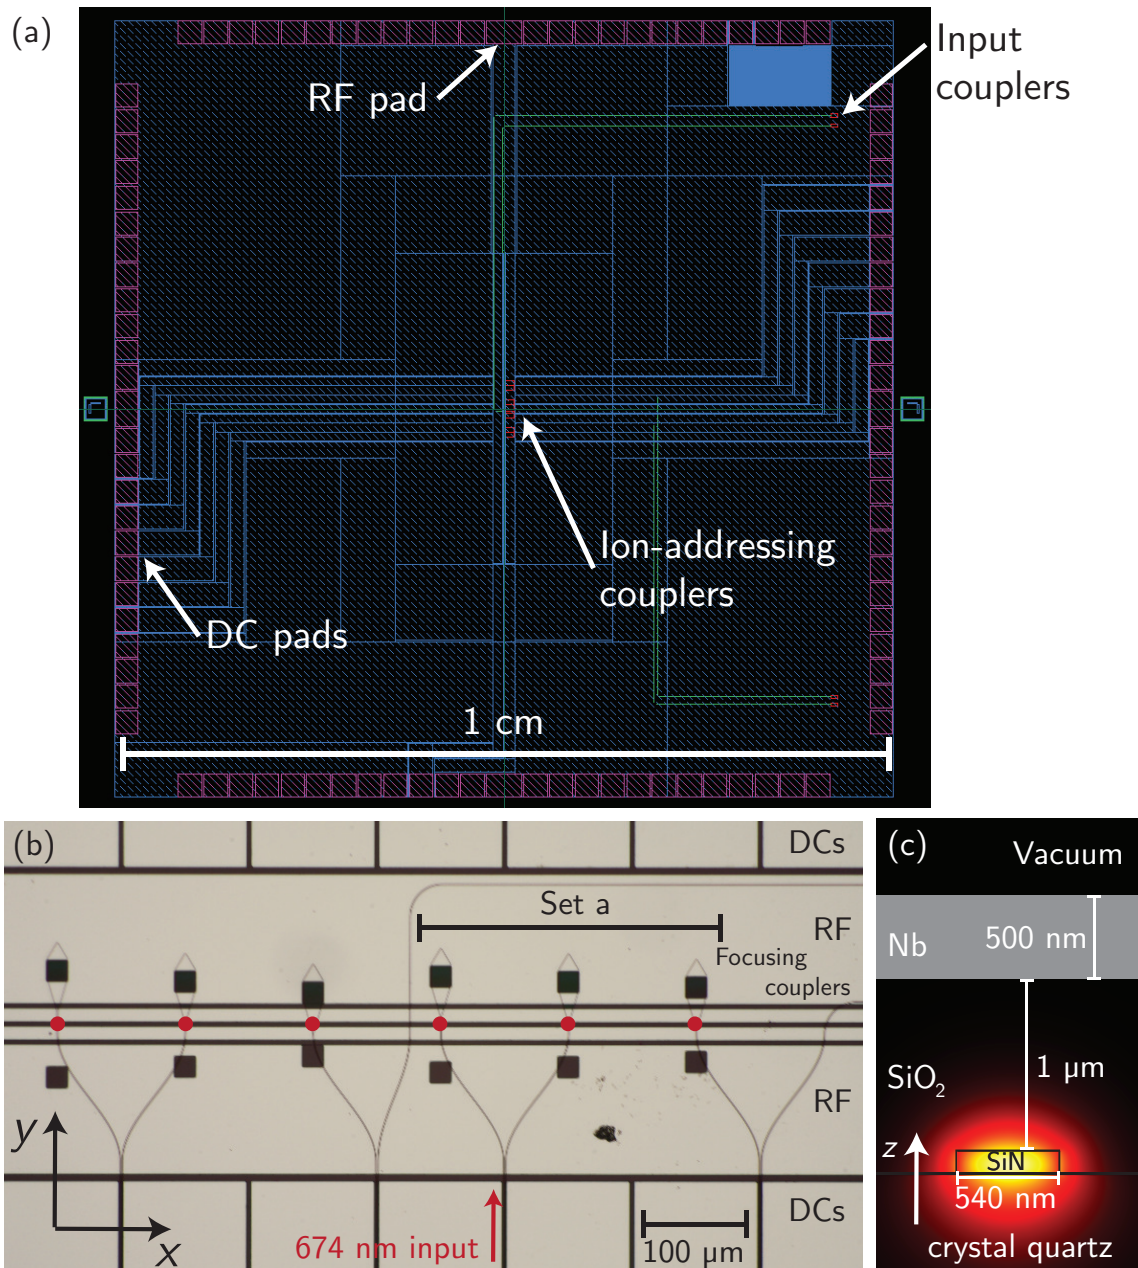


Figure 4-1: (a) Optical micrograph of the designed ion trap with integrated waveguides and couplers underneath at multiple trap zones; waveguides and couplers are visible via topography transfer to the metal. Ions are trapped at one of the positions marked by the red dots, $50\ \mu\text{m}$ above the electrodes, with appropriate potentials applied to the DC and RF electrodes. (b) Simulated electric field mode profile of the single quasi-TE mode (field oriented predominantly horizontally) waveguide used for routing. The crystalline quartz substrate and PECVD SiO_2 form the cladding for the SiN core.

Methods.

The ion trap design is as presented previously [SKC12], except for openings in the RF electrodes (introduced symmetrically about the trap axis) to allow the beams from the focusing couplers to emit through the chip surface. Although the couplers were present on only one side of the trap axis, the 20 μm -square openings in the RF electrodes of the ion trap (Fig. 4-1) were introduced symmetrically around the trap axis to prevent walk off along y of the trapping pseudopotential minimum; along z , 3D simulations indicate that these openings resulted in the RF null moving up away from the electrode from the initial 50 μm by only 1 μm .

4.2.4 Grating design

The design of the couplers is similar to that of those presented in the preceding chapter, although here a square aperture was used and the radii of curvature were not matched to the taper length. Following Fig. 4-2(a) and (b), a taper first expands the mode of the SM waveguide to a larger size laterally; the taper is nonadiabatic and results in curved phase fronts with radius of curvature at the end of the taper approximately equal to the taper length. Subsequently a grating consisting of a series of lines approximately along x with spatially varying period $\Lambda(y)$ and grating strength $\alpha(y)$ emits the light at an angle approximately -38° from normal; backwards emission is preferred to prevent emission into multiple diffraction orders. The lines are parabolic, with curvature radius chosen to focus the beam emitted in free space along the direction transverse to propagation (along x in Fig. 4-2), accounting also for the divergence introduced by the non-adiabatic taper. Due to the orientation of the couplers with respect to the trap axis, and that multiple ions in a given trap zone arrange themselves along the trap axis and can be repositioned with DC fields across many microns along the axis, the couplers were designed to focus only along x ; this eases requirements on alignment between the waveguide and trap metal features.

In these devices, the periodicity ranged from 290 - 310 nm and the duty cycle from 0.1 to 0.4 over the course of the grating, so as to produce an approximately Gaussian amplitude profile along y , while maintaining a constant angle of emission. The grating lines had a curvature radius of 39.6 μm , chosen to produce a beam focused along x near 50 μm above the waveguides.

4.2.5 Device fabrication

The 1 cm^2 die that formed the chip was written three times on a 3-inch crystalline quartz wafer, chosen as substrate for its high thermal conductivity at low temperature and its relatively low optical index ($n \approx 1.54 - 1.55$ at $\lambda_0 = 674$ nm) which serves to keep the optical mode well confined in the SiN core. An Oxford-100 plasma-enhanced chemical vapor deposition (PECVD) tool depositing SiN at 300°C was used to create the SiN film. Following the HSQ resist spin on and softbake at 85°C , to prevent sample charging during e-beam lithography, a thin layer of conductive polymer (E-SPACER 300Z) was spun on top of the HSQ. This layer is necessary in this case as the substrate is insulating; without this grounded conductive layer, charges accumulated

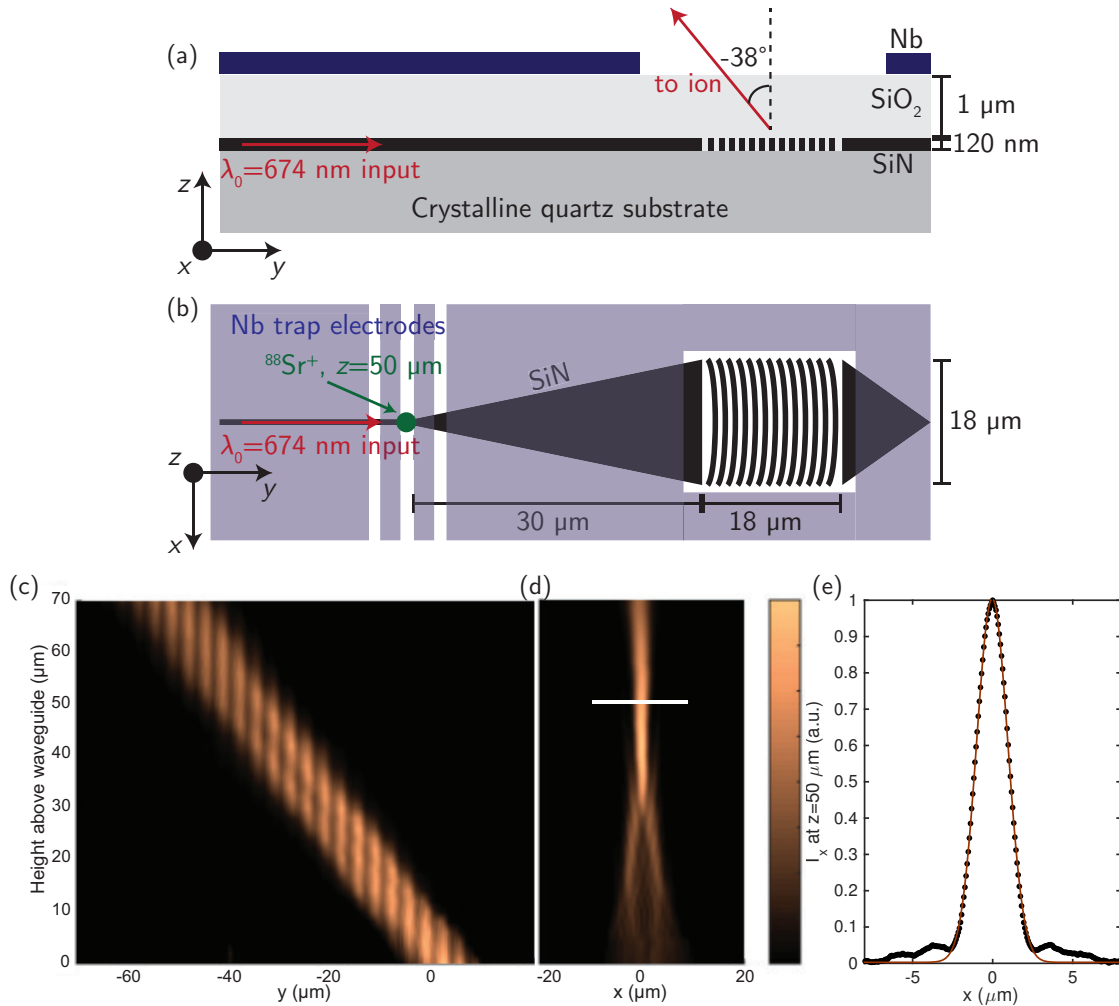


Figure 4-2: (a) and (b) Cross-sectional schematics of the designed focusing grating coupler as integrated with the trap electrodes in the $y-z$ and $x-y$ planes. (c) and (d) “Knife-edge” profiles of the emitted intensity along x and y from 0-70 μm above the waveguide layer, showing collimated emission along x and focusing along y ; striations visible in the y -data are interference artifacts owing to reflections in the imaging system. (e) Beam profile along x near the focus, at the designed trap height of about 50 μm along the section labeled by the white line in (d), the fit (brown line) to the measured points indicates a predominantly Gaussian beam with $1/e^2$ radius $w_0 = 2.0(1)$ μm, with uncertainty arising from the pixel-length calibration.

in the substrate would result in uncontrolled deflections of the electron beam, unlike when using doped Si substrates (as in the previous chapter).

The waveguide fabrication and cladding was done in the cleanrooms at MIT, and subsequently the Nb deposition and trap electrode patterning was performed by staff at Lincoln Laboratory. Electron beam exposure defining the waveguides was performed with an 125 keV e-beam lithography system (Elionix F-125). Following exposure the conductive polymer was rinsed off with DI water and the HSQ was developed in a room-temperature 1% NaOH, 4% NaCl solution for 4 minutes, and further rinsed with acetone and isopropyl alcohol. The pattern was transferred to the nitride film via reactive ion etching (RIE) using CHF_3 and O_2 gases. The same PECVD tool as for the nitride was then used to deposit the SiO_2 cladding (also at 300°C).

Alignment marks written in the nitride were used to spatially reference the photomask for Nb during contact lithography after sputter deposition of the metal film. Following the trap electrode lithography and RIE of Nb in SF_6 , the individual die were diced from the wafer (leaving the independent test structures used for the grating measurements in Fig. 4-2 intact), mounted, and the trap electrodes wirebonded (these steps also by the staff at Lincoln).

Although electron-beam lithography was used to define the waveguides and gratings in this work, the minimum gap size in the grating design here is 30 nm, within resolution limits of current 14-nm CMOS processes. Furthermore silicon nitride waveguides with losses below 1 dB/cm in the visible have been fabricated photolithographically for some time [DMR⁺04]. As such we expect it should be possible to produce the same devices in a CMOS process.

4.3 Grating optical properties

The couplers and waveguides were characterized independently of experiments with trapped ions in separate test structures on the same wafer as the trap-integrated devices. The emission from the ion-addressing couplers was characterized via imaging the emission using the same method as described in the previous chapter.

Fig. 4-2(c) and (d) show the resulting intensity profiles of the emitted light along the y and x directions, showing a collimated beam emerging along y , and focusing along x primarily to a spot with a diffraction-limited minimum $1/e^2$ -radius of $w_0 = 1.8 \pm 0.1 \mu\text{m}$ at $42 \mu\text{m}$, and a slightly expanded waist of $2.0 \pm 0.1 \mu\text{m}$ at the $50 \mu\text{m}$ trap height (profile shown in Fig. 4-2(e)). As the discrepancy in the actual z -position of the focus with respect to the target is less than the Rayleigh range along this dimension, the effect on beam waist is small. The simulated efficiency of these couplers is 32%, calculated from a COMSOL simulation of the grating cross-section as the upwards-radiated power divided by the incident waveguided power.




Figure 4-3: Micrograph of input couplers used for input coupling to free-space beams, with inset showing SEM image of curved grating lines near center.

4.3.1 Input grating couplers

To avoid the need for a direct fiber interface to the chip in the vacuum chamber, for these first experiments a free-space coupling approach was used to input light to the device. Another grating device was designed to form the input coupler, with an emitting region of $40 \times 40 \mu\text{m}^2$, designed to emit (and hence couple to) a free-space beam of approximately $30 \mu\text{m}$ diameter at -37° from normal. The angle was chosen primarily for convenient optical access into the vacuum chamber. This coupler was designed simply with a uniform periodicity and duty cycle, and was not optimized for high-efficiency mode matching. Full 3D simulations of the structure in CST Microwave Studio, and a numerical overlap calculation with an ideal Gaussian indicated an expected coupling loss of 10 dB, which of course does not account for the deviations from ideal Gaussian of the actual incident beam in the lab.

Regarding the beam-pointing instability advantages with this approach, we note that since we presently in-couple from a free-space beam, beam-pointing instability would still appear here in coupling variations. Nevertheless, in-coupling goes as the overlap between the approximately Gaussian profile corresponding to the input grating and the Gaussian input beam profile, and since the overlap between two Gaussians of variance σ^2 with one displaced by d with respect to the other is, as a function of d , a Gaussian with variance $2\sigma^2$, for a given $d < \sigma$ the intensity variation at the

ion would be $\sim 2\times$ lower than if the same beam were directly incident on the ion. If furthermore we take into account that the input beam waist w_g when grating-coupled can be chosen to be larger than that of a focused addressing beam w_i directly incident on the ion, supposing comparable beam displacements in both cases, the coupler would offer lower variations by a factor of $w_i^2/2w_g^2$. These approximate considerations indicate that we can expect improvement in pointing stability with the present coupling, though the full benefit of our approach in this regard would be realized with direct fiber-coupling.

4.4 Individual ion experiments

The full ion-trap device was tested in a cryogenic vacuum setup similar to one described previously [SKC12], with the chip at approximately 4K; after loading, ions could be trapped in the present system for over 6 hours with Doppler cooling. A magnetic field of about 6 G was applied perpendicular to the trap surface along z to break the degeneracy of the Zeeman sublevels; the relevant levels are illustrated in Fig. 4-4.

The input beam incident on the input couplers described above could be aligned a few ways, certainly more robustly than was done for this first chip. The present devices included “output” couplers intended to generate beams that would be detectable outside the vacuum chamber and allow optimization of the input coupling, but in the long waveguides leading to these on the die tested there were defects in the waveguide that prohibited use of these output couplers. Instead, it turned out to be possible to tell when alignment was close simply by looking by eye, via a mirror placed in such a way as to allow sight to the trap chip in the chamber, and noting when light was being sent in through the 40 μm openings in the metal electrodes. When this was the case some light would be visible at the edge of the chip, as it propagates through the crystal quartz substrate and then scattered at the edge of the die. Even with the low power emitted by each focuser, the beams generated were also visible by eye (when looking from the proper direction), and this allowed slightly finer tweaking of the alignment, before ultimately doing so based on the ion signal itself.

Each Zeeman sublevel (of either the $5S_{1/2}$ or $4D_{5/2}$ manifold) shifts by an energy given by

$$\Delta E_{m(s,d)} = m\mu_B B g_{j(s,d)}, \quad (4.1)$$

where $g_{j_s} = 2$ and $g_{j_d} = 1.2$ are the Lande g -factors of the $S_{1/2}$ and $D_{5/2}$ states, μ_B is the Bohr magneton, and B is the magnetic field, and m the magnetic quantum number of the particular Zeeman level. Three transitions are possible from each ground state and hence six in total, and from the equation above the shifts in all transitions can be calculated as a function of magnetic field.

By fixing a length of a pulse of 674 nm light applied and tuning the optical frequency (via the r.f. input to an acousto-optic modulator), the frequencies of these transitions can be found, as when the light is resonant with one the ion will be excited from its ground state. The top panel of Fig. 4-5 shows such a spectrum, obtained




Figure 4-4: Relevant level structure of $^{88}\text{Sr}^+$.

with light input into the focusing grating couplers. In this case, the ion was not initially optically pumped, and was equally likely to have begun in either ground state, and hence the minimum value for ground state probability is 0.5. The red circles represent the positions of the transitions with $\Delta m = \pm 2$, the green crosses those with $\Delta m = \pm 1$, and the blue diamonds those with $\Delta m = 0$ (each plotted at an arbitrary height on the vertical axis), for a field of 6.1 G, from which the carrier transitions were identified.

A fit to a finer-resolution scan around the transition near 80 MHz (and with optical pumping pulses initializing the ion into the $m = -1/2$ state) is shown in the bottom panel of the same figure, together with a fit to the transition probability for Rabi oscillations [HR06]:

$$P_{gs} = 1 - A \frac{\Omega^2}{\Omega_{\text{eff}}^2} \sin^2 \left(\frac{\Omega_{\text{eff}} t}{2} \right), \quad (4.2)$$

where A is an amplitude coefficient allowing for imperfect optical pumping, and the effective Rabi rate accounting for detuning is

$$\Omega_{\text{eff}} = \sqrt{(2\pi(f - f_0))^2 + \Omega^2} \quad (4.3)$$




Figure 4-5: (top) Spectra across the $5S_{1/2} - 4D_{5/2}$ manifold, with points marking the expected transition positions for a 6.1 G magnetic field applied (bottom) fine scan over the $\Delta m = 1$ transition near 80 MHz used for the measurements below, with optical pumping.

4.4.1 Coherent manipulations

Coherent operations here utilized the $\Delta m = -1$ transition starting from the $m = -1/2$ ground state. The ion is optically pumped into the $5S_{1/2}, m = -1/2$ state with six 50 μs -long pulses emitted from the focusers (at $\lambda = 674$ nm), each followed by quench pulses at 1033 nm (see Fig. 4-6a), which allow the $4D_{5/2}$ state to be quickly emptied and hence this pumping to occur rapidly without waiting for the second-scale spontaneous decay. The probability that the electron remains in the S orbital after a pulse on the qubit transition is measured by the presence or absence of scattered light when the ion is illuminated with light near resonant with the $5S_{1/2} \rightarrow 5P_{1/2}$ transition at 422 nm, with 1092 nm light also incident during readout to repump out of the $4D_{3/2}$ state (occasionally occupied during decay from the $5P_{1/2}$ state). As labeled in Fig 4-6(a), the qubit and pumping frequencies were routed to the ions via the integrated waveguides and couplers; in this work the other wavelengths present were in free-space beams.

Fig. 4-6(c) shows the probability that a single ion remains in the ground state after a 674 nm pulse of varying length resonant with the $\Delta m = -1$ transition, with each point representing the average probability inferred from 450 repetitions. Each repetition consists of the pumping pulses to initialize the ion into the $m = -1/2$ ground state, followed by a 674 nm pulse of length given on the x -axis, followed by illumination with the 422 nm laser beam (propagating in free space). If the number of counts detected on the PMT is above a threshold, the state is judged to remain in the ground state and vice versa; the average of these repetitions gives the ground state probability.

With the ion near the beam center, Rabi oscillations in this probability with $t_\pi = 33.2 \mu\text{s}$ are observed (black circles), and with the ion displaced by $7.5 \mu\text{m}$, low probability of excitation is observed (blue squares). The ions in this experiment were not cooled to the motional ground state, and thermal occupancy of motional modes contributes to decay in Rabi contrast with increasing pulse length; nevertheless the fidelity of the first π -rotation is 99%. We verified also that Rabi oscillations with comparable π -times could be observed with light from the couplers at all three trap zones in Set a (Fig. 4-1), illuminated through cascaded 50:50 splitters from a single waveguide.

The profile of the beam emitted from the focuser addressing the ion was measured by translating the ion along the trap axis (x), and measuring the Rabi oscillation π -times at various displacements; since the Rabi rate

$$\Omega_r \propto 1/t_\pi \propto \sqrt{I}, \quad (4.4)$$

with I the optical intensity, this corresponds to a measurement of the beam profile along this direction. The points in Fig. 4-6(d) are well fit by a Gaussian (blue line), indicating an intensity profile with $w_0 = 2.0 \mu\text{m}$ (drawn as a red line). This verifies that the light reaching the ion is predominantly in the focused beam designed.

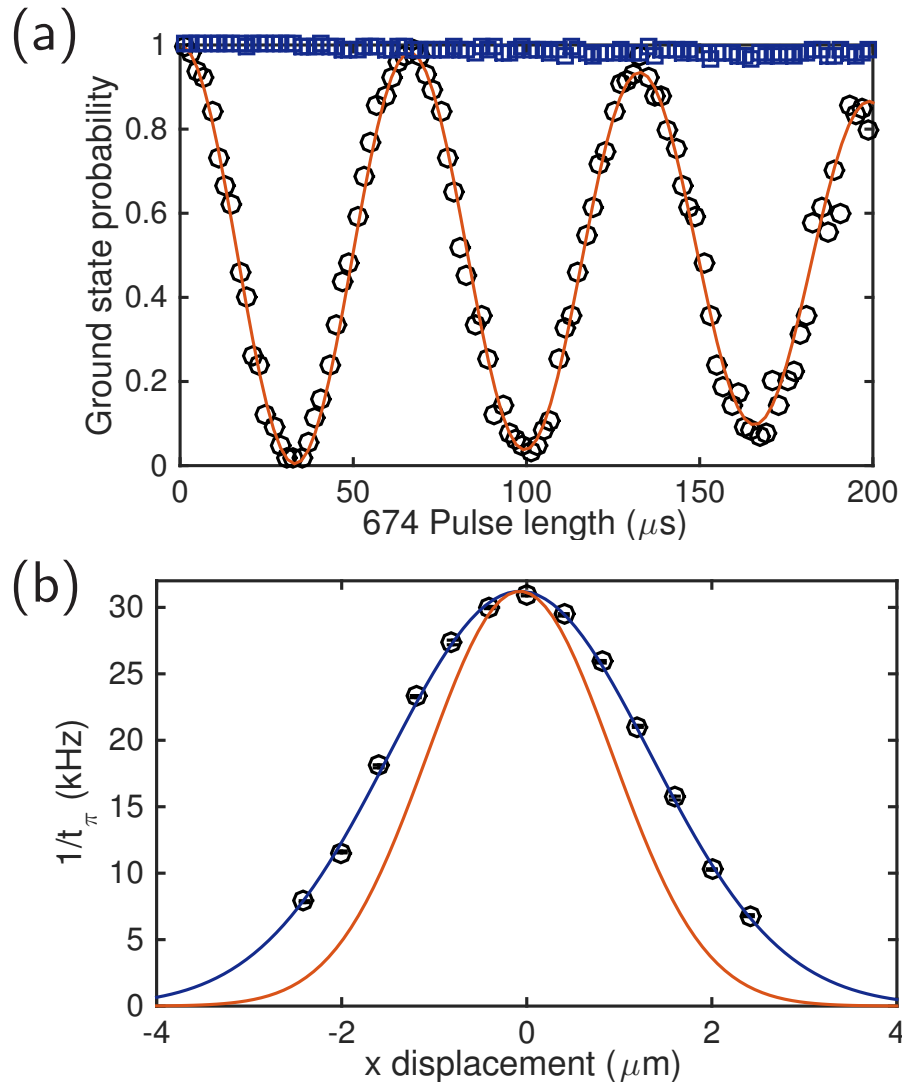


Figure 4-6: (a) Rabi oscillations on the $5S_{1/2}, m = -1/2 \rightarrow 4D_{5/2}, m = -3/2$ transition obtained near the focus of the grating coupler. Each point represents the probability that the electron remains in the ground state after a pulse of varying length over 450 repetitions, and the line is a fit to a Rabi oscillation with Gaussian amplitude decay, from which the first π -rotation's fidelity is determined to be 99%. Blue squares are with ion displaced by $7.5 \mu\text{m}$ along the trap axis, showing low excitation rate away from the focus. (b) Rabi rates vs. ion position as ion is scanned through the focus along the trap axis, with Gaussian fit (blue line) indicating optical intensity profile with $w_0 = 2.0 \mu\text{m}$ (red line).




Figure 4-7: Sequence of EMCCD images of 422-nm fluorescence from a chain of 5 ions, with the middle ion aligned to the grating coupler’s focus and occasionally entering the dark D state due to the addressing 674 nm beam; the sequence spans 2 seconds with frames evenly spaced.

4.4.2 Individual addressing

That this beam could individually address ions was qualitatively observed with 5 ions trapped in the same well. The 674 nm light in the focused beam introduces some probability of occupying the $4D_{5/2}$ state, and the 422 nm light also incident on all ions trigger collapses of the wavefunction into either the bright or dark states, easily imaged on an electron multiplying charged coupled device (EMCCD) camera. Quantum jumps between bright and dark states [SNBT86] occurred only in the center ion aligned to the focus of the center coupler in Set a (Fig. 4-1). This is illustrated in the sequence of images in Fig. 4-6(b), spanning 2 seconds, with the inner 3 ions each separated by about $7 \mu\text{m}$.

4.4.3 Crosstalk quantification and comparison to existing experiments

Particularly for individual addressing in linear ion chains, crosstalk between neighboring ions is an important potential error source [KZI+10, CMBK14, SNM+13], and the simple individual addressing afforded by the ability to tightly focus short wavelength radiation [NLR+99] is a significant advantage of optical in relation to microwave approaches [WOC+13].

We quantified crosstalk errors that would result on a neighboring ion using the




Figure 4-8: Black and red points: imaged intensity of coupler emission along y (see text for description of measurement); and blue diamonds: intensity relative to beam center inferred from cross-talk errors measured with ions variously displaced along the trap axis (with corresponding values of ϵ_x as defined in the text labeled on the right vertical axis), demonstrating crosstalk errors of order $10^{-3} - 10^{-2}$ and below beyond $\pm 7.5 \mu\text{m}$ from center, and of order 10^{-4} past $\pm 12.5 \mu\text{m}$.

waveguide grating for addressing by displacing an ion a known distance from the focus, and measuring the probability of excitation when a pulse of length equal to the π -time at the focus, $t_{\pi 0}$, is applied; this we define as the crosstalk error, consistent with previous work [WOC⁺13]. This probability of excitation is $\sin^2(\Omega_d t_{\pi 0}/2)$, with Ω_d the Rabi rate at the displaced position; for $\Omega_d t_{\pi 0} \ll 1$, this probability and hence crosstalk error ϵ_{\times} is approximately

$$\epsilon_{\times} = \left(\frac{\Omega_d t_{\pi 0}}{2} \right)^2. \quad (4.5)$$

Letting Ω_0 represent the Rabi frequency at the center and noting that $\Omega_0 t_{\pi 0} = \pi$ by definition, for single-photon transitions as used here when Rabi frequency is proportional to the square root of the local intensity, we see that the crosstalk error is proportional to the ratio of the beam intensity at the displaced position I_d to that at center I_0 , where

$$\epsilon_{\times} = \left(\frac{\pi}{2} \right)^2 \frac{I_d}{I_0}. \quad (4.6)$$

CCD intensity measurements and ion crosstalk error measurements were observed to correspond in our device. Both the measured intensity profile near the beam center and along the trap axis, measured with a CCD in the imaging setup used for the data in Fig. 4-2, as well as values of I_d inferred from ϵ_{\times} measured with the ion (blue diamonds), are plotted together in Fig. 4-8, showing good correspondence. To obtain sufficient dynamic range, the intensity measurements are taken at two exposure times, with dark frames subtracted; points from the short and long exposure times are shown in black and red, respectively.

Excitation of higher-order spatial modes in the grating region contributes significantly to the observed deviation from the Gaussian profile (gray line) beyond about $2.5 \mu\text{m}$; this is for the same reason as discussed in Appendix A, describing our observations on mode distortions in the gratings with arcs not well matched to the incident waveguided phase fronts. Optimization of the taper and arcs to tailor the transverse field profile in the grating can reduce the crosstalk errors at displacements of around $3\text{-}5 \mu\text{m}$ (as shown in the previous chapter), an important fact since this is a range typical for many ion trap experiments.

For the $^{88}\text{Sr}^+$ ions used here, at a 1 MHz axial trap frequency the two-ion inter-ion spacing is $4.3 \mu\text{m}$. Although our later devices improved significantly on the results observed in the present trap-integrated devices discussed here, we note that relative intensities of $\sim 1\%$ at $5 \mu\text{m}$ -displacements, as achieved with the grating device used in the integrated ion trap chip, are already comparable to those in existing experiments with individual addressing. For example, the individual addressing beam (addressing quadrupole transitions in $^{40}\text{Ca}^+$) used in ref. [SNM⁺13] is stated to result in maximum $(\Omega_d/\Omega_0)^2$ values of 3%, corresponding to crosstalk errors of a few percent for the $\sim 4 \mu\text{m}$ spacings used. Likewise, the recent demonstration in [DLF⁺16] uses counterpropagating Raman beams for stimulated Raman transitions, one of which addresses all ions and the other which is focused to individual ions. Crosstalk errors there are given as $< 4\%$. Both of these setups used sophisticated bulk optics to reach

such tight focuses within a vacuum chamber, but this residual intensity remains a challenge.

We note that in the case of Raman transitions, if *both* sidebands are focused to individual ions, in this case since the Rabi frequencies $\Omega_{\text{Ram}} \propto E_1 E_2$, where $E_{1,2}$ represents the field magnitude of either addressing beam. If both beams have equal intensity, we have $\Omega_R \propto I$, and hence the crosstalk error would in fact be

$$\epsilon_{x\text{Ram}} = \left(\frac{\pi}{2}\right)^2 \left(\frac{I_d}{I_0}\right)^2, \quad (4.7)$$

and hence significantly easier to reduce if both Raman beams are individually addressed (thanks to David Lucas for pointing this out).

Though these devices reach performances comparable to present experiments, as operation fidelities increase it will be desirable to use components allowing crosstalk at such distances to the $10^{-3} - 10^{-4}$ level or below so it is not a dominant noise source. Ch. 3 presented devices that reach the 10^{-3} level in relative intensity, as well as indication that by simply increasing the device area the same metric could be a few parts in 10^{-4} . It thus seems that integrated waveguides offer a significant advantage on this metric (and crosstalk errors could be much lower if Raman gates with both beams focused are used).

4.4.4 Optical losses

Using a first-principles calculation of the Rabi frequency [Jam98], the $33 \mu\text{s}$ π -time observed, given the measured beam dimensions from the focusing coupler, is consistent with a power of 300 nW being emitted from the grating coupler, 39 dB lower than the ~ 2.6 mW incident on the input coupler. After accounting for the 6 dB designed intensity reduction owing to the two 50/50 splitters in the optical path, the system losses total 33 dB. A number of sources contribute to this loss. Propagation loss in the waveguides was measured in independent test structures to be 6 dB/cm, dominated by material loss; this waveguide loss in our sample contributes 5 dB over the 0.85 cm over which the light is routed on chip. We note that the deposition here was not optimized for loss, but PECVD SiN has been demonstrated elsewhere with material loss as low as 0.1 dB/cm in the red and < 1 dB/cm at as low as 470 nm [GJG⁺08]. The coupler's simulated efficiency of 32% (calculated as the upwards-radiated power divided by the input power in waveguide, from a frequency-domain simulation of the grating) corresponds to a loss of 5 dB, and is due to the approximately vertically symmetric structure of the grating which results in about 50% of the input power being emitted towards both $\pm z$, reflection off of the oxide-vacuum interface, and the finite length and maximum grating strength in the device (18% is not emitted by the end of the grating, and the inverse taper at the end of the grating is included to prevent reflection and re-emission). Due to an incomplete etch of the SiN waveguide layer, however, the coupler's loss may be as high as 8-9 dB as fabricated. Waveguide bends in the path are estimated to contribute 3 dB as well due to the incomplete etch. The remaining 16-17 dB is likely due to the input coupler (simulated efficiency




Figure 4-9: Waveguide loss measurements on a similar sample, before and after metal deposition, showing negligible loss addition from the electrodes. In this case, the material was LF PECVD-deposited SiN, on a fused silica substrate, and the measured waveguide loss as patterned was 4.5 dB/cm.

10 dB for a perfect Gaussian mode impinging) and any excess loss from the splitters on chip.

None of these losses are fundamental, and can be significantly reduced. Waveguide material optimization as mentioned can reduce waveguide loss to a level lower than achievable coupler losses. This, together with optimization of the free-space coupler should bring total loss to about 15 dB (10 dB from input coupler, 5 dB from focusing coupler). With more substantial changes, fiber coupling directly to the chip [TBB⁺02, GHSN⁺13], should allow improvements of input coupling loss to about 2 dB. And incorporation of a bottom reflecting layer, using silicon as discussed in the previous chapter or a metal layer [RGMZ⁺13], or the use of a multi-layer grating [WPK⁺15], could approximately double the focuser efficiency and reduce focuser loss to about 2 dB; therefore, ultimately we expect the total power efficiency can be increased by almost 30 dB.

4.4.5 Polarization purity

The field of the beam is expected to be polarized predominantly along the trap axis. Due to the symmetry of the structure and the SM waveguide feeding the taper and grating (whose dominant E -field component along x is even about z , and whose components along y and z are odd about z and 0 at the center of the waveguide), at focus




Figure 4-10: SEM image of fabricated waveguide cross-section under metal electrodes, on quartz. The image is taken at a cleaved facet, from a region of the same wafer on which the ion trap device tested above was made.

the polarization in principle should be purely along x , the trap axis. By inserting a polarizer into the microscope used for the beam-profile measurements in Fig. 4-2, we measure the power in the transverse field component orthogonal to x over the whole beam to be 2% of that in the dominant polarization component, and a minimum polarization impurity at center, consistent with the symmetry argument above, of $< 5 \times 10^{-4}$ in relative intensity. Along with these orthogonal transverse components, which result from the taper and grating design, away from the beam center there should additionally be longitudinal field components (pointing along the propagation direction); as these are localized to the focal region, our present polarization measurement is not sensitive to them, though these longitudinal components arise when any Gaussian beam is tightly focused and are not due to the grating design. We leave a full characterization of these field components orthogonal to the dominant one along x and their effects on relative transition rates of the $5S_{1/2} \rightarrow 4D_{5/2}$ sublevels for future work.

4.4.6 Photo-induced charging and stability

Finally, we note that although light propagating through the trap chip surface may result in photo-induced charging of the dielectric [HBHB10], we found that the compensation voltages (applied to the DC electrodes to compensate any stray fields developing during the course of operating the trap) were stable over days, suggesting negligible charging from the 674 nm light; possible charging at shorter wavelengths, if eventually problematic, could be overcome for example by coating the electrode openings with a transparent conductive layer [EWA⁺13].

4.5 Conclusion and future work

This work demonstrates the possibility for large-scale nanophotonics integration within planar ion traps, allowing focusing at ion sites, flexible and parallel routing in complex geometries, robust alignment stability, and utilizing scalable planar fabrication. We anticipate a number of avenues for future work. The use of lithographically defined nanophotonic waveguide devices is readily extensible to more complex optical arrangements, at the other wavelengths required for full ion control as well, and co-design with more complex traps; the CMOS-compatibility of both the planar ion traps [MEB⁺14] as well as the optics [OMS⁺12] furthermore suggests an approach for fabrication of eventual large-scale systems. Fiber coupling directly to the trap chip in a scalable fashion, either still with grating couplers [TBB⁺02], or direct fiber-waveguide interfaces [GHSN⁺13], presents a significant technical challenge, but would be required both for a scalable, parallel optical interface to the chip, and to realize the full benefit of beam-pointing stability with this approach. Integration of electro-optic modulators [XPT12], or controlled ion movement through the beams formed by the grating couplers [LKOW07, dLM⁺16], with the waveguides and gratings here would further enable parallel encoding of quantum operations to multiple sites, in addition to addressing. Given the demonstrated practicality of scaling such planar

photonic systems, their ability to operate across a wide range of wavelengths, and the robust alignment and focusing allowed by integration within the trap chip, we expect this approach will significantly reduce the complexity of optical systems required to implement nontrivial QIP with large ensembles of trapped ions.

4.5.1 Direct fiber coupling

Direct fiber coupling would be preferable to the free-space coupling used here to allow multiple beams (including at different wavelengths) to be directed to ions without the need for repeated beam alignments, and to fully realize the beam pointing stability possible with this approach. In future devices this could be done a number of ways. A variety of adiabatic waveguide couplers based on tapered fibers have been demonstrated [GHSN⁺13, TNT⁺15], as well as those based on butt-coupling a standard SMF to a tapered waveguide [YHM⁺14, BJPK⁺15]. The latter may be more straightforwardly scalable, even if not quite as efficient as the adiabatic approaches. The fibers can be attached via a low-temperature compatible adhesive to the chip (one example attachment is in [HOT⁺10]), and hence this appears straightforwardly implementable.

4.5.2 Waveguide losses

Work at Lincoln Laboratory to develop a more scalable foundation for fabricating such structures as were demonstrated in this chapter is underway. LPCVD material there has been observed to allow guiding in SM waveguides with losses of 0.1 dB/cm at $\lambda = 1092$ nm, 0.7 dB/cm at 674 nm, 3.8 dB/cm at 427 nm, but 19.5 dB/cm at 405 nm. It thus appears likely the silicon nitride itself could be used for all but the photoionization lasers for $^{88}\text{Sr}^+$.

For lower wavelengths, Al_2O_3 (alumina) has been observed to be a candidate core material, with propagation losses < 4 dB/cm observed in films at wavelengths as low as 250 nm [AWB⁺10]. Such measurements are carried out in films deposited via atomic layer deposition (ALD), and which have also been shown to be etchable via RIE. Hence alumina, though it has a slightly lower index of refraction, appears a potentially interesting material for shorter wavelength operation, as necessary for some ion species, and it will be important in future work to explore this and other possible materials, as well as the influence of photo-induced damage in either the core or cladding in limiting short-wavelength operation in nanophotonic structures. Similar techniques to those used to reduce the effects of solarization in silica fibers at short wavelength [CSW⁺14] may be employed in the SiO_2 cladding; the susceptibility of different core materials to solarization may be interesting to study.

Of course, the capability demonstrated in this chapter relies on passive optics; while this itself can be a powerful addition to trapped-ion systems, especially given the capabilities for gates encoded via the controlled motion of ions [LKOW07, dLM⁺16], more general capabilities would be allowed by active devices for detection and modulation, which we discuss in the following two chapters.

Chapter 5

CMOS integration and silicon avalanche photodiodes

The optics described in the previous two chapters are formed via planar fabrication with materials (in this case Si_3N_4) that can be patterned within CMOS processes. This chapter focuses more explicitly on ion-trap devices fabricated within existing CMOS foundries. Motivations for this work include leveraging the established manufacturing infrastructure in CMOS foundries to reliably create multi-layer traps as required for many ion experiments, but which require significant development for individual groups to achieve, e.g. in [AUW⁺10]; the ability to integrate complex circuitry to carry out classical computations required for general computations efficiently and with low delays on chip; and the possibility to leverage CMOS fabrication capabilities to achieve large-scale trapped-ion systems with integrated optics, as described earlier in Chapter 2, and developed in the latter half of this chapter in the discussion of integrated CMOS detectors.

This chapter first describes the implementation of a CMOS ion trap and its characterization, as published in [MEB⁺14] and proceeds to discuss subsequent work on integrated Si single-photon-counting avalanche diodes (SPADs) for integrated readout of ions' quantum state via fluorescence.

5.1 CMOS-fabricated ion trap

Fig. 5-1 shows a schematic cross section of patternable layers, aside from the various implants controlling devices in the Si substrate, in the 90-nm CMOS process (with IBM/Global Foundries identifier 9LP) used for the devices presented here.

5.1.1 Design and implementation

Our aim was to implement an ion trap design provided by Amira Eltony, which specified the dimensions of the RF and DC electrodes, and the gaps between them. The drawn design has to be compatible with the set of design rules specified by the foundry in their “Process Design Kit” (PDK), which ensures that none of the drawn




Figure 5-1: Cross-sectional “back-end” stackup of the 90-nm CMOS process used. For the trap presented here, the m5 layer was used as the ground plane, and the top Al layer for the trap electrodes; the intermediate layers m6-m8 were used only in the regions where vias were placed to connect the ground plane to the trap ground electrode.

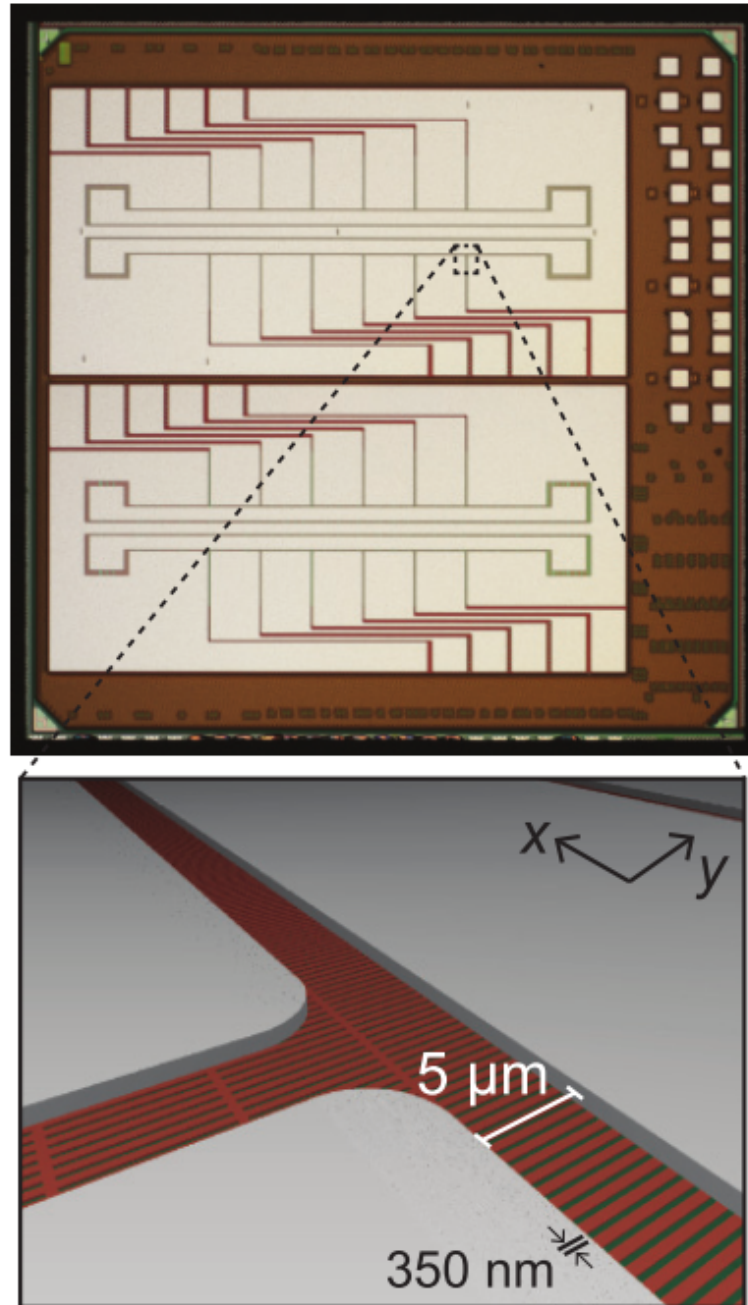


Figure 5-2: Micrograph and close-up perspective of GDS showing meshed ground plane.

features are likely to result in structures that adversely affect yield of other devices (for example structures that may detach), or are below the resolution limits of the process. In this case the relevant design rules were these resolution limits, together with density requirements (which dictate minimum and maximum areal densities of features in metal layers, related to the requirements of the chemical-mechanical polishing applied to the metal layers).

Devices were fabricated on $3 \times 3 \text{ mm}^2$ die (Fig. 5-2) on a shared multi-project wafer produced in a 90-nm CMOS process. This process is primarily utilized for integrated digital circuits, and the trap die was one of many designs fabricated in parallel on the same wafer. This is a roughly 10 year-old process, and certainly a newer process could have been used; however the additional cost of the newer nodes was not necessary for the features made here, and the higher density, shallower implants can make implementing APD devices as discussed later on in fact more challenging. In terms of the trap itself though, we expect similar structures can be implemented in a range of process nodes.

The process allows for patterning of 8 copper interconnect layers, along with the top aluminum pad layer (Fig. 5-1). This $1.3 \text{ }\mu\text{m}$ thick pad layer was used for the trap electrodes, and a copper layer (m5) approximately $4 \text{ }\mu\text{m}$ below the aluminum layers bottom surface and $2 \text{ }\mu\text{m}$ above the silicon substrate was used to form a ground plane under the extent of one of the traps. Due to metal density constraints arising from chemical- mechanical polishing steps applied to the copper layers, this ground plane was patterned as a mesh of 600 nm strips separated by 350 nm along the x direction and $10 \text{ }\mu\text{m}$ along the y direction (see Fig. 1).

This geometry for the ground plane was chosen to be compatible both with the maximum density (80% within any given “checking box” of area $280 \times 280 \text{ }\mu\text{m}^2$) of patterns in this layer, which of course precluded a uniformly full pattern, as well as with the minimum wire (and spacing) width of 140 nm, and the minimum “enclosed area” of $0.36 \text{ }\mu\text{m}^2$ (for comparison, in this 90-nm process the corresponding dimensions for the polysilicon transistor gate layer are 80 nm minimum width, 140 nm minimum spacing. This process is fairly old (90-nm nodes became available roughly around 2004-2005), and current nodes at 14 nm can of course reach much smaller feature sizes.

In the Aluminum pad layer, the minimum feature width was $3 \text{ }\mu\text{m}$, and the minimum gap width was $1.6 \text{ }\mu\text{m}$, both easily compatible with the design – here the RF electrodes were each $60 \text{ }\mu\text{m}$ wide, and the center ground $50 \text{ }\mu\text{m}$, with spacings of $5 \text{ }\mu\text{m}$ between electrodes. Also, the Al wires leading to each DC wirebond pad from the DC electrodes were $15 \text{ }\mu\text{m}$ wide, with $15 \text{ }\mu\text{m}$ gaps between them. Furthermore there is no maximum density limit on this layer, as no CMP is applied to the Al layer, and hence no fine meshing of the trap electrodes is required. We note that by default a polymer encapsulation is deposited over the top of the chip, and an additional layer, labeled “LV” in this particular PDK and serving as the “pad opening”, defines where this encapsulation is excluded. This layer was simply drawn over the entire area of the ion trap to leave the Al as the top-most layer everywhere.

Metal vias connect the copper ground plane in the m5 layer to the center electrode of the trap through the upper metal layers m6 to m8; as illustrated in Fig. 5-3(c), each




Figure 5-3: Images of mask layout for CMOS trap. (a) Area near RF electrode bond pads; purple lines mark the “exclude” shapes blocking the generation of metal fill shapes [Orc12] for layers about the ground plane in the gaps around the RF electrodes. (b) Ground plane mesh and (c) Via stack connecting the Al pad layer to the m5 layer. The four 5 μm squares are vias between Al and the topmost layer, and the vias between the lower layers are formed as a more fine array of either 280 or 140 nm square shapes (standard for the process).




Figure 5-4: Micrograph of a diced and wirebonded CMOS trap, with inset showing the image of a pair of ions trapped.

via was approximately $25 \times 6 \mu\text{m}^2$, consisting of process-standard via arrays between the metal layers. These vias were placed symmetrically at a few points on the trap, as labeled in Fig. 5-2. Since it was expected the vias would be polished and not result in any disturbance of the profile of the top Al layer, one was placed as well right at the center of the trap. The vias were visible under a microscope and were observed in the ion-trap experiments to result in excess scattering of the 422 nm light, indicating that they did disturb the Al layer profile slightly, and it would help such experiments to keep the vias away from the trap site by at least the beam waist of the beams used.

Each trap (individually $2.5 \times 1.2 \text{ mm}^2$) was intended to be diced out of the full die and tested independently. An example of a diced chip, with the various DC and RF contacts wirebonded is shown in Fig. 5-4. Here since just the center trap zone was being tested, the outer two DC electrodes on either side were shorted to ground as labeled. We also point out that the various beams for ion manipulation and readout were propagating along the diagonal from lower left to upper right in this image, and hence the wirebonds situated to either opposite corner allowed these beams to propagate without clipping. The widths of the wirebond pads in this case were chosen conservatively to be $150 \mu\text{m}$.

Corners were rounded to $5 \mu\text{m}$ radii (as illustrated in Fig. 5-2, and again with a discretization of 5 nm) on this iteration given concerns about possible breakdown due to fringing fields. However this appeared not to be necessary, since the second generation trap (with integrated APDs) discussed below did not have rounded corners




Figure 5-5: Leakage current in DC HV measurement on CMOS trap chip. Current noise in the high-voltage source-measure unit could have been reduced by use of longer integration times but the present noise level was sufficient to verify the lack of breakdown.

and still had no difficulty with breakdown with comparable potentials applied.

5.1.2 CMOS trap characterization

To assess the likelihood that the 100 V-level RF potentials applied to the trap electrodes would induce dielectric breakdown, DC breakdown testing was carried out initially to make sure the trap electrodes did not break down at the surface, or that breakdown would not occur between the RF electrodes and the ground plane $4\ \mu\text{m}$ below. This characterization was performed at room temperature; since the measurement was not in vacuum, a high dielectric strength fluid (FC-770 Fluorinert fluid from 3M) was used to coat the area around the electrodes tested to prevent air breakdown (which at atmospheric pressure can breakdown at just $\sim 3\ \text{V}/\mu\text{m}$) from limiting the measurement. No breakdown was observed for voltages up to 200 V, applied between one of the RF electrodes and one of the DC electrodes (or ground), with all other electrodes floating. Furthermore leakage currents were below 15 pA even at these high voltages, indicating high quality dielectrics. A sample current trace obtained during such a measurement is shown in Fig. 5-5.

Devices fabricated without a ground plane, both in this process and in the Micron process showed significant variations in the RF power delivered to the chip (monitored via the RF power reflected from the chip in the Lincoln setup) as a function of temperature. This may be due to the particulars of RF conductivity (and loss) at low temperatures in doped silicon substrates [PG61], and related to but not the same as observed in experiments on undoped Si substrates [NLK⁺14]; but in addition, without the ground plane stray light from the visible beams involved can excite carriers in the substrate which alter the RF loss and hence the RF potential applied to the chip, or more directly generate stray fields that destabilize the trap. Traps with a ground plane exhibited no such behavior and allowed stable ion trapping for over an hour with trapping (equivalent to the best traps otherwise tested in the same vacuum




Figure 5-6: Measurements of heating rate, performed by the Lincoln Laboratory group, in CMOS foundry trap, via relative strengths of motional mode sidebands following some delay after ground state cooling. These measurements were taken at 8K chip temperature, 1.3 MHz axial mode frequency, and 50 μm trap height.

chamber at the time).

The heating rate of the axial motional mode at 1.3 MHz was measured to be approximately 80 quanta/s at a chip temperature of 8.4 K (Fig. 5-6), approximately a factor of 10 larger than comparable geometry traps tested at the same temperatures [CS14]. This increase may be due to the relatively large surface roughness measured on this chip, of 35 nm r.m.s. (Fig. 5-7), as compared to the 2 nm seen on single-layer Niobium traps otherwise fabricated at Lincoln; naively one would expect the increased surface area to result in a larger number of surface impurities that may contribute to heating, and quantitative studies exist as to how the fine local environment around some adsorbate may affect its contribution to field noise [LLC16], but the reason for the heating rate is very much uncertain. Nevertheless the heating rate is comparable and with improvements to the top Al surface (which is not treated with CMP), which may be as simple as reducing the thickness (and hence also the roughness), very slightly modified processes should reduce heating rates to those seen in the best traps in any fabrication process. Additionally we note that there is of course variation in the heating rates observed even in the same trap, as summarized in Fig. 5-6, and in addition a later trap in essentially the same process was observed to have a lower heating rate of around 20 quanta/s (see below on traps with integrated APDs). Hence




Figure 5-7: Example surface profile, via atomic force microscopy, of the Al surface of the first CMOS foundry trap; the measured profile resulted in an approximately 35 nm root-mean-square roughness.

motional heating rates and trapping performance in CMOS foundry traps appear to be comparable to those observed for comparable trap heights in the best planar traps generally.

For further details relating to the trap design, we refer to Amira Eltony’s thesis [Elt15], and for additional details pertaining to the ion trap characterization performed by the Lincoln Labs group, we refer to the publication [MEB+14].

In general these results indicate the possibility of using CMOS foundry processes to create ion trap structures with performance comparable to custom-fabricated planar traps, but with straightforward possibility utilizing the multiple metal layers patternable with high resolution to create much more complex trap structures. The possibility for integrated optics of various kinds in similar processes [OMS+12], as well as of course sophisticated electronics, we believe will make possible the practical implementation of interesting ion systems. The work presented in the next section pertaining to integrated silicon photodetectors is an initial step along this direction.

5.2 Integrated silicon avalanche photodiodes

A second-generation CMOS trap was made with the intent of exploring integrated silicon avalanche photodiodes for fluorescence readout. In Chapter 2 we discussed the potential merits and challenges associated with integrated detectors of any kind; in this section we first discuss the performance of CMOS APDs observed by other

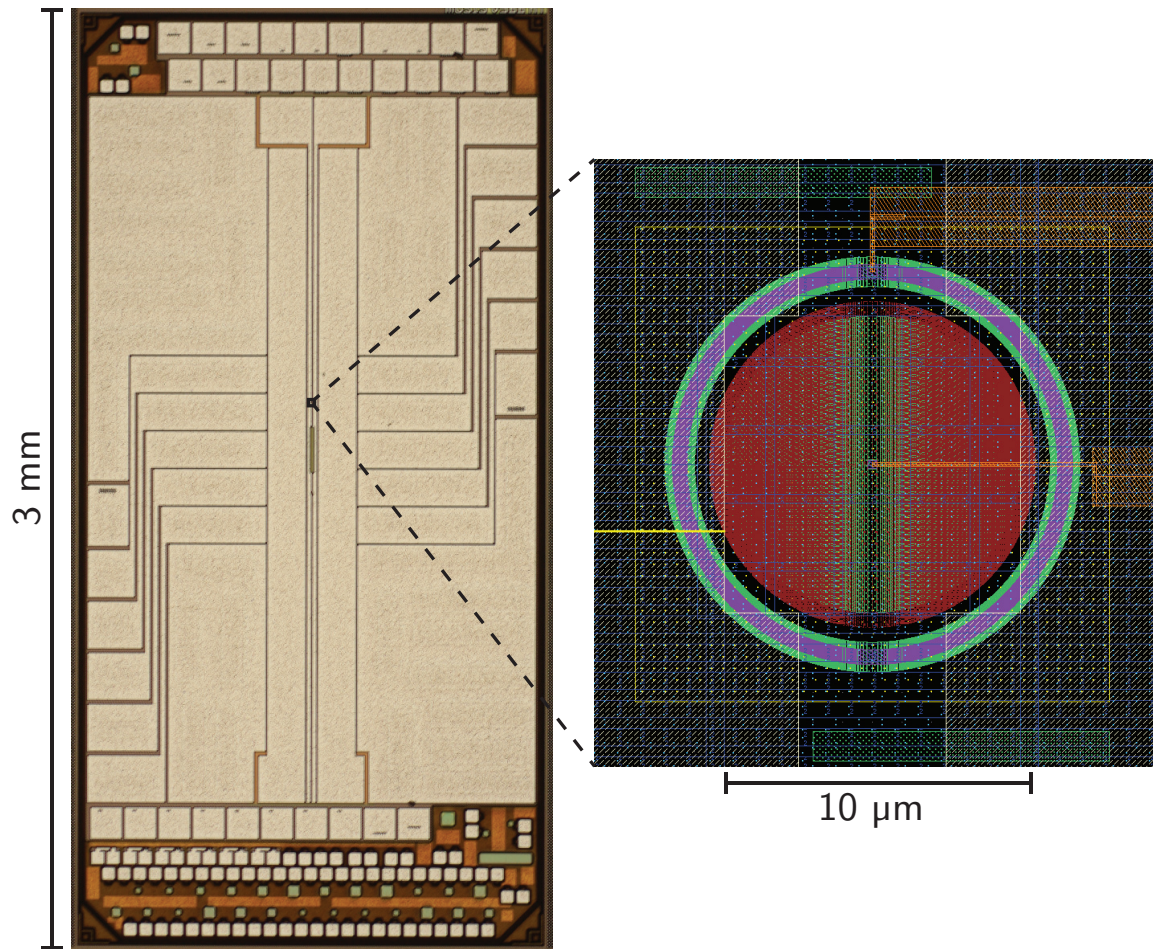


Figure 5-8: APD trap chip micrograph and inset with mask layout of STI GR device.

groups in various contexts and argue based on previous results that they could allow high quality ion readout, and then present our own observations on CMOS APDs in a 90-nm process very similar to that used for the CMOS trap above. The process designation here was 9HP, which is a “Bi-CMOS” variant of the 9LP process with additional process steps for SiGe devices, which were not used here.

5.2.1 Trap-integrated device layout and design

A micrograph of the trap chip with APD devices integrated with the ion trap is shown in Fig. 5-8. The APDs here are surface-illuminated devices with circular active regions in the silicon substrate (designs described in the section below, and with mask layout for one of five variants placed beneath five different trap sites represented in the figure inset); the same ground plane as used in the first CMOS trap shields most of the trap area from the silicon substrate, but just above the APD active areas the mesh is relaxed to a pattern with $2\ \mu\text{m}$ -square openings meshed by 500 nm-thick wires, to allow $\sim 85\%$ of the light to pass through (calculated simply from the metal

fill factor over the active region and neglecting diffraction effects) while still shielding the potential on the APD from the trap site.

The large 100 μm -square pads above and below the ion trap in Fig. 5-8 are contacts for the 5 APD devices placed under the trap, as well as for a counter circuit designed by Michael Georgas also located under the trap and intended for interfacing to an APD. The results discussed below on characterization of the APD devices were obtained from the independent test structures placed below and addressed via the smaller pads visible at the bottom of the device, which allowed simple probing of a larger number of variants (some of which were drawn identically to those under the trap). We proceed to discuss the designs of the APDs and their realized characteristics before coming back to the presently observed behavior of the integrated device.

5.2.2 CMOS APD designs and room temperature DC characteristics

APD devices are typically formed in the substrate via a junction composed of a shallow implant, often a highly p-type doped (p+) region, and a deeper n-type well region around; at the interface a built-in electric field accelerates generated electron-hole pairs to either contact (Fig. 5-9).

A typical need in the design of avalanche photodiodes operating above the breakdown voltage is a “guard ring” (GR) structure around the edges of the junction to somehow reduce fringing fields in the silicon which would otherwise lead to breakdown at the junction edges at a lower voltage than required to reach breakdown over the whole device area [SN06]. A variety of approaches to achieve this practically have been demonstrated [DBSR⁺11]. One solution is an oxide guard ring, in CMOS modern processes formed of the oxide used for “shallow trench isolation” (STI) which is used to isolate the channels of nearby transistors; while this solution is fairly robust [LRC12] and also allows for high density packing of APDs for a sensor array, the way the STI is patterned results in a Si-Oxide interface with a high density of defects which has generally resulted in high dark count rates on the order of 1 MHz [FHE06] (here, though counts are not periodic and count rates should strictly speaking be given as counts/s and not in Hertz, we nevertheless use Hz with this understanding).

The use of a GR formed of an additional lower-concentration p-type implant can allow low dark-count rate operation (e.g., [RGF⁺03]), and in a 130-nm process the use of such a structure allowed 100 Hz-level dark counts at room temperature [FLC⁺10]. Deep graded junctions with carefully tailored field profiles have allowed impressive performance, with peak QEs of 70%, and broadband operation, with < 50 Hz dark count rates [WGH12], with similar structures also implemented in a 90-nm CMOS imaging process [WRGH11].

In our own work, we focused primarily on structures based on STI GRs (cross-section shown in top panel of Fig. 5-9) – devices with this guard ring were initially laid out following simulations performed by Sabareesh Nikhil Chinta and with Jason Orcutt on the first CMOS trap chip. Other structures studied additionally on the




Figure 5-9: Layers involved in main APD designs attempted. shallow n+ implants are automatically placed where bp (p+ implant) is not present, to form the n-side contact. Mask layer names used in the figures are as follows: n3 denotes the p-well implant used for “triple-well” pFETs; sblk refers to “silicide-block,” which blocks the formation of an optically lossy metal/silicon contact layer at the substrate surface; rx and pc refer respectively to the body silicon and polysilicon, with “fill” shapes excluded over the APD area.




Figure 5-10: Dark I-V Characteristics of main APD designs attempted (for $5\ \mu\text{m}$ -diameter active region variants), showing premature breakdown in p-type guard ring devices (with STI) and high leakage in devices without STI.

second CMOS chip aimed at imitating the devices from [FLC⁺10] (bottom panel of the same figure), and devices like those in [FLC⁺10] but without the STI altogether. In the figures, the “nw” label refers to n-well which extends over the entire device’s extent, “bp” is the layer that defines the p+ implant at the p+ contact (in this case the inner circle), and an n+ implant is placed by default in the body silicon where p+ is not present (i.e. in this case in the outer ring around the device – see also in the inset of Fig. 5-8. “n3” designates the p-well implant associated with what are known as triple well nfets, which have a channel in a p-well implant isolated from the main substrate; this was the implant used as the p-type guard ring here, and in [FLC⁺10] albeit in the latter in a 130 nm process. “rx” finally is the layer that defines where the STI is *not*, i.e. where the body silicon extends to the top of the substrate.

Dark current profiles for devices of $5\ \mu\text{m}$ diameter formed on the same chip of these three designs are shown in Fig. 5-10. It was immediately clear that whereas the STI GR devices had a breakdown voltage near 12 V reverse bias, the p-type guard ring devices turned out to exhibit a lower breakdown near 10 V, which since the central area of the junction is formed of the same implants, suggests premature breakdown at the edges. And the devices with no STI exhibited high leakage currents, making




Figure 5-11: Dark and illuminated I-V characteristics of STI GR devices of 10 μm and 15 μm diameters. The photocurrents measured at low reverse biases indicate a quantum efficiency of 15%, and the independence on device area indicates that the devices are indeed capturing almost all of the light incident from the SMF.

them unusable as photodiodes and certainly as Geiger-mode devices. It is unclear why devices with ostensibly the same lateral cross section as those of [FLC⁺10] exhibit premature breakdown, but it may be due to different concentrations of the triple-well implant in the 90-nm vs. 130-nm process. On the other hand this behavior is similar to that observed in [DMR⁺04], suggesting perhaps some variation in processes.

Though it should of course be possible to reproduce similar devices to the other GR structures cited above, whether in this process itself or in one with different implants, our characterization focused on the STI GR devices given their good breakdown performance. Dark and illuminated IV curves of these devices are shown in Fig. 5-11, illuminated IVs being taken with a measured power of light from a 405 nm laser diode incident in a single-mode fiber with mode field diameter of approximately 4 μm positioned nearly above the active area.

From these measurements the quantum efficiency (the fraction of incident photons

that lead to charge carrier collected as current) is calculated as:

$$\eta_Q = \frac{I h\nu}{P e}, \quad (5.1)$$

with I the photocurrent and P the incident optical power, and ν the optical frequency, the e the elementary charge. This evaluates to 15% given the low-bias photocurrent observed in the present devices at 405 nm (Fig. 5-11), comparable to the results in [FLC⁺10]. In carrying out these measurements, coupling losses between fiber patch cords when using threaded metal mating sleeves supplied by Thorlabs (as for example when measuring guided power) could vary by a few dB, given the short wavelength. The use of unthreaded ceramic mating sleeves allowed for more reproducible couplings (< 1 dB variability), which was helpful for establishing confidence in fiber-guided powers (which can ultimately be best checked by measuring the output power from the fiber addressing the photodetector on a free-space power head). In Geiger mode the total photon detection probability (PDP) would be $\eta_Q\eta_P$, where η_P is the voltage-dependent probability that a generated e-h pair leads to an avalanche pulse.

5.2.3 STI-GR device characterization and low-temperature measurements

The important questions for operating these devices for ion trap readout was whether the dark count rate would be low enough, and the device would continue to operate at low temperatures in pulsed mode.

Quenching circuit configuration, pulsing, and cryostat measurements

A variety of quenching circuits are possible for ending the avalanche process, which if the overvoltage on the diode was maintained would be a persistent effect [CGL⁺96]; here we use passive quenching in the “current mode”, where a series resistor develops a voltage in response to current flow in the junction that opposes the applied bias, and the current pulses are read out via the voltage on a smaller “sense” resistor. Since we were interested in low temperature behavior, the devices were tested in a flow cryostat system, in which the nature of the heatsinking in the two-terminal probe contacting the device required that one terminal of the device be grounded.

This resulted in the use of a quenching circuit configuration shown in Fig. 5-12; here, the voltage across the sense resistor R_s is proportional to the current flowing through the junction, and the “stray” capacitance is chosen to be in excess of any parasitics in the circuit (namely the 10s of pF stray capacitances associated with the fact that the device was contacted through a two-terminal probe connected to a coaxial cable) to ensure that diode discharge current flows through the sense resistor and not through the diode’s own capacitance [CGL⁺96].

At room temperature, pulse sequences like those shown in Fig. 5-13 could be observed (with amplitudes corresponding to voltages after amplification with an SRS preamplifier). With the threshold set at approximately 40 mV, using a Keysight




Figure 5-12: Schematic of passive quenching circuit used.

pulse counter (53220A, Universal Frequency Counter), the count rate could be measured as a function of applied voltage, and is plotted in Fig. 5-14 for the STI GR devices of 5, 10 and 15 μm -diameters.

If these pulses were due to thermal generation, we would expect dark count rates to decrease substantially (exponentially) with temperature. However we see, from the data in Fig. 5-15, that as temperature decreases, little decrease in the observed dark count rate is seen for a given overvoltage (the decrease in the breakdown voltage evident from the onset of pulsing in this plot is due most likely primarily to the increase in mobility at lower temperatures).

Additionally, we note that the quantum efficiency near 90 K as measured by DC photocurrent at low reverse bias was found to be higher than at room temperature by almost 60% (i.e. the QE was observed to be almost 25% at these temperatures). This may be due either to some effect of carrier freeze-out increasing the depth of the junction and the fraction of light absorbed, or due to the higher mobility at low temperature.

Such behavior had been separately seen in a previous measurement on similar devices (same design, but on the first CMOS trap tapeout), where the photocurrent at a single bias point was measured as a function of temperature, at $\lambda = 674$ nm. Here the DC photocurrent at -2 V bias was monitored as the cold head fell in temperature (blue points), periodically realigning the fiber input over the device to maintain maximum photocurrent. A clear increase of nearly $4\times$ was observed near 40 K, perhaps due to the effect of freeze-out in the n-well implant, and a resulting increase in the depletion region's width. Red points were measured as the cryostat increased in temperature, with considerably less confidence in the fiber alignment due to a rapid changes in temperature in this range. The in-fiber power was constant to within 20% for this measurement, and the dark current was well over two orders of




Figure 5-13: Pulse traces observed in Si APD with current-mode quenching, using a 65 k Ω resistor, read out on a 150 Ω resistor and amplified (From 14 Dec 2015).




Figure 5-14: Dark count rates at room temperature for STI GR devices of various diameters.




Figure 5-15: Dark count rates with 10 μm-diameter STI GR APD at various temperatures as measured in the cryostat and using the Keysight counter.




Figure 5-16: DC photocurrent measurement as a function of temperature at 674 nm for -2 V applied, for a STI GR device from the chip with the first CMOS trap. Room temperature QE was approximately 3% owing to the long absorption length compared to the junction depth; carrier freezeout in the n-well implant and resulting thicker junction may be responsible for the higher photocurrent observed at around 40 K. Blue points were taken as the sample cooled down from room temperature.

magnitude below any of these photocurrents values. These measurements gave initial indication that at least the DC photoresponse would persist at low temperatures, and even increase (which as we noted above was later confirmed also in the blue).

Afterpulsing and bunching behavior

Returning to the observation that the DCR did not decrease with temperature as measured with the device operating in Geiger mode, this lack of a decrease in DCR (and even apparent substantial increase below 155 K) turns out to have a more subtle interpretation, as can be seen in the oscilloscope traces obtained at near 91 K and plotted in Fig. 5-17. Here it turned out that there were ms-scale periods of rapid pulsing separated by periods of no pulsing, which we interpret as a series of afterpulses following an initial generation and avalanche event. The afterpulses are due to carriers from the original avalanche that become trapped at any defects present in or near the junction, which may be released after the initial current pulse has been quenched and re-trigger a new avalanche.

This interpretation as to the origin of these afterpulse bunches is supported by the fact that when the quench resistor value is increased, the duration of the bunches was observed to clearly reduce, as shown in Fig. 5-17, which shows oscilloscope traces for quench resistor values of 65 and 520 k Ω , with bias and temperature nearly equal otherwise. With a higher quench resistor, the RC time constant determining the voltage recovery on the diode after any given avalanche pulse is increased (the relevant time here would be $R_Q C_s$, with C_s unchanged from 220 pF, giving respectively $RC =$

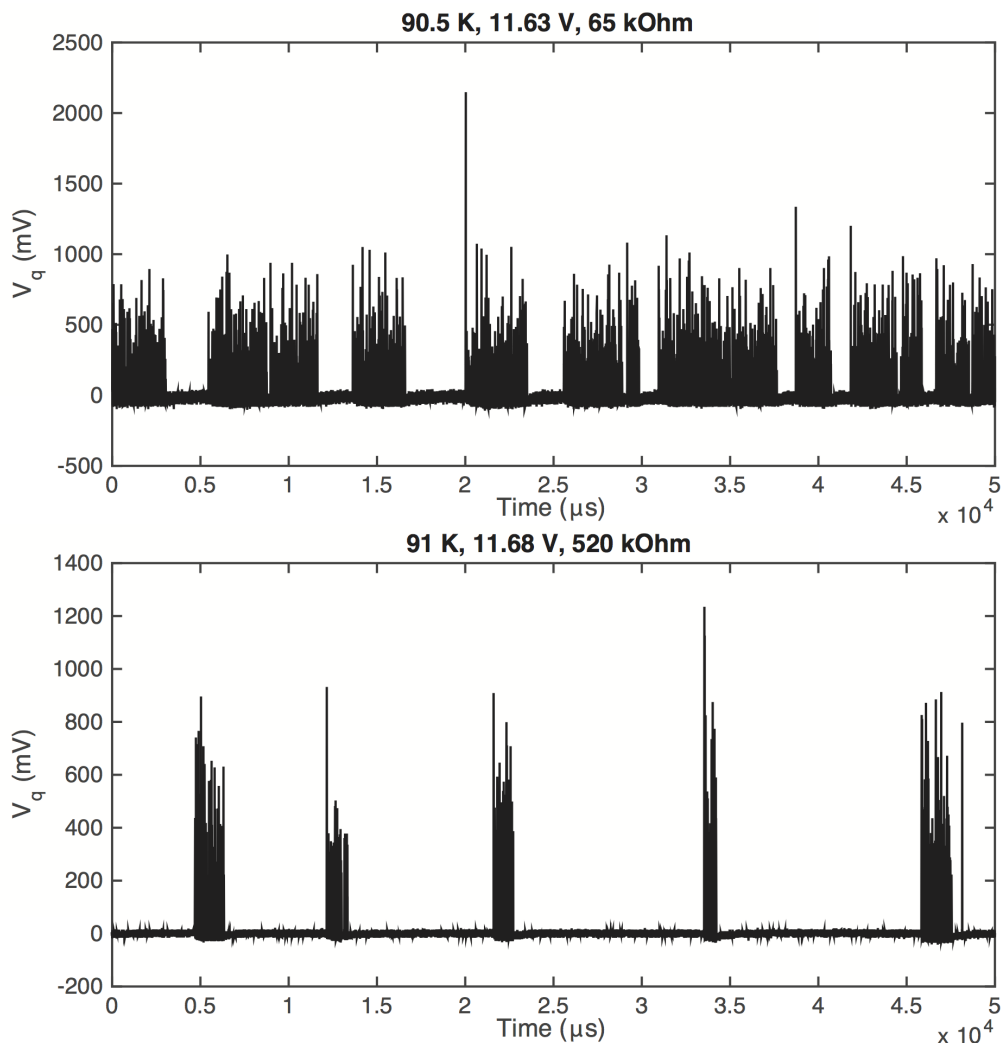


Figure 5-17: Observed afterpulse bunches from dark counts near 90 K and dependence on quench resistance.

14.3 and 114 μ s for these resistor two values). As a result, for a given lifetime of any relevant trap states, the probability that subsequently released carriers generate a new avalanche is reduced since with a larger R_Q they are more likely released when the voltage has not recovered to near the original overbias condition.

These bunches can be seen to last for timescales of 100s of microseconds to few milliseconds depending on the quench resistor value, and in both cases in excess of the voltage recovery time, and the end of the bunch appears to occur when all populated traps happen to evacuate within a voltage recovery time without retriggering a new avalanche. In general this bunching behavior can be expected to occur when the voltage recovery time is comparable to the trap lifetime, and is seen only for a relatively narrow range of temperatures (roughly 20-30 K); when the trap lifetime becomes significantly longer, the bunches are unlikely to ever terminate and the device pulses persistently, as was seen below about 80 K in these devices. Indeed this puls-

ing can become so frequent that the amplitude of the pulses decreases dramatically, and at 4 K for example it was difficult to even observe pulses. On the other hand, when the trap lifetime reduces, the bunches last for less time and eventually are not distinguishable from independent generation events.

Low-temperature afterpulsing has been reported in other contexts in Si SPADs [RLA⁺07], and it is unclear if even with a less defect-ridden GR structure as is obtained with STI GRs [FHE06] it would be more straightforward to operate without pulse bunching at temperatures below about 80 K. Nevertheless, even with the current devices, by implementing an active-quenching circuit with a “hold-off” time of order the trap lifetime, it should be possible to reduce the duration of the pulse bunches at a given temperature, and work on these circuits is in progress as of this writing.

The pulse bunches in the present devices operating with passive quenching by no means rule out the low-level light detection required for state detection. Despite the apparently large dark count rates persisting to low temperatures as in Fig. 5-15, the rate of occurrence of bunches at 90 K corresponds to dark generation rates of order ~ 200 Hz, in the 10 μm -diameter device, and their appears to be enough time during which the device is “silent” that the same devices may still be used at this temperature for low-level light detection for kHz-level photon arrivals, which we discuss below.

Ion fluorescence signal expectation and low-level light measurements

At saturation, a photon is emitted on the $5S_{1/2} \rightarrow 5P_{1/2}$ transition in $^{88}\text{Sr}^+$ every 16 ns, corresponding to a total scattering rate of 63 MHz. Accounting for then the solid angle subtended by the 10 μm -diameter device given the 50 μm trap height ($\eta_c = 0.0025$) and the 80% transmission owing to the metal wire mesh above the APD active region (Fig. 5-8), we expect a photon arrival rate from the ion on the APD of 130 kHz. In practice, the Lincoln group has observed that the actual count rates observed on PMTs, accounting for the collection efficiency of their optics (4%) and the PMT QE (20%) is about a factor of 10 lower than would be expected based on this saturation scattering rate, which may have to do with some dynamics associated with the repumper beams present during readout. While this issue is not fundamental and should be solvable, we would hope to be able to distinguish photon arrivals with rates at the 10-20 kHz level as a result.

Using a fiber feed-through into the cryostat used for the pulsing measurements shown above, we input light corresponding to these levels, using a heavily attenuated 405 nm laser diode. Fiber-guided light was launched into free space through a fiber collimator and after a few inches recoupled into fiber by an opposing collimator, which allowed insertion of free-space ND filters (whose attenuation was independently measured) to reduce the light intensity. Free space attenuation was preferable to fiber attenuation since especially at these short wavelengths, significant variability in connector losses was observed (on the level of 1 dB), which reduces confidence in the total attenuation of independently calibrated attenuators.

Using a combination of ND filters, the total power emitted from the fiber in the vacuum chamber was reduced to 10 fW, corresponding to 20 kHz photon arrivals for the 405 nm light used here. With this light incident on the detector, a scope trace

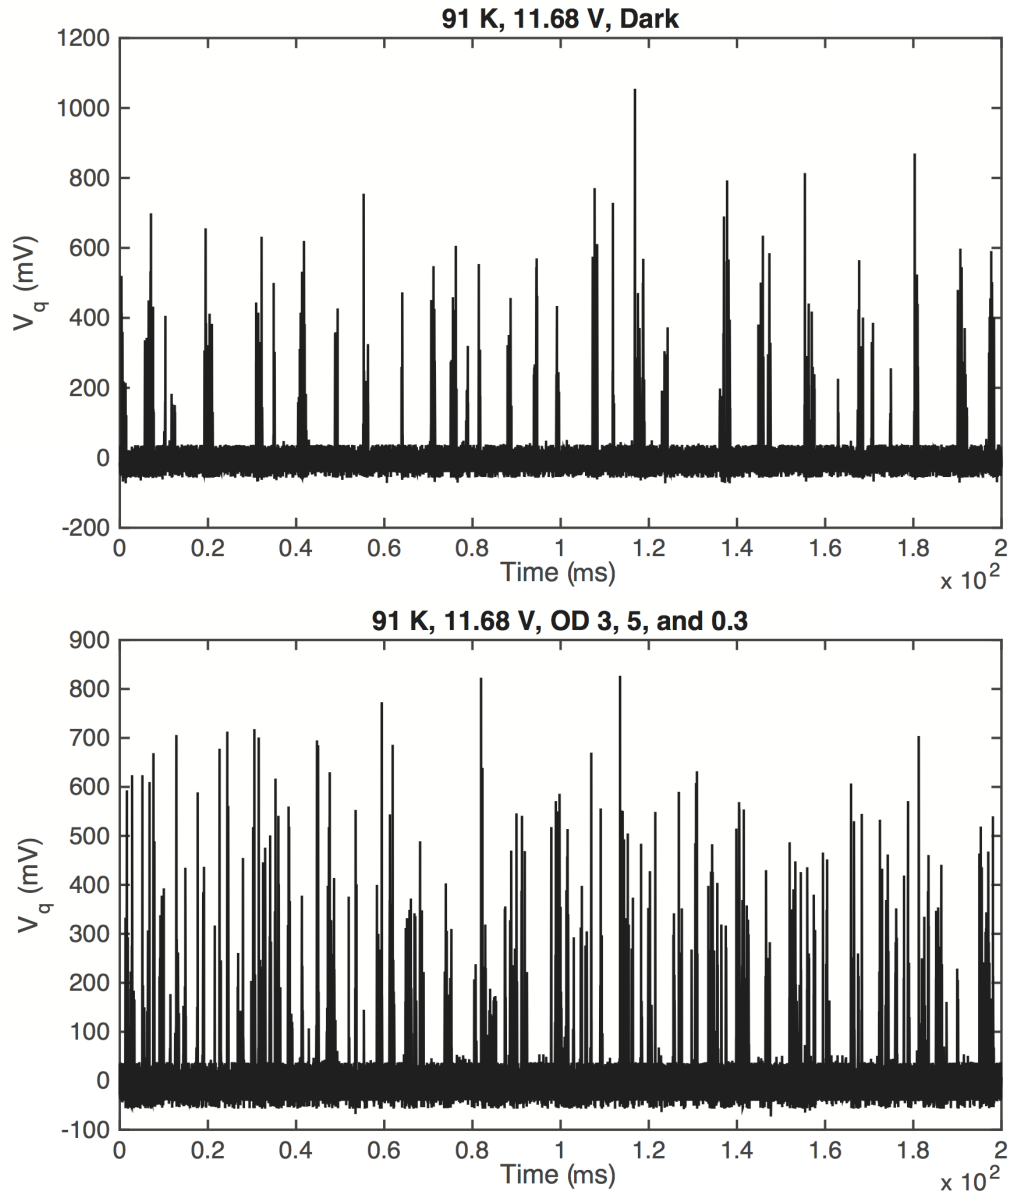


Figure 5-18: Dark counts (top) and response to low light 405 nm illumination at 90 K (bottom panel); the optical power and combination of ND filters used here corresponded to a power and photon arrival rate on the detector of 10 fW, or 20 kHz.

corresponding to the bottom panel in Fig. 5-18 could be observed, in contrast to the dark trace in the top panel, in both cases again using the 520 k Ω quench resistor.

These measurements indicate that even with the reduced scattering rates apparently observed so far in the Lincoln setup, the sensitivity of the present device appears to be sufficient even with passive quenching that at 90 K it should be possible to use the present devices for ion state detection. However, to achieve short readout times and high fidelities, reduction of the afterpulsing behavior, as well as of course achieving the theoretical maximum scattering rate, will be necessary. Afterpulsing behavior can be significantly alleviated with the use of active quenching circuits [CGL⁺96], though again it seems that at low temperatures below ~ 100 K SPADs have not been well explored previously.

5.2.4 Full device characterization

This section briefly describes the early observations made as of this writing on the integrated ion trap/APD device, primarily by the Lincoln group.

Ions have been trapped in the CMOS APD trap, and heating rates of 20 quanta/s were observed with the chip at 8 K for the same parameters as the CMOS trap presented above. The lack of breakdown in this trap suggests rounded electrode corners were not strictly necessary in the earlier device.

Pulses could be observed in the 5 μm -wide device at 4 K, and even with the device biased above breakdown the ion remained stably trapped. Furthermore the heating rate was measured and found to be negligibly affected by the pulsing of the APD (19(2) quanta/s without pulsing, 24(1) quanta/s with pulsing).

The setup used for these measurements had to be modified to accommodate local heating of the chip to allow operation near 90 K. Afterpulsing very similar to what was observed in the cryostat measurements on independent APD devices shown above were observed in the trap-integrated device in the Lincoln setup with the chip at these temperatures, and furthermore the pulsing behavior was seen to be independent of the RF applied to the trap chip; on the scope traces of the pulsing, RF pickup with amplitude 10x lower than the pulse amplitudes was visible, perhaps owing to the exposed wirebonds in the present devices, but the pulsing behavior corresponding to signal readout was unchanged, indicating that even in the presence of the RF trap potentials the APDs should operate as described above.

Characterization of the full trap/APD device is in progress, and ions have now been trapped at 90 K with the APD operating as well. Observations so far suggest the combined trap and APD are able to operate together, and the device's signal to noise ratio is sufficient for ion state discrimination; however, results so far are preliminary and future work will aim to verify that the APD can indeed operate in a way similar to that observed in the independent device measurements above alongside the trap and detect an ion's signal.

5.3 Conclusion and summary

We expect the ability to fabricate ion traps with good performance in standard and fairly accessible CMOS foundries may in the short term make it more feasible for groups to design and conduct experiments with complex, custom-designed multi-layer structures.

In the long run, the ability to integrate optical functionalities, such as the integrated readout explored here, should allow experiments at a scale not yet possible. Though fast and high-fidelity integrated readout is likely to require more than just working detectors to be feasible (structures to quickly bring ions closer to the chip before readout, e.g.), these results indicate that CMOS APDs are a possibility, able to operate alongside the RF trap structures and to achieve sufficient signal to noise ratios at intermediate temperatures. The afterpulsing issues that limit the detectors' performance, even with the present device structure, should be possible to address via use of an active quenching circuit, primarily through the use of a "hold-off" time after an initial pulse during which the bias voltage is brought down below the breakdown voltage [CGL⁺96]. In suppressing the probability that a subsequently released trapped charge triggers another avalanche, this should lower somewhat the minimum temperatures at which the detector can be expected to measure low-light signals.

Also, in future work it should be possible to improve this performance and reproduce the high-performing CMOS SPADs achieved by other groups; since the best have been achieved in CMOS imaging processes with implants tailored to the needs of the APDs, this may benefit from some degree of process customization. However, even given this flexibility, the possibilities for high quality SPAD operation at 4-10 K remain a question, and will be important in assessing the relative merits of the integrability and yield of Si SPADs as compared to the very high performance of individual superconducting nanowire single photon detectors for the readout problem [SVL⁺16].

Chapter 6

Visible electro-optic modulation

The waveguide optics presented in Chs. 3 and 4 were passive, and did not allow switching or encoding of pulses in a parallel fashion – at present this has to be done off chip. While this may still allow interesting systems relying on on-chip electronics to move ions in and out of beams to encode gates [dLM⁺16], integrated modulators would greatly add to the capability of the system.

A number of approaches may be taken to realizing integrated optical modulators within the type of platform developed above. This chapter discusses some of the possibilities, and discusses the features of the hybrid SiN-LiNbO₃ waveguides we have begun to explore. Initial observations on these waveguides both at infrared and visible wavelengths are presented, together with the first integrated electro-optic modulators at visible wavelengths demonstrated in such a platform.

6.1 Approaches to modulation

Typical electro-optic modulators, as used in fiber-optic telecommunications systems, usually utilize waveguide Mach-Zehnder interferometer structures (for intensity modulation) formed in waveguides with relative arm lengths tunable via an electro-optic effect, the dominant material being Lithium Niobate [WKYY⁺00]. These are typically doped-core, low-index contrast waveguide structures (like the fibers themselves) in bulk LiNbO₃, with comparable mode-field diameters to SMF, but allow fast modulation (multiple GHz bandwidths) in cm-length devices.

A number of approaches can be taken to try to integrate electro-optic functionality in general high-index contrast platforms. Bulk silicon exhibits no χ_3 nonlinearity, but in silicon photonics at IR wavelengths, modulation of the free carrier density in the semiconductor can allow tuning of the index via free carrier dispersion [SB87], which has allowed development of modulators based either on Mach-Zehnder structures (e.g. [GRSV07]), or in various more compact resonant devices (e.g. [XSPL05, TSAS⁺14, SOW⁺13]).

In the platform considered in this thesis, the need arises for similar functionality at visible wavelengths. Here, we can indeed chose amongst a few possible core materials which would exhibit electro-optic effects, as summarized in Ch. 1 as well; this include

thin films of GaN [XPR⁺11], AlN, and LiNbO₃ itself. The different materials have relative merits in terms of optical properties and processing, and may play roles in various settings; but given the availability of thin-film LiNbO₃ samples currently [PHSG12], and LiNbO₃'s large electro-optic coefficient (an order of magnitude larger than both GaN and AlN), our work focused on this material.

6.2 Silicon nitride/Lithium niobate hybrid waveguides

Thin films of LiNbO₃ fabricated via processes similar to those used for silicon-on-insulator wafers used in the semiconductor industry have recently become commercially available, and a few groups have worked on high-index contrast photonics in this material recently. No RIE chemistry for LiNbO₃ is known, and in many cases direct argon ion etching is used. Early on these processes were slow and resulted in very rough sidewalls [GPR⁺07], but recently smoother sidewalls in high- Q resonators have been achieved [WBL⁺14] (with the highest Q s of order 10^6 [WBW⁺15]), still with slanted sidewalls rougher than those achieved via RIE which significantly complicates design of integrated devices beyond disk resonators. However it appears these processes may mature further, in which case general waveguide devices written directly into LiNbO₃ may become practical.

Given the relative difficulty of directly patterning LiNbO₃, however, another possible solution is to employ hybrid ridge waveguide structures, in which a separately patterned material partially guides a mode that otherwise is guided in a uniform LiNbO₃ thin film [CXWR14, RMK⁺13]. It may seem that the patterned material should have a higher index than the LiNbO₃, so as to prevent leakage into slab modes radiating away from the waveguide, but it is in fact possible to use lower-index materials as the patterned guiding layer (e.g. SiN), since the requirement truly is that the guided mode of the hybrid waveguide should have a higher index than that of the *slab* modes of the thin-film LiNbO₃, which owing to the overlap with the cladding can be lower than that of LiNbO₃ itself – it need only be larger than the top cladding (likely SiO₂) to ensure this. This approach should enable using the straightforwardly patternable SiN to define interferometer and resonator structures in LiNbO₃, and we argue this should allow interesting electro-optic devices in the visible, which would include intensity and phase modulators as we aim to work towards for the systems discussed in this thesis.

In the last few years a few groups have considered such structures in the IR, using a bonded Si layer as a patterned core material [CXWR14, WSD⁺16, WVAA⁺16]; or more along the lines here, with a chalcogenide glass film deposited on LiNbO₃ and directly patterned [RPC⁺15]. All such structures so far have been studied in the IR.

6.2.1 Simulated waveguide properties

The effective indices of the fundamental mode of a SiO₂-clad LiNbO₃ film, as a function of film thickness, is shown in Fig. 6-1 for both $\lambda=674$ nm and 1550 nm




Figure 6-1: Effective indices of fundamental TE-polarized LiNbO_3 thin film mode, plotted for $\lambda=674$ and 1550 nm, as a function of slab thickness and assuming a SiO_2 cladding.

(the latter because these devices may be of interest in the IR as well), and in Fig. 6-2 a simulated mode profile for a 200 nm-thick SiN film, patterned to a waveguide width of 620 nm, on a 200-nm LiNbO_3 film. The effective index here is 2.01 (whereas the slab mode refractive index is 1.95). The guiding is essentially due to the evanescent field of the LiNbO_3 slab mode, which is more strongly drawn into the SiN than the SiO_2 , allows a higher index mode and hence guiding as a result. Since SiN is straightforward to pattern directly and can indeed be directly deposited and patterned above the LiNbO_3 film (in contrast to the other hybrid waveguides pursued previously [CXWR14, RMK⁺13]), this appears an attractive way to form waveguides despite the relatively low index of SiN.

A drawback of any hybrid waveguide approach and in particular using a lower index patterned material is that the index contrast relevant to bend losses, or also to the efficiencies of grating devices, for example (patterned in the SiN) is that between the ridge-guided mode (here about 2.01) and the slab modes (1.95). This results in significantly higher bend losses, as shown in the calculations in Fig. 6-3 for a few different values of the LiNbO_3 film thickness and for SiN ridge dimensions used in the devices below – these calculations were done using the mode solver with PMLs written by Milos Popovic. The waveguide mode shown in Fig. 6-2 hence requires a bend radius of many 10s of microns, though this is not certainly not prohibitive for modulator devices alone, and we note that if devices of such cross-section prove compelling they could be integrated with more straightforward high-index-contrast waveguides where EO functionality is not required. As shown in Fig. 6-3, this bend loss is also a strong function of the LiNbO_3 thickness.

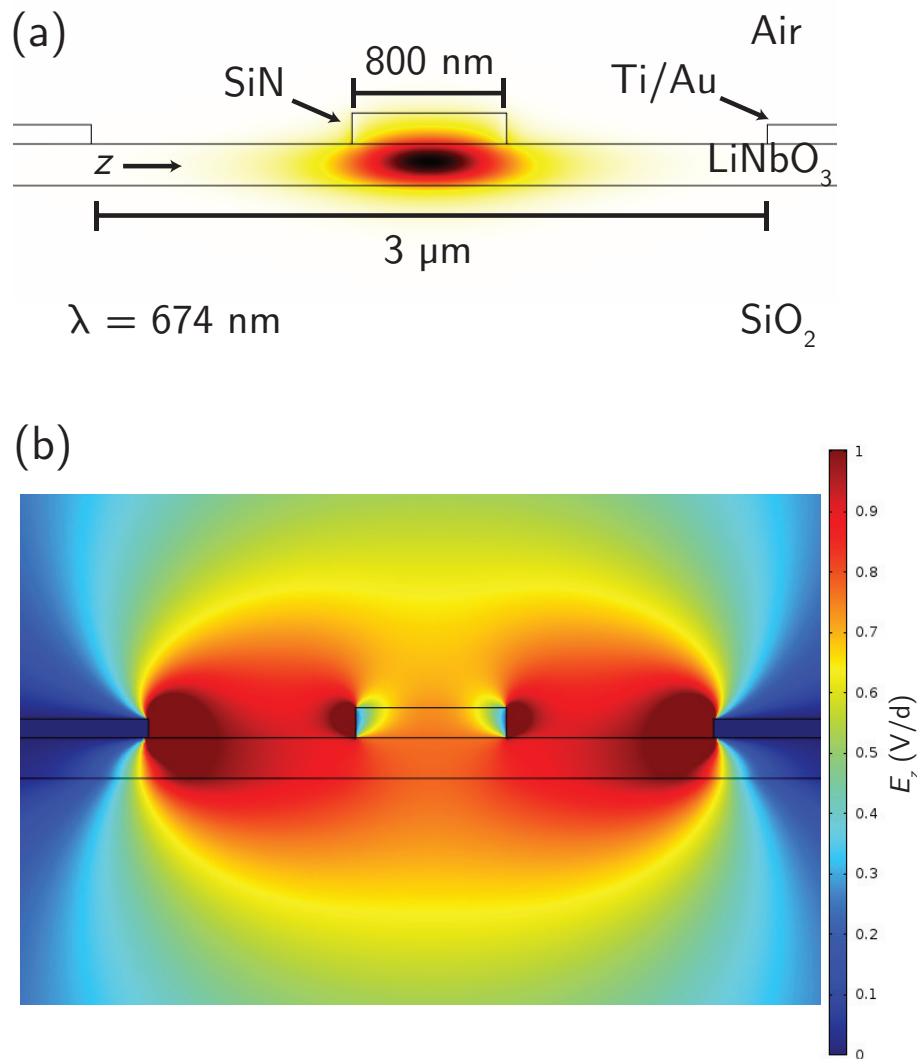


Figure 6-2: (a) Simulated mode profile at $\lambda = 674$ nm of a hybrid SiN-LiNbO₃ waveguide, with 160 nm SiN thickness with 215 nm LiNbO₃, and 800 nm ridge width. Here z refers to the direction horizontal in the page and aligned with LiNbO₃'s axis of max EO coefficient (r_{33}). (b) Static simulation of electric field applied across the electrodes, showing the horizontally-directed component as a fraction of V/d , with $d = 3$ μm. Field in the LiNbO₃ is sensitive to the DC permittivity of the SiN ridge, here taken to be $\epsilon_r = 7$.




Figure 6-3: Simulated bend losses at 674 nm for fundamental modes of various SiN dimensions above a LiNbO₃ uniform film, for an 800 nm-wide waveguide in a 160 nm-thick SiN layer ($n = 1.95$) on a LiNbO₃ film of various thicknesses. The 215 nm thickness is the approximate realized value in the devices fabricated and characterized experimentally below.

6.2.2 Approximate expected electro-optic properties and modulator designs

To get an approximate feel for the expected EO properties of such waveguides, we consider an x -cut thin film. This is defined such that LiNbO₃'s extraordinary optical axis, by convention referred to as the z -axis and the one with the largest EO coefficient, is oriented parallel to the film. Electrodes are on the surface of the LiNbO₃ separated by some distance d_e between which a voltage V is applied to tune the index, as shown in Fig. 6-2. In this case the propagation direction is perpendicular to the z -axis, and the E-field (both the optical mode's field and that from the voltage applied to the electrodes) is aligned with the crystal's z -axis.

The relative permittivity for an optical wave polarized along this axis by a DC field applied along the same axis is

$$\epsilon_r = n_e^2 - n_e^4 r_{33} E_0, \quad (6.1)$$

where $n_e \approx 2.2$ is the extraordinary index at 674 nm, $E_0 \approx V/d_e$ is the applied field magnitude, and r_{33} is the EO tensor element describing the effect along this axis, which has a value of approximately 31 pm/V in LiNbO₃ [Tur66].

For a small modification of the permittivity the change in index is approximately

$$\Delta n_e \approx \frac{\Delta \epsilon_r}{2\sqrt{\epsilon_{r0}}} = -\frac{n_e^3 r_{33} E_0}{2} \quad (6.2)$$

The phase shift experienced by light at wavelength λ over some propagation length L , assuming the waveguide index is close to the material index (true for high confinement waveguides), is

$$\Delta \phi = k_0 L \Delta n_e, \quad (6.3)$$

Calling V_π the voltage required to effect $\Delta \phi = \pi$ as required for full intensity modulation in a Mach-Zehnder configuration when one arm is tuned, the product $V_\pi L$, which describes the tradeoff between device length and necessary voltage for modulation, is

$$V_\pi L = \frac{2\pi d_e}{k_0 n_e^3 r_{33}} \frac{1}{\Gamma_{\text{opt}} \Gamma_{\text{elec}}}, \quad (6.4)$$

where Γ_{opt} is the optical mode's confinement factor in the LiNbO₃ and Γ_{elec} is a factor scaling the field seen by the mode; i.e. the relevant field is $E = \Gamma_{\text{elec}} V/d_e$. As seen in Fig. 6-2(b), owing to electrostatic boundary conditions around the SiN ridge, this results in a lowering of the field seen by the optical mode that increases with the difference between the static permittivity of the SiN core relative and that of the cladding.

Setting both these factors to 1, for $d_e = 2.5 \mu\text{m}$ comes to 0.5 V-cm, using $r_{33} = 32$ pm/V at 674 nm. This indicates that for a 1 mm-long MZ device, a 5 V swing on one arm should be expected to switch between the bright and dark state at the output waveguide (and if both arms are modulated in a push pull configuration such that only a $\pi/2$ shift is required on either one, this voltage is further reduced by a factor

of 2).

This gives a sense of expected scale of MZ devices based on such waveguides (as well as the effects of lower EO coefficients or more distant electrodes, etc.); device lengths can be further reduced by employing resonant devices as mentioned above, but reaching many orders of magnitude in extinction appears more challenging at first with such approaches, though some clever schemes may be envisioned eventually [MOR13]. 40 dB-level extinctions with MZ structures have been achieved [HSM⁺15], though few groups appear to have studied some of the various approaches to achieve very high extinctions, in the 60-80 dB range as would be desirable here (aside from cascading more than one modulator in sequence) in modulators of any architecture so far.

6.3 Fabrication

A sample of *x*-cut, 215 nm-thick (on average – the specification was 200 nm, as in the designs above, but the delivered film was with this slightly larger thickness) LiNbO₃, on a 2 μm-thick oxide layer above a bulk LiNbO₃ was provided by the company NANOLN. The 3-inch wafer supplied was diced into smaller rectangular pieces for processing; PECVD deposition using the same Oxford-100 system used for the SiN waveguides above was performed with the chip at the standard 350° temperature, and no damage to the piece from stress or fracture were observed. 15 minutes of deposition were observed to result in a 165 nm-thick film on a silicon monitor piece, and 18 minutes was used for the sample discussed below. The same HSQ resist and E-Spacer solution as described in Ch. 3 was used for e-beam lithography, which went smoothly – the use of LiNbO₃ substrates presented no obvious difficulties in these respects.

6.3.1 Electron-beam proximity effect on LiNbO₃

In writing structures to ascertain the necessary dosages, since a 200 nm-thick layer of nitride would have little impact on electron backscattering given that the range is in the 10s of microns, these dose calibration structures were attempted on pieces of the wafer which did not have SiN deposited. However, we found that adhesion of the HSQ directly to the LiNbO₃ was poor, and many waveguide and grating line features had lifted off after development. This held true both for a sample simply rinsed in solvents prior to resist spinning and one that was piranha cleaned. No issues with adhesion were observed when the resist was spun onto a SiN layer above the LiNbO₃ film.

An issue that arises in writing on LiNbO₃ substrates is that, as the mass density of the material is almost a factor of two higher than that of Si, electrons are backscattered from within the substrate much more strongly and hence the electron range is lower. To quantify the strength of this, a commercial software package (TRACER, running on the MTL computers used for e-beam lithography) was used to carry out Monte-Carlo simulations of electron trajectories in the substrate, with results showing




Figure 6-4: Electron range plots for Si and LiNbO₃ substrates, respectively. Radius on the x -axis refers to the distance from the beam center along the surface of the substrate, i.e. perpendicular to the electron beam.

energy density delivered to the resist as a function of radius from the beam center plotted in Fig. 6-4. Electrons on LiNbO₃ are clearly scattered more strongly within shorter radii, and at approximately 10 μm the delivered dosage for a given beam current is approximately a factor of 10 higher than that when using a Si substrate.

This results in the fact that, around a dense feature like the tapers to the grating couplers, within about 10 μm of the boundary of such features, a relatively high dosage is delivered to the surrounding areas not intended to be exposed. As a result, and owing to the finite contrast of the HSQ used and as developed by the NaOH/NaCl mixture used throughout this thesis, some residue of partially-exposed resist was visible around the optimally exposed samples (SEM images in Fig. 6-5, and even for features underexposed by nearly a factor of two.

We note in addition that owing to this larger proximity effect at length-scales close to those of the gratings and waveguides, automatic proximity effect correction (PEC) was used on LiNbO₃ writes; this can be carried out in the BEAMER software used to define the e-beam writes, and weighs dosages in features of varying density to equalize delivered doses to different parts of the pattern. This helped expose features uniformly, but had a negligible impact on the residue around dense features. This residue however appears to have had a minimal effect on the optical performance, as presented below.

All of these issues would be essentially eliminated by using a Si substrate with a thin film of LiNbO₃; when the sample used in this chapter was purchased, NanoLN was not supplying such samples, but they are available now and would simplify some

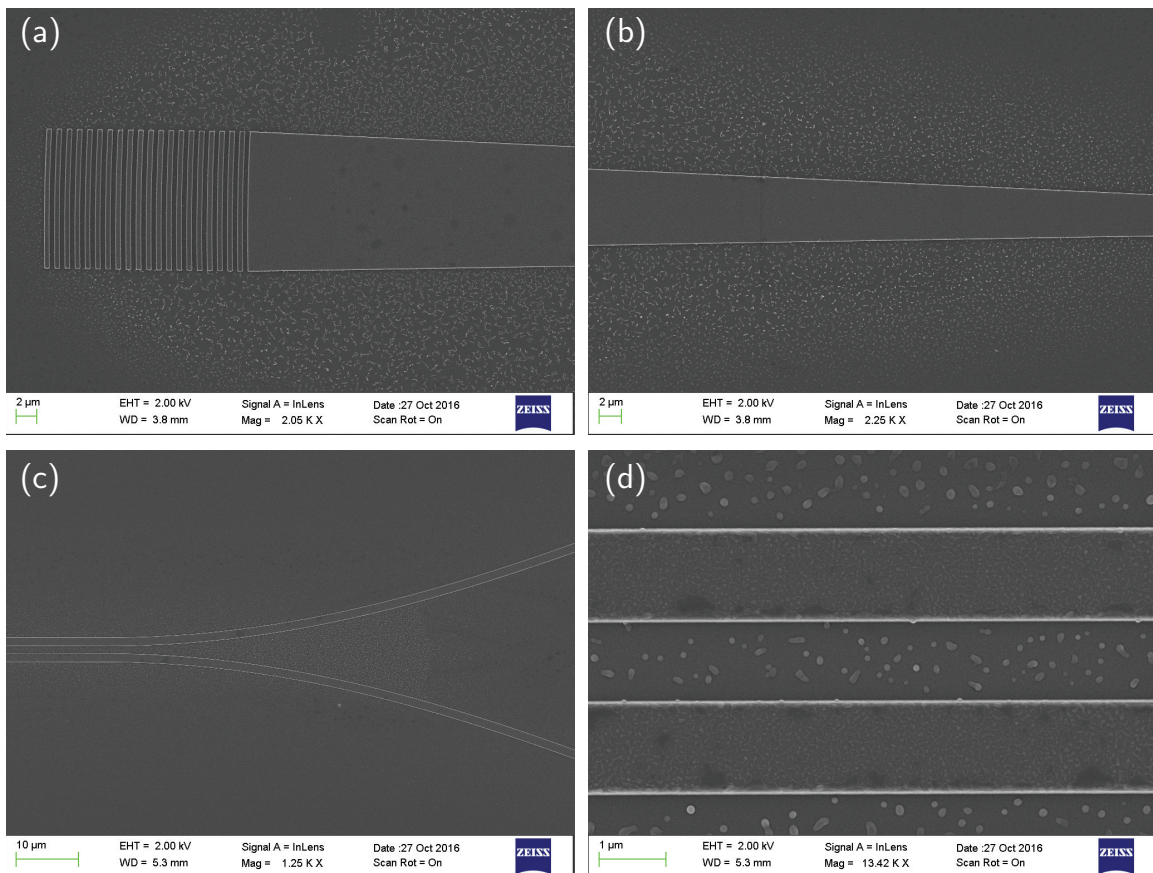


Figure 6-5: SEMs of written features showing artifacts from proximity effect. These artifacts are visible, for example in (d) as the ~ 50 nm-scale features around the waveguides.




Figure 6-6: Mode profiles at 674 nm on the 215 nm-thick film, with 175 nm SiN deposited above and patterned to 700 nm and 1.2 μm widths for SM WGs at 674 nm and 1550 nm, respectively.

of this processing.

6.4 Sample 1: uncontacted waveguide devices

A first chip was designed and fabricated to understand passive waveguides in this platform. This section describes the designs and device characterization performed on this first chip; following this a second chip was designed and fabricated with electrodes aligned and patterned as well, as discussed in the following section.

6.4.1 Designs included and brief background on ring resonators

A variety of designs were included on the chip written, including waveguide loss test structures (fiber couplers with varying lengths of SM waveguide between) for both 674 nm and 1550 nm. Test structures for directional couplers intended to function as 50:50 splitters, for use in Mach-Zehnder Interferometer (MZI) structures, were included as well at both wavelengths. At 1550 nm, where commercial tunable lasers across a wide range (~ 100 nm) and mode-hop free are available, ring resonator test-structures were included as well, together with full MZI structures, as shown in the optical micrograph of Fig. 6-7. These rings were chosen conservatively with a radius of 150 μm , and the couplers are offset vertically to prevent any light coupled into unconfined slab modes of the thin film (as opposed to the ridge modes guided by the SiN) from appearing at the output for the ring resonator structures. Both the ring resonators and MZIs are of interest as elements in electro-optic modulators; rings, or resonant devices generally offer a path to compact modulators, and in this context also are convenient for characterizing waveguide loss.

Fiber grating couplers

Fiber couplers on these devices were designed with a constant period and duty cycle set to 50%; the strongest possible grating was desirable to maximize the light emitted,




Figure 6-7: Optical micrograph of fabricated ring resonator and MZI devices on LiNbO_3 ; rings in the image have diameters of $300 \mu\text{m}$.

because the index contrast is too low to emit the majority of the guided power within a length corresponding to the fiber mode diameter. The period was chosen to emit at the desired angle (usually 15%), and the width of the grating was chosen according to the specified mode field diameter of the optical fiber according to the mode-matching guidelines described in Ch. 3. The taper was set to be $200 \mu\text{m}$ long, simply as a conservative choice ensuring little leakage into other slab modes.

Since much of the mode, especially for the 674 nm modes, resides in the LiNbO_3 , the perturbation effected by etching a grating into the SiN is not as large as in a simpler waveguide structure. Nevertheless, for the film thicknesses here, 2D cross-section simulations indicate that the gratings at 674 nm, over $6 \mu\text{m}$ (approximately the fiber mode field diameter) emit upwards 20% of the power incident on the grating from the waveguide, indicating 20% as an upper bound on these grating efficiencies (since we were not optimizing these closely, full modal overlap simulations were not performed). As described below in the discussion of sample 2, the realized coupling losses at 674 nm were approximately -9 dB (corresponding to 13% efficiency).

Ring resonator background

To give some background for the results shown below, we briefly describe here the relevant characteristics of ring resonators. The rings support resonant modes at wavelengths where the round-trip length is a multiple of the guided wavelength. Such resonances are observed via dips in the transmission from input to output centered on the resonant frequencies. The line shapes of these dips can be described by a coupled-

mode theory, derived in a simple form used here in [Meh12]. The transmission here is mathematically equivalent to the *reflection* from a Fabry-Perot cavity with light input through a partially-transmitting mirror and the opposing mirror perfectly reflecting, with some internal loss; this analogous reflection can be written as:

$$R = \frac{1/\tau_0^2 + (\omega - \omega_0)^2}{\left(\frac{1}{\tau_0} + \frac{1}{\tau_e}\right)^2 + (\omega - \omega_0)^2}, \quad (6.5)$$

where τ_0 is the “intrinsic” loss time constant (for the field amplitude, and owing to loss due to waveguide loss, in this case; this would limit the ring Q in the absence of any coupling to the bus WG mode), and τ_e is the “extrinsic” loss time constant reflecting the coupling to the bus WG mode (or due to the partially reflecting mirror in the FP case). Each time constant is associated with a $Q_{0,e} = \omega_0\tau_{0,e}/2$ (with $\omega_0 = 2\pi/\lambda_0$ the resonant angular frequency), the combination of which gives the total “loaded” quality factor $Q_{\text{tot}}^{-1} = Q_0^{-1} + Q_e^{-1}$. When the the intrinsic loss rate equals the loss rate owing to coupling to the bus waveguide, the light directly transmitted through the bus waveguide interferes destructively with that coupled out of the ring in steady state, and the transmission from input to output approaches 0 on resonance; this is known as “critical coupling”, and is the condition one would hope to approach for high-extinction modulation (though the amount of extinction is very sensitive to the relative loss rates and for robust high-extinction modulation other structures may be preferable).

The coupling between the bus waveguide and the ring is controlled by the minimum gap g between these waveguides, and is mediated by the overlap between the exponentially-decaying evanescent fields extending away from the guiding core in the LiNbO₃ slab. Owing to the exponential decay of this field, except for extremely small gaps, the coupling strength ($\propto 1/\tau_e$) depends exponentially on the inverse of this g ; hence sweeping g approximately linearly (as was done here, between 100 nm and about 2.9 μm over 15 variants) allows the possibility for reaching close to critical coupling in at least one fabricated structure for a wide range of possible waveguide losses, varying linearly on a dB/cm scale (exponentially in $1/\tau_0$).

Finally, we briefly note that the intrinsic Q_0 can be related to the propagation loss in the ring (given by the field decay constant α , defined such that intensity decays proportional to $e^{-2\alpha L}$). The group index of the guided mode relates the decay per unit length to decay per unit time, allowing us to write $Q_0 = 2 \cdot \frac{\omega_0}{\alpha} \frac{n_{\text{gr}}}{c} = 2k_0 n_{\text{gr}}/\alpha$.

The bend radius for the 674 nm test structures, owing to the lower confinement in the SiN (mode profiles for both wavelengths for the dimensions employed in design are shown in Fig. 6-6) and given the realized LiNbO₃ film thickness of 215 nm, was chosen to be 300 μm ; the required radius is a very strong function of the film thickness, and could be significantly reduced if a thinner film was used, which should be done in the future for lower wavelengths.

6.4.2 Characterization of IR ring resonators and MZIs

Immediately after fabrication, which ended with an etch of the SiN, leaving the waveguides unclad, the ring resonators for 1550 nm-wavelength appeared to have low Q s, and very weak transmission. Clear resonances could only be observed with the smallest gap included of about 100 nm, resulting in a transmission spectrum around 1510 nm of that shown in the bottom trace in Fig. 6-8. The approximate Q s indicated by the linewidths, as well as the transmission maxima between resonances (the input/output couplers are spaced by about 1 mm) suggest a waveguide loss on the order of a few 100 dB/cm.

This extremely high loss was found to drastically decrease with moderate temperature treatment. A 1 minute bake on a 300° C hotplate in ambient significantly reduced the loss (see the black trace in Fig. 6-8), which was then further decreased after a 10 minute bake at 340° C (blue traces in Fig. 6-8); the linewidths of these resonances indicate an intrinsic quality factor Q_0 of order 100,000, corresponding to a waveguide loss of 3 dB/cm near 1510 nm.

The significant reduction in loss observed with temperature treatment up to 340° C is unlikely to be due to any effect of the temperature on the SiN, since this material was deposited at 350°. The other possibilities are that either some severe roughness in the remaining HSQ on the sample was resulting in the loss, or else there may be some effect in the LiNbO₃ film under the nitride, where the electron-beam was incident during lithography, which results in excess loss. A 5 s dip in buffered oxide etchant was performed to etch away the residual HSQ (BOE is known to etch HSQ resist prior to high temperature baking around 350° C significantly faster than SiO₂), but little effect on Q was observed after this etch; a fit of a resonance profile of a mode near 1515 nm is shown in the red points and curve in Fig. 6-9. Hence the most likely explanation is that the dominant contribution to the very high loss initially observed was from some damage, or built up free charge in the LiNbO₃ was contributing to loss that was at least partially annealed out during the hotplate treatment.

After a further 70 h anneal at 350° C, Q s were not significantly affected, but it was observed that at higher wavelengths (near 1550 nm) Q s were higher (the angle from normal of the coupling fibers was reduced to achieve peak coupling at a longer wavelengths). Finally, after a 400° bake, the highest Q s were observed around 1590 nm of 350,000 (blue curves in Fig. 6-9). The fact that the Q s are generally higher at larger wavelengths may be related to hydrogen absorption in SiN, which is known to result in an absorption peak around 1520 nm [HKL⁺87].

Though no electrical contacts were included on this sample, we verified tuning by simply landing DC probes on opposing sides of the ring resonators, as in the micrograph in Fig. 6-14. The spacing was roughly 30 μ m, and the transmission spectra for various voltages applied are shown in the top panel of the same figure. The linearity of tuning of the center wavelength with voltage (Fig. reffig:ringtuning, bottom right panel) indicates that the EO effect is responsible. The probes here were landed in such a way that the field was applied primarily along the LiNbO₃'s z axis, i.e. the one with the largest EO coefficient (r_{33}). The observed tuning indicated that the waveguides were behaving roughly as desired, and motivated fabrication of




Figure 6-8: Ring resonator transmission spectra after anneals. Rings 12, 11, and 9 had minimum ring-bus coupling gaps of 100 nm, 300 nm, and 707 nm, respectively.




Figure 6-9: Measured spectra of ring resonators at IR wavelengths labeled and after annealed conditions as labeled, together with fits of the form described in the text. Rings 9, 8, and 7 had minimum ring-bus coupling gaps of 707 nm, 918 nm, and 1140 nm, respectively.

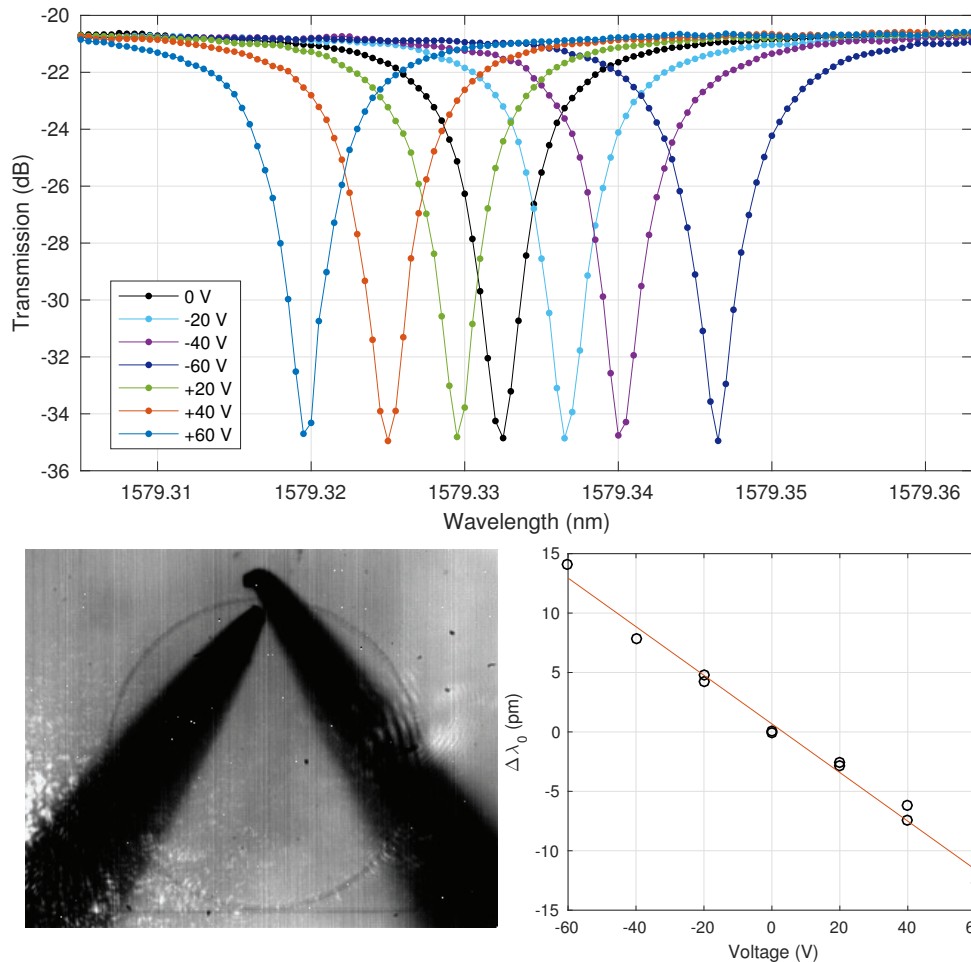


Figure 6-10: Transmission spectra of a high- Q ring as a function of applied biases across probes landed on an otherwise uncontacted sample; probe positions are as shown in the micrograph of the device in the lower left panel. From fits to the measured spectra, an approximately linear dependence of the center resonance wavelength on the applied voltage is inferred.

a second sample (following section) with integrated contacts to enable lower-voltage and faster tuning.

6.4.3 Visible waveguides

Stable guiding of 674 nm light could be observed through the straight WG test structures written for these wavelengths, with approximately $10 \mu\text{W}$ coupled into the guided mode. Measuring the fiber-fiber power transmission as a function of waveguide length gives an estimate of the waveguide loss as 4 dB/cm (Fig. 6-11).

In $\lambda = 674 \text{ nm}$ test-structures with bends, significant bending loss of around 1 dB per bend was observed. From simulations, a loss of 0.1-0.2 dB/bend was expected (for the 215 nm-thick LiNbO_3 film, and for $250 \mu\text{m}$ -radius bends); given sensitivity to film thickness, variations across the wafer could account for the discrepancy. Thinner films clearly enable tighter bends, as indicated in Fig. 6-3, which will be desirable for more robust and smaller-scale devices in the visible.

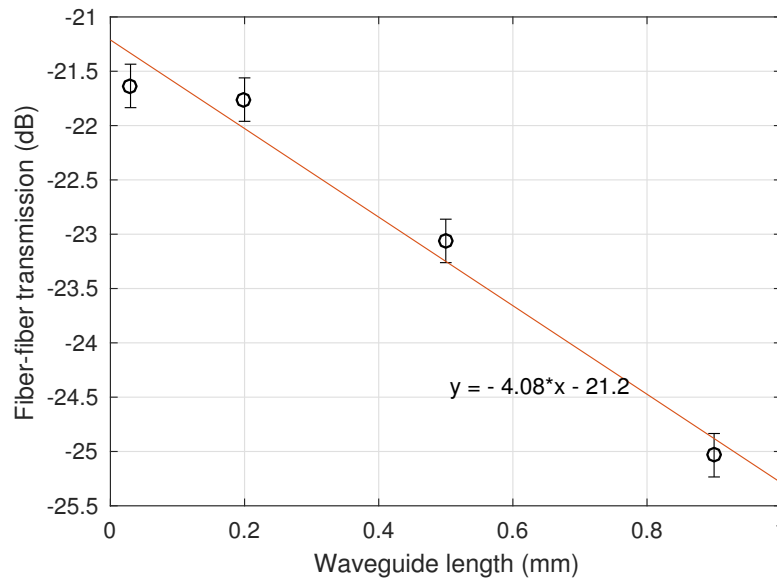


Figure 6-11: Waveguide loss measurement at 674 nm, indicating a loss of 4 ± 2 dB/cm with uncertainty reflecting 95% confidence intervals on the slope.

6.5 Sample 2: contacted MZ and ring devices

Given the functioning devices on the first sample and the knowledge gained about the waveguides in this platform, a second was patterned with provisions for aligned electrodes. Here we describe the fabrication of this sample followed by characterization of tunable structures at both 674 and 1550 nm.




Figure 6-12: Alignment mark written in SiN layer used for registration of the contacts; lines forming the outer cross are $5\ \mu\text{m}$ -wide and extend for $500\ \mu\text{m}$ from center, and the lines of the inner cross are $500\ \text{nm}$ -wide.

6.5.1 Fabrication

Here the silicon nitride fabrication proceeded the same as previously, except that the realized SiN thickness here was $160\ \text{nm}$. Four alignment marks written in this layer at each corner of the piece. Following etching of the SiN, the chip was baked at $350\ \text{C}$ for 1 hour and briefly optically tested. Then about $650\ \text{nm}$ of PMMA was spun on (7% PMMA in anisole) at 2500 rpm and baked for 90 s at $180^\circ\ \text{C}$ to give a thickness of about $650\ \text{nm}$, and exposed to pattern contacts around the waveguides. PMMA was chosen for this step since as a positive resist it lends itself to lift-off definition of electrodes (only the electrode area need be exposed); additionally it is capable of spinning layers a few microns thick, much thicker than HSQ which can typically be spun to thicknesses up to $\sim 250\ \text{nm}$. E-spacer was spun over this layer as usual in addition. The PMMA was then exposed (at a much lower base dosage, at approximately $420\ \mu\text{C}/\text{cm}^2$ at the center of electrode regions (which extended over larger dimensions than the $10\ \mu\text{m}$ electron range). This sample was developed in a mixture of 2:1 IPA:MIBK for 135 s, followed by immersion in IPA for 60 s and subsequent rinsing in IPA. Electrodes were evaporated with a $10\ \text{nm}$ Ti adhesion layer first, followed by $90\ \text{nm}$ of Au, after which lift-off was performed in NMP at $100^\circ\ \text{C}$. After about an hour soaking in the NMP, the sample was removed and the unexposed PMMA could be easily rinsed off with flowing NMP, leaving contacts.

Imaging the $120\ \text{nm}$ -thick SiN registration marks written in the first lithography step, after deposition of the thick PMMA was possible but contrast was difficult to achieve, and artifacts from charging (perhaps due to the E-spacer) made it difficult to image the alignment marks with the SEM for more than a few minutes; nevertheless alignment errors of about $100\ \text{nm}$ were achieved on the present sample. The combination of large and fine features in the alignment mark used (Fig 6-12) was essential to locating the features.

The need to perform this alignment step and uncertainty about imaging the alignment marks was the primary reason no top-cladding was deposited on this sample; SiO_2 could be used as the cladding layer (and as discussed below is likely necessary for future devices), but a $\sim\mu\text{m}$ -thick SiO_2 layer would increase the difficulty of imaging

the alignment marks in the SiN. This will be discussed briefly again at the conclusion of this section.

6.5.2 674 nm Mach-Zehnder modulator

Fig. 6-13 shows a micrograph of a MZ modulator connected in a push-pull configuration (the EO axis is vertical in this image), with the outer two electrodes shorted (owing to absence of a top-cladding, waveguide crossings of the electrodes were avoided); the waveguide and electrode dimensions relevant to this were shown in Fig. 6-2. With 54 μ W incident on the input coupler, and inputting and outputting at the ports as labeled on the micrograph, the voltage-dependent transmission is plotted in Fig. 6-13b, indicating a low V_π of 3.0 V near 0 V bias (and slightly lower at higher bias points), and extinctions of \sim 20 dB, which could be limited either by nonuniformities in the couplers or differential loss in the two arms. However, the fiber coupler loss was approximately 9 dB/coupler on this sample, indicating, given the peak transmissions, an insertion loss on-chip of approximately 9 dB. The bulk of this was due to the fact that the codirectional couplers, intended to be 50/50 splitters, appeared to be significantly off in power coupling (significantly more light was output at the top port as labeled in the micrograph, than at the input in these devices); the coupling constant is very sensitive to both the realized gap between the two co-propagating waveguides and the LiNbO₃ thickness in this region. Nevertheless, on opposite ports, extinction is maintained in spite of non-ideal splitting ratio, at the cost of insertion loss.

The extinction on the ports measured would be in principle infinite if the splitters, and the losses in each arm, were identical. The use of compensation of any imbalance between the arms can be used to increase this extinction [KKY⁺07], as can simply cascading multiple MZIs, though of course only if the insertion loss is manageable.

We note that to observe these DC transmission characteristics, devices were baked again on a hotplate at 350° C following the contact lift-off step. Prior to this bake, the DC tuning at low voltage was observed to be much lower, and a leakage current of a few 10 nA was observed to flow for \sim 20 V biases applied. This current was eliminated following the hotplate bake, again suggesting the influence of trapped charges perhaps owing to the e-beam lithography.

Considering the optical confinement factor in the LiNbO₃ (82% simulated) and the electric field within the LiNbO₃ (\sim 80% of V/d), a somewhat higher V_π of 4.5 V would have been expected, based on an EO coefficient of $r_{33} = 31$ pm/V [Tur66]. The reason for this apparently stronger than expected tuning is at present unclear, as is the reason for the reduced V_π at higher voltages, evident in Fig. 6-13(a).

To verify that this tuning was not due to some slow drifts in the device, dynamic modulation on the heat-treated sample was verified by applying a \pm 1.5 V square wave to the device at 10 kHz and monitoring the output with a photodiode (Fig. 6-13c). Owing to the signal level here, a 1 M Ω load was used to obtain detectable voltage, which given a load capacitance of approximately 20 pF limits the rise/fall times to about 20 μ s, consistent with the observed waveform. This verifies that the dynamics responsible for the tuning are operating on timescales at least this short. Together

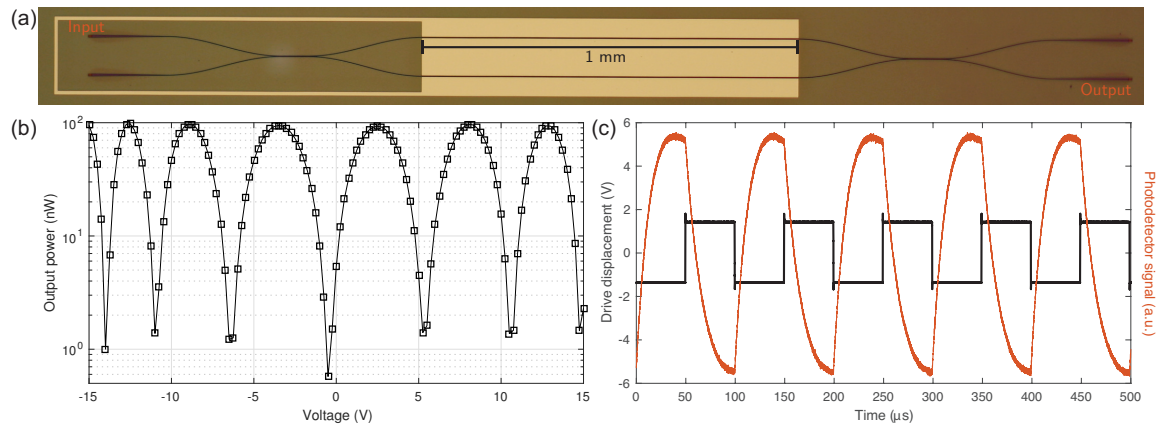


Figure 6-13: (a) Optical micrograph of MZ modulator operating at 674 nm. (b) MZ transfer function, exhibiting $V_\pi = 3$ V near 0 bias and ~ 20 dB extinction, and (c) Low-frequency modulation response to a ± 1.5 V 10 kHz square-wave drive signal; bandwidth limited in the setup used here by the scope input impedance ($1\text{ M}\Omega$) and capacitance.

with the approximate agreement between expected and observed V_π at low voltage, this increases confidence that the EO effect is in fact responsible for the modulation.

The reason for this apparently stronger than expected tuning is at present unclear. Along some axes, at low frequencies LiNbO_3 can have larger EO coefficients due to piezoelectric strain effects which do not impact the high-frequency performance; although such effects are understood to be minimal in bulk samples along the z axis as used here [Tur66], perhaps the effect of interfaces in the thin-film sample introduces some field-induced strain in this case. To rule out the role of any additional trapped charges or defects from e-beam lithography, it would be interesting to study devices annealed at higher temperatures before the contact definition, and devices with contacts patterned photolithographically instead of with e-beam lithography.

Higher-speed measurements would be more conclusive on these fronts, and we should expect modulation bandwidths in the GHz range. Such measurements were attempted using a microwave spectrum analyzer to measure weak RF tones at the modulation frequency with a $50\ \Omega$ load; however, these measurements were precluded by the fact that after some time under the AC field, electrical breakdown was observed across the center and one of the outer electrodes, which manifested as visible damage to the electrodes and waveguides, following which given the large resulting differential loss in the modulator arms very little modulation was possible. This is most likely due to the lack of a cladding and breakdown through air.

Approximately $5\ \mu\text{W}$ of power were in-coupled, resulting in guided intensities of order $5\ \text{kW}/\text{cm}^2$ in the waveguide. Despite the potential for photorefractive damage in LiNbO_3 , stable DC transmission characteristics were observed over at least 20 minutes; slight shifts in the voltage (~ 0.1 V could be observed over this time, but no resolvable drop in the peak transmission. Thus, despite the limitations of this first device and setup and the potential for improvements on this preliminary result, the




Figure 6-14: Micrograph of a contacted ring resonator device (contact spacing $5 \mu\text{m}$), and transmission spectra for $0, \pm 5 \text{ V}$ applied via DC probes. Fits to each measured spectrum (plotted as the lines) indicate an intrinsic Q of 140,000 and a loaded Q of 100,000. Minimum ring-bus coupling gap here was $1.24 \mu\text{m}$ (in contrast to the devices in Fig. 6-9 and Fig. 6-14, here the coupling occurs with the field polarized along the ordinary axis, with higher index, resulting in stronger coupling for a given gap).

low V_π achievable and extinctions achievable even with unbalanced splitters indicates significant promise for visible wavelength modulation in such a platform.

6.5.3 EO tunable infrared rings

Finally, on the second sample, rings designed for IR operation were included as well, to demonstrate more efficient tuning than could be achieved with the DC probes as used for tuning the uncontacted sample. On this chip, the Q s achieved after hotplate bakes even before contact deposition were somewhat lower than on the first sample; the HSQ resist used on this sample had just reached its expiration date, and occasional defects around the waveguides were observable under SEM, which may have played a role. However, intrinsic Q s of 150,000 were achieved with contacts, and

rings with contacts spaced by $5\ \mu\text{m}$ and around approximately 1/5th of the ring allow tuning across more than a linewidth with only 5 V applied (Fig. 6-14). The plotted transmission is normalized to a coupling loss of -10.5 dB/coupler (which was -8 dB at the peak coupling wavelength near 1550 nm).

The data indicate that with this electrode configuration, a voltage swing 3 V is sufficient to tune one linewidth of the current rings. This tuning voltage could be trivially reduced by a factor of two by contacting the bottom half of the device (not done here as this would have required wiring over the ring waveguide to achieve a shift of the same sign on both halves of the ring if only two contacts were to be made). Of course as the Q s increase the necessary voltage is reduced as well. These results in all simply indicate the practicality of modulation and general electro-optic interactions in the IR, also with low voltages.

6.6 Conclusion

The work presented in this chapter takes first steps towards implementing integrated modulators operating in the visible and potentially capable of interfacing with the waveguide optics considered in earlier chapters.

We build on previous results on similar hybrid waveguides in LiNbO_3 [RPC⁺15] by demonstrating operation at visible wavelengths, which have allowed MZ modulators with low-frequency $V_\pi L$ products of 0.3 V·cm and extinctions of 20 dB. Additionally the waveguide losses in the IR attainable after heat-treatment have enabled quality factors up to 340,000, roughly a factor of 3 higher than those achieved in the previous hybrid LiNbO_3 ridge waveguides, either in chalcogenide glasses or with SiN ridges [RPC⁺15], and close to being comparable to the ~ 1 million Q s achieved in suspended LiNbO_3 microdisks [WBW⁺15]. The samples were not annealed above 400° C here, and it is likely that higher temperature annealing would reduce both any remaining loss in the LiNbO_3 , as well in the PECVD-deposited SiN at 1550 nm.

The hybrid waveguide approach has some attractive features – since the bulk of the mode is guided in an unpatterned film, the waveguides should be capable of reaching very low loss, at least from sidewall roughness. This will be particularly valuable for short wavelengths, where direct Ar^+ ion-milled LiNbO_3 SM waveguides still exhibit ~ 3 dB/cm even at 1550 nm owing to sidewall roughness [WXA⁺16]. Uniformity in loss will also be a major factor in the achievable extinction in MZ modulators; lower insertion losses, as well as the higher extinctions required for trapped-ion systems, will be possible with optimized directional couplers, and perhaps with alternate modulator designs altogether relying e.g. on distributed Bragg reflectors or interferometric techniques [MOR13].

Photorefractive damage in LiNbO_3 will pose a difficulty for short wavelength operation, though it can be mitigated by doping with MgO [BGT84] Prospects for applying these ideas at shorter wavelengths will be discussed in the next, final chapter. Broadly speaking though, the work presented here indicates that hybrid waveguides appear promising for low-loss electro-optically active waveguides in the visible.

Chapter 7

Conclusion, future work, and outlook

Ideas and techniques in the fields of integrated optics and atomic physics have developed largely independently over the last few decades. In recent years it seems that there may be significant opportunity for them to cooperate, e.g. [TTdL⁺14, GHY⁺14]. The work presented here has tried to bring aspects of these fields together to address particular needs of trapped ion systems.

This work demonstrated waveguiding optics integrated within planar ion traps for the first time, and shows how waveguide-based grating devices can generate free-space beams with precisely tailored beam profiles. In addition, CMOS processes were used to create planar ion traps with performance comparable to the single metal layer traps more commonly used, and the possibility of using integrated Si APDs for fluorescence detection readout in the blue was demonstrated. Finally, initial results on a potential route to visible light modulation were presented.

We have argued that the combination offers significant potential for practically achieving interesting operations with many ions. Furthermore the needs of the ion system raise new questions and avenues for work in integrated optics, and indeed many of the components considered here may find application also in different areas. In the remainder of this chapter we describe avenues for future work on the kinds of optical devices studied here, still focusing on the needs of the trapped-ion system, and then conclude with some discussion of the impact of such devices on trapped-ion systems in terms of scale and errors.

7.1 WG devices for blue and UV wavelengths

The waveguide devices studied here at $\lambda = 674$ nm could of course be implemented at other wavelengths provided low-loss waveguide materials. Silicon nitride itself should be applicable at shorter wavelengths, but it is unclear from the literature just how low in wavelength it supports low-loss propagation. Though as mentioned in Ch. 2 some modeling work predicts an absorption edge as low as 290 nm [Pal98], measurements on LPCVD Si₃N₄ waveguides fabricated and tested at Lincoln Labs have generally




Figure 7-1: Losses in SM Si_3N_4 waveguides fabricated and measured at Lincoln Laboratories, numbers provided by John Chiaverini and Jeremy Sage.

shown a strong increase in loss around 410 nm (Fig. 7-1). This rapid increase in loss with decreasing wavelength has a dependence on wavelength much stronger than λ^{-4} and is most likely due to material absorption rather than any kind of scattering. While experiments examining the effect of Si:N ratio and deposition conditions are at a preliminary stage as of this writing, prospects for SiN waveguides operating at 405 nm for photoionization of neutral strontium, or the shorter wavelengths required for the $S \rightarrow P$ transitions in other ions down to 313 nm for ${}^9\text{Be}^+$ for example, are uncertain.

7.1.1 Alumina as a core material for blue/UV wavelengths

Amorphous alumina, as can be obtained by atomic layer deposition (ALD) is an alternative, which one group has observed to allow guiding with no sharp absorption edge until approximately 250 nm [AWB⁺10]. In that work, losses below 4 dB/cm were demonstrated in thin films for wavelengths above this value. We are pursuing this material as a guiding core for shorter wavelengths; films recently deposited at MIT via thermal ALD by Gavin West showed losses below 0.5 dB/cm at 405 nm (and at 632 nm). The index of such films is lower than SiN; for example, crystalline sapphire at 632 nm has an index of 1.76, and amorphous ALD-deposited films can have lower indices near 1.67 [AWB⁺10], though this can be increased via annealing or growing at higher temperatures [ZJL⁺07, JSH⁺03, GFEG04]. Nevertheless this index is sufficient to create gratings still with oxide cladding, particularly at shorter wavelengths, though the anneal can help significantly due to the sensitivity of grating strength on index contrast. Furthermore, RIE of amorphous Al_2O_3 has been reported, using a variety of gas including the CF_4/O_2 mixture also often used for SiN, to allow definition of SM ridge waveguides at 1550 nm with losses on the order of 0.1 dB/cm [BAWP07]. These observations together with the previous work indicate that low-loss waveguide structures at short wavelengths should be achievable in such films, and it will be very

interesting to explore the possibility for this material for UV integrated optics for the ion application and otherwise. In the UV, both the core and the cladding may experience optical damage and resulting loss, though similar passivation techniques as have been explored for optical fibers may be applied to minimize the impact of this [CSW⁺14].

7.1.2 Materials for electro-optic modulators

LiNbO₃ has an absorption edge around 400 nm, limiting the use of the devices discussed in the previous chapter. Photorefractive damage may impede use even at slightly higher wavelengths, depending on power and temperature. Realizing such devices in the blue and UV may require consideration of significantly different materials, but one candidate apparent right away is Barium Borate (BBO), also a commonly used nonlinear crystal, which is optically transmissive down to $\lambda = 210$ nm, and is also known to exhibit a high damage threshold for UV radiation. Its relatively low indices ($n_e \approx 1.55$, $n_o \approx 1.68$) mean hybrid waveguides may not need thin-film BBO; a higher-index core material (alumina would be well matched) on a bulk substrate would support a guided mode. However, the mode overlap, together with the lower electro-optic coefficient of BBO ($\sim 10\times$ lower than LiNbO₃) will present challenges for efficient modulation in compact structures. Nevertheless these materials present definite possibilities, and approachable challenges for good device design.

It is possible that other materials, or techniques to induce electro-optic effects (strain engineering, or poling of glasses) can be brought to bear on the problem as well.

7.1.3 Characterization of waveguide nonlinearities and damage thresholds in the visible

As discussed in Ch. 1, certain nonlinearities, like self-phase modulation, may have an impact on quantum operations using guided light. The relative lack of data on nonlinear coefficients in the visible for potential core materials makes it challenging at present to quantify these effects, likely to play some role for high guided powers. Characterization of these nonlinear coefficients is likely to be of interest for applications beyond the present one; an interesting possibility for assessing self-phase modulation, with the use of just a CW, high power laser at the wavelength of interest, is to monitor the splitting ratio of an unbalanced MZI as a function of guided power; self-phase modulation here would induce a power-dependent relative phase between the two arms (carrying unequal powers owing to the imbalance). Damage thresholds, in samples with good heat conduction away from the waveguides, would be likewise an important parameter for assessing the extent to which a single bus could carry the light for operations on multiple ions.




Figure 7-2: Schematic of multi-layer trap device (only ground plane shown for simplicity), together with waveguides on a silicon substrate.

7.2 Direct fiber attaching, and silicon substrates

A few technical details related to the input coupling and the substrate may require investigation. As discussed at the end of Ch. 4, direct fiber coupling to the trap chip would likely be required for stable coupling of multiple wavelengths in multiple input waveguides, and to take advantage fully of the stability offered by this approach to addressing. A number of established approaches for such couplings were referenced there, and given these known techniques this appears to be a quite practical problem, e.g. [BJPK⁺15]. A few companies in fact provide such fiber alignment and attach services, and the Lincoln group is pursuing a chip with multiple inputs attached as such.

Eventually, the kinds of devices presented here in Chs. 3–5 would be combined, using waveguide optics on silicon substrates together with multiple metal layers, as illustrated (with simply a single ground plane between the waveguide layer and trap electrodes for simplicity) in Fig. 7-2. A potential issue in doing this is that the downwards-directed light will excite electron-hole pairs in the Si, which will result in photoconductivity of the Si as was problematic generally in Si traps without ground planes. Though in this case the ground plane should still screen the RF potentials from these carriers, the necessary openings in the ground planes to allow light to pass through may yet allow some leakage near where these carriers are excited, potentially difficult for stability. Additionally, even if the RF is sufficiently well shielded, these e-h pairs may diffuse somewhat before recombining, and hence result in fluctuating stray fields that may couple to the ion. The extent to which either of these are indeed problems will benefit from experimental study. If prohibitive, a transparent conductor as the ground plane would allow optical access without compromising the shielding around the grating locations; alternatively placing the ground plane underneath the waveguide layer and using it as the reflecting layer would also avoid the problem. Both are solutions in principle but not trivial to implement, and how necessary either would be remains to be seen.

7.3 Short-term ion experiments

Aside from their utility in large-scale systems, and independently of future additional integration of more complex systems including electronics, detectors, and modulators, the waveguide and grating optics used for ion addressing here should be helpful in small-scale experiments in the near term. This section briefly suggests a few possibilities.

7.3.1 High-fidelity, fast two-qubit gates

Experiments demonstrating 10^{-3} level infidelities have relied on Raman transitions addressed with beams focused to relatively large spots, with waists around order of $25\ \mu\text{m}$ [BHL⁺16, GTL⁺16]. In both cases, optical power delivered to the ions (5 mW in each Raman beam in the Oxford experiments, 10s of mW in the NIST experiments) and hence detuning of the Raman beams was limited. In the NIST experiment, for example, the largest known contributor to error was spontaneous Raman scattering (4×10^{-4}), which could be reduced for a given gate time by increasing the detuning and power; as discussed in Ch. 1 the decrease in scattering error for large detuning goes proportionally to the available intensity, and hence with beams focused to μm -scale spots and generated on-chip this error should be possible to reduce to a negligible level, while at the same time not increasing fluctuations from beam-pointing instabilities which would with bulk optics accompany tighter focusing. Hence, if a chip with low-loss waveguides at the appropriate wavelengths can be fiber-coupled stably and light delivered with low (compared to the $\sim 20 - 30$ dB gain in intensity from focusing) total loss to the ions, it appears that such optics could play a role in improving fidelities and also lowering required optical powers and/or gate times for two-qubit gate implementations. This is discussed in somewhat more detail later on in this chapter.

For gates on optical qubits [BKR08b], in which operations should be fast compared to the spontaneous decay, the higher intensities using such optics may also help increase operation fidelities, and furthermore low-loss waveguides are more easily available at these longer wavelengths.

Additionally, proposals to bypass trap frequency limitations in geometric phase gates via continuous pulse shaping (e.g., [GRZC05, PMGL⁺16]) will rely on even higher intensities than used in the slower gates implemented so far. Such gates have yet to be implemented experimentally with high fidelity in any setting, but it seems that the need for high intensity here also would motivate consideration of integrated focusing optics.

7.3.2 Transport gates in a 2D geometry

Transport gates with ions are an interesting approach to encoding operations on qubits, relying on moving ions through a beam waist of constant intensity in time rather than modulating the beam in time [LKOW07, dLM⁺16]. Carrying out such gates in a 2D geometry with free-space optics presents major challenges, given that

the waist would diffract significantly over the chip area, and plus the geometry of addressed sites would be significantly constrained by the free-space beam propagation. Beams propagating out of the chip and focused to tight spots may be favorable, both by alleviating this constraint on geometry, and again due to the tight focusing (the extent of which becomes decoupled from the chip size unlike in the case of beams parallel to the surface propagating over a long length). Highly focused beams would additionally allow faster gates in such an approach by minimizing the distance necessary to traverse. It would be interesting to explore the possibilities for these kind of operations using controlled ion motion, which may benefit from both the electronics achievable in CMOS processes for shaping the pulses encoding the movement, and the optics discussed here.

7.3.3 Optical addressing of microwave operations

As discussed in Ch. 1, for single qubit operations, microwave radiation directly driving transitions between hyperfine states have been capable of faster and higher fidelity operations than laser-based gates, and with less technical complexity. For two-qubit gates, the need for strong field gradient to couple to motional modes has prevented fast gates and the multiple Watt-level MW powers required are daunting [HSA⁺16], but they are nevertheless capable of high fidelities. Hence there is a compelling case that microwave gates may generally be preferable in the long run for single-qubit gates, but the difficulty in addressing operations to individual ions is a challenge for such approaches.

Even using microwaves for some coherent operations it may be favorable to take advantage of the straightforward addressing afforded by focused laser beams with high-purity profiles allowing low crosstalk to address operations to individual ions via laser-induced AC stark shifts, related to the approach in [SNM⁺13]. Clock states, which exhibit small differential AC stark shifts in response to light detuned by more than the ground state splitting, may still be so addressed by a small detuning on a quadrupole transition [AKB⁺07], if such transitions are available (e.g. in $^{43}\text{Ca}^+$). Such an approach may simplify the practical implementation of microwave gates on multiple qubits.

7.4 APDs and devices for ion movement

Though the results so far point to the possibility for the already fabricated integrated SPADs to enable ion state discrimination around 90 K, and with a relatively primitive passive quenching circuit, this work leaves plenty of room for future development. A variety of other APD structures, whether in unmodified electronics or imaging processes, should allow lower dark count rates.

However at lower temperatures (4 K), even with ideal SPAD structures, the limits imposed by afterpulsing present a question that appears unanswered in the literature so far. Development of appropriate quench circuits, or extremely pure SPAD structures with few active deep traps, that may be capable of operating at such tem-

peratures, presents an interesting challenge and one that would likely have significant use beyond trapped-ion systems as well.

But as also alluded to earlier, integrated SPADs alone are unlikely to be sufficient for meaningful readout in large-scale ion systems, due to the difficulties of obtaining sufficient solid angle as well as spatial discrimination between ions. As a result, trap structures capable of ion movement, for example to quickly bring particular ions closer to the surface, before a readout step, may be essential, whether readout relies on silicon SPADs or superconducting nanowire photodetectors. Ion traps capable of such movement have been demonstrated, so far allowing displacement of a linear trap's RF null by a fraction of the ion-electrode separation [VRMK16, HDA⁺09], or lateral displacements in point Paul traps [KHC11]. Large vertical displacements on the order of the trap height do not appear to have been explored, and may rely on structures where the ion is translated laterally with DC voltages along a direction over which the trap electrodes are tapered to reduce the RF trap height. Such structures may pose interesting problems for trap design, applicable in the long run either to integrated APDs or superconducting nanowire single-photon detectors [SVL⁺16].

7.5 Errors and scale

It is likely that most of the physical qubits in a practical quantum information processor will be dedicated to allowing diagnosis of errors that will inevitably occur; and the total size of a system required to implement fault-tolerant computation will be a strong function of the error rates in each individual qubit. For example, for the surface code considered in [FWH12], the analysis there shows that to maintain a target logical error probability of 4×10^{-5} , a physical error probability per operation of 1×10^{-3} as opposed to 5×10^{-3} allows an encoding for each logical qubit $12 \times$ smaller in terms of number of the number of physical qubits. This impact on scale can be more dramatic as yet lower logical error rates are targeted.

Our work has focused on scalable technology, so a concluding statement on errors may seem out of place. Nevertheless, it seems to us that going forward in quantum information processing in whatever physical modality, implementing the apparatus around the qubits necessary to control and measure them in a scalable fashion, and yet without introducing new sources of noise and decoherence, will prove a crucial challenge.

7.5.1 Noise/infidelity sources and magnitudes in this approach

How does the approach considered here likely fare in this regard? We can anticipate certain sources of added noise here. As discussed in Ch. 2, self-phase modulation in waveguides may lead to appreciable phase-shifts at the ions as intensity in the waveguide is modulated – or, since a $1 \mu\text{A}$ current in a wire $50 \mu\text{m}$ away from a trap site would result in a magnetic field of order a few $10 \mu\text{G}$ at the ion location, integration of electronics, depending on the magnitudes of currents involved, may result in appreciable magnetic field noise. Or, charging of dielectrics that may occur from the

passage of UV wavelengths may result in DC field drifts and induce micromotion. In each case, though, depending on the magnitude of the problem in a given implementation, solutions to such problems can be envisioned. Phase can be actively stabilized on-chip, magnetic field noise may be mitigated by proper arrangement of structures carrying high current or through the use of some magnetically shielding layer within the chip, and charging mitigated through use of transparent conductive materials.

Potential advantages in spontaneous Raman scattering/beam-pointing stability

On the other hand, addressing optics as pursued here can be expected to alleviate some sources of error in quantum gates. Spontaneous Raman scattering occurring during two-qubit gates can be a significant limitation to gate fidelities. In fact, though many different factors contribute appreciable error, this scattering was one of the largest and made up roughly half the total infidelity in the highest-fidelity gates reported in [BHL⁺16] and [GTL⁺16]. As summarized in Chapter 2, higher intensities in the Raman beams (as achieved most easily by tight focusing) allow a gate to be enacted in a given time using a detuning that increases approximately as $\Delta^2 \propto I$, I being the intensity at the ion location. The corresponding decrease in error from inelastic scattering goes as $1/I$ (again for a constant gate time). This means that by focusing even to a $10\times$ lower spot size (without increasing beam-pointing fluctuations), the contributions from this error in both experiments (roughly 5×10^{-4}) should be reduced to the 5×10^{-5} level, lower than other present errors. In any given atom, as detunings are increased potential couplings to other excited states would have to be considered, but given that the detuning for a given gate time scales roughly with \sqrt{I} , for a $10\times$ increase in I the effect on detuning is modest.

Compared to present experiments, our approach, assuming waveguides operating with low loss at the relevant wavelengths, should allow straightforward focusing to $100 - 1000\times$ smaller spots. This further focusing could of course allow faster gate times, or more operations in parallel. In addition, errors from beam-pointing instabilities can certainly be brought down significantly; though these appear to be minor contributors ($\sim 10^{-4}$ or lower) to the high-fidelity gates as reported, they would likely be a concern if tighter focusing were employed without sufficiently more stable optics.

Errors from crosstalk

Another consideration, discussed in detail with regards to observations on the grating devices in Chapters 3 and 4, relates to crosstalk errors on ions nearby one to which an operation is targeted. For a $5 \mu\text{m}$ ion spacing (corresponding to ~ 1 MHz axial frequency in a linear trap), the gratings in Ch. 3 allow relative intensities of approximately 2×10^{-4} ; this is roughly an order of magnitude lower than the relative intensities at similar ranges for experiments with individual addressing [DLF⁺16, SNM⁺13]. It still corresponds to an appreciable crosstalk error as defined in Ch. 4 of $\sim 5 \times 10^{-4}$, for resonant single-frequency addressing. Though this can be reduced almost by an orders of magnitude with grating aperture size as suggested in Fig. 3-14, a more dra-

matic enhancement comes if we use two-beam Raman addressing with both beams addressed. As discussed in Ch. 4 would result in this relative intensity translating into a crosstalk error instead proportional to the relative intensity squared, translating for the 2×10^{-4} relative intensity to $\epsilon_x = 1 \times 10^{-5}$ (such crosstalk errors were not achieved in [DLF⁺16] despite the use of Raman gates since there only one beam was addressed).

The relative intensity here is for 674 nm light, and would be reduced using a shorter wavelength (as would be used for Raman gates). Additionally, Stark shift addressing, as may be employed as suggested above for microwave single-qubit operations, can enable similar favorable scaling with relative intensity, and hence it appears that such optics should allow addressing errors low enough to consider not adding additional motional steps to separate ions before individual operations.

Factors not directly affected – e.g. motional mode heating

As mentioned above however, we emphasize that many different error sources come into play in such operations and many, like fluctuations in trap frequencies, laser coherence and intensity noise, are unaffected by the kinds of devices here. Although the heating rates observed in the CMOS traps (as low as 20 quanta/s at 1 MHz) are comparable to those observed in single layer traps (~ 5 quanta/s [CS14]), further work is required to reduce these. In considering two-qubit gate-fidelity targets, for a phase gate coupling to the center-of-mass mode of two ions, the error corresponding to a 10 quanta/s heating rate for a 100 μ s gate would be of order $\sim 10^{-3}$ [Bal14]. Of course, this is subject to the caveat that stretch modes of two ions can be used for these gates as well, which by virtue of their insensitivity to E-field noise common to both ions, exhibit substantially lower heating than the COM mode. Nevertheless, we can note that to target a heating contribution to total error of 10^{-5} , for a 100 μ s gate we would require a heating rate on whatever mode is used of ~ 0.1 quanta/s. Given that stretch modes can exhibit heating rates over an order of magnitude below that of the COM mode [KWM⁺98], this may not entail a full two orders of magnitude reduction; though the lower heating rates on stretch modes have not been characterized in our CMOS traps so far, the single-ion heating rates are on a level that it would be reasonable to expect may contribute errors of order 10^{-4} on a stretch-mode.

Integrated fluorescence detection

With regards to our work on integrated photodetectors, the dark counts and afterpulsing, together with the small solid angle, indicate that we are some ways from demonstrating high-fidelity readout with integrated detectors. Nevertheless, for the sake of getting a rough feel for how such devices may perform in terms of detection errors, we can assume that afterpulsing can be minimized at relevant temperatures to a point that it becomes possible to think strictly in terms of count rates (the bunches are of short enough duration compared to the mean inter-arrival time). Then, supposing devices behave as some of the best demonstrated at room temperature [WGH12], with 50 Hz dark count rates and 30% QE at 420 nm, we can ask what the read-

out fidelity can be expected to be for a given readout time and detector geometry. Supposing $\tau = 100 \mu\text{s}$ spent per readout, we would have an $\langle N \rangle_{\text{dk}} = 0.005$ (using the terminology of Ch. 1); assuming Poissonian statistics, for a given $\langle N \rangle_b$ we can therefore use the $\langle N \rangle_{\text{dk}} = 0$ points in Fig. 1-4 as a guide to expected infidelity. With η_{QE} and η_{SA} the QE and solid angle collection efficiencies, the expected $\langle N \rangle_b = 63 \text{ MHz} \times \tau \eta_{QE} \eta_{SA} = 2100 \eta_{SA}$ for the assumed τ . For a $10 \mu\text{m}$ -diameter detector and a $50 \mu\text{m}$ trap height, $\eta_{SA} = 0.0025$ implying $\langle N \rangle_b \sim 5$, and a readout error of order 10^{-2} . As can be seen in Fig. 1-4, by doubling the radius and thus roughly quadrupling $\langle N \rangle_b$, we can reduce the error rates to the 10^{-9} level (where effects other than signal-to-noise ratio would doubtless limit the fidelity – in the case of optical qubits, this would be the spontaneous decay probability of the D -state, which for $100 \mu\text{s}$ in $^{88}\text{Sr}^+$ would be $\sim 5 \times 10^{-4}$).

This discussion assumes much, but is included to indicate rough numbers that can be expected; the effects of scattered photons (effectively increasing the dark rate) or the possibility of reducing readout time by increasing area or decreasing trap height can be estimated from these considerations as well. Such detectors should therefore allow readout fidelities competitive to those with bulk optics, but we reiterate that the motivations for integrated detectors relate also to issues of scaling and the possibility for fast feedback.

7.5.2 Concluding remark, integration with isolated qubits

But considerations as to particular error sources directly impacted by integrated devices as explored here may turn out to be of secondary importance. More significant in the long run may be the fact that ions' nature, as relatively isolated and pure quantum systems, should allow integration of such optical and electrical devices as discussed here, without significantly coupling to the qubit state. This capability integration into larger systems while maintaining relative isolation has been suggested as an advantage of ions since some of the early discussions [CZ00], but systems capable of truly benefiting from it have not so far been implemented; and while careful measurement of noise sources introduced will have to be carried out as larger systems are made, the heating rates observed in the CMOS trap and trap with APD operating discussed here present initial indication that this expectation may be borne out. That is to say, despite ions' relative isolation from solid-state devices, complex solid-state optoelectronic systems when properly designed may still interface with atomic ions in fashions as extensible as with solid-state qubits, with the ions' isolation ensuring relatively small or predictable and cancellable effects on coherence. We hope that this work lays a basis to develop systems for interesting quantum information processing that may finally take advantage of this basic property of ions.

Appendix A

Motivation for constraint on curvature radius in grating design

To justify the approach to radius of curvature chosen in the devices presented in Ch. 3, motivated by minimizing the intensity profile's sidelobes along x , we briefly describe here the observations that led us to this constraint.

In designs where the radius was not matched to the radius of the phase fronts expanding through the taper (yet still focusing to around $50\ \mu\text{m}$ along z), both in measurements and in simulations of the devices, in the waveguide plane distortions in the transverse field profile were observed, which generally developed a minimum in intensity at the center of the waveguide ($x = 0$). This happened to a comparable degree for taper lengths ranging from $30\ \mu\text{ms}$ to $150\ \mu\text{ms}$, so was not simply due to the adiabaticity of the feeding taper, but was instead due to the fact that the grating is again not a weak grating, and the approximation that it only affects the amplitude of the field propagating along the grating length breaks down.

These observations are summarized in Fig. A-1, which shows the simulated and measured emission intensity profiles in the waveguide plane for the device formed on Quartz and discussed in Chapter 4. The minimum developing over the course of the grating in x is clearly visible in both cases, and a diffraction calculation based on this mode profile, with a reasonable phase-curvature assumed (to produce focusing as observed) gave sidelobe profiles in reasonable agreement with those observed.

The $30\ \mu\text{m}$ -taper expanding to the $18\ \mu\text{m}$ -wide waveguide is highly non-adiabatic, and under the suspicion that higher order modes excited in the taper were somehow interfering in a way that would not result from a mode profile closer to that of the fundamental mode illuminating the waveguide, we tested similar devices made with much longer tapers. 2D simulations of $150\ \mu\text{m}$ -long tapers indicated that, in the absence of grating lines, the mode should distort minimally in the $18\ \mu\text{m}$ -long grating region, but nevertheless a similar distortion was observed in measurements of long taper devices (Fig. A-1, bottom panel). These devices had radii chosen still to focus near $50\ \mu\text{m}$ above the chip. These observations indicated strongly that the taper itself was not responsible for the distortion, and pointed to distortion by the curved lines themselves, motivating the choice of radii described in Chapter 3.

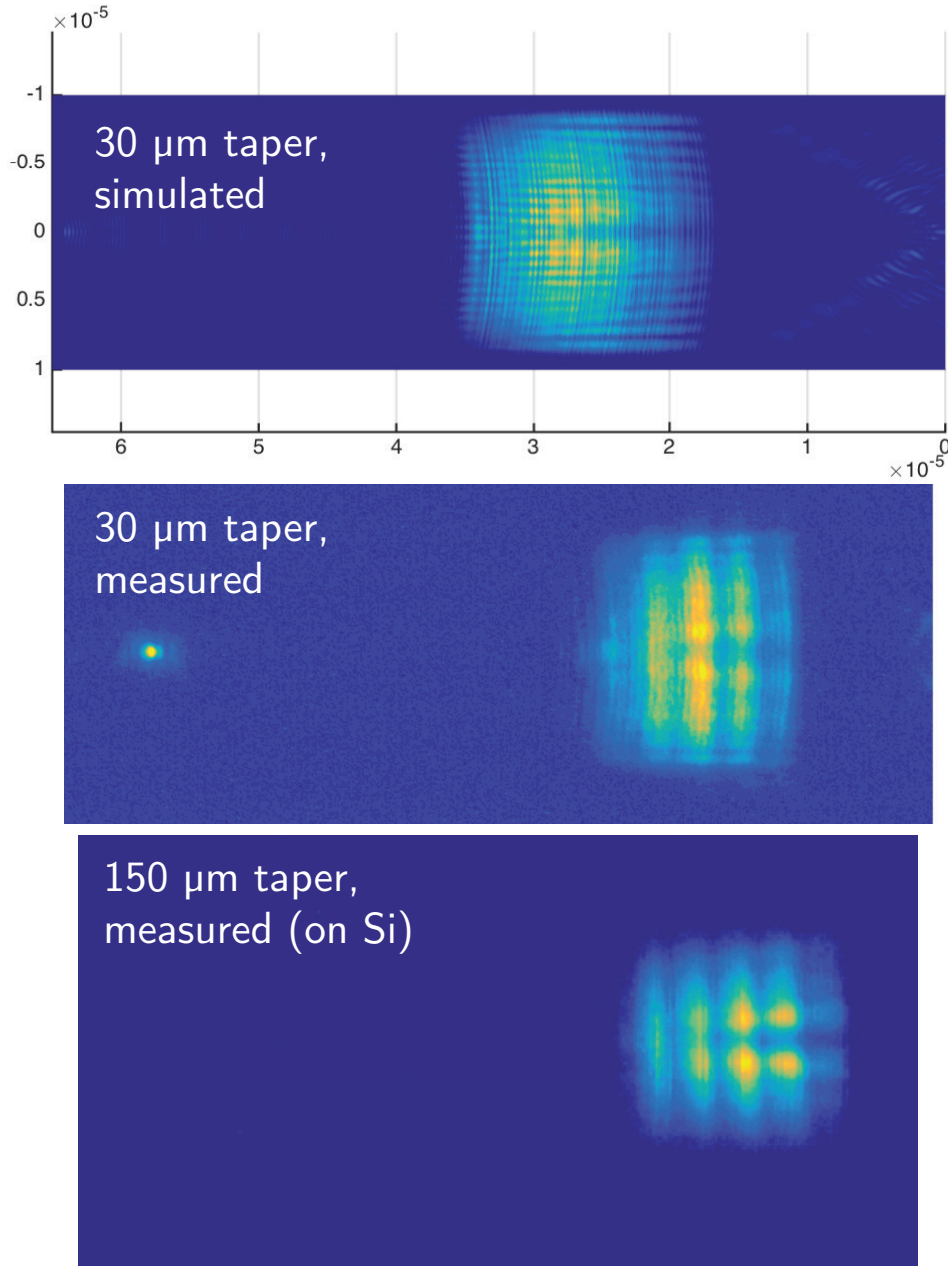


Figure A-1: Simulated and measured emission profiles in waveguide plane, showing mode distortion in grating region. The top image shows the simulated profile from the same device as discussed in Ch. 4 from a full 3D simulation in CST microwave studio, and the middle image is the measured emission profile. Using a much longer taper (150 μm), as shown in the bottom image, a similar distortion is observed.

Appendix B

E-beam write order and focusing grating sidelobe profiles

A perhaps helpful detail observed in fabricating the gratings presented in Ch. 3 concerned the order in which shapes constituting the grating design were written by the e-beam tool, which was found to play a significant role in determining the sidelobe strength and symmetry.

GDS files describing the pattern to be written were generated from within Cadence Virtuoso, consisting of a series of rectangles placed via scripts. These gds files were imported to software called LayoutBeamer, running on the e-beam facility computers, which processes the gds design file into a set of instructions to the e-beam tool, referred to in this software suite as a “.con” file. This file contains the order in which the beam should scan over the various shapes in the design. In general no “healing” was used for the designs presented in this thesis – “healing” would refer to removing overlaps between shapes in the same layer, or grouping together adjoining shapes into larger shapes with fewer vertices, but was not used as it, at least as involved in the current software version, led to less symmetric arrangements of individual shapes.

Within LayoutBeamer it is possible to view animations of any of the various write fields, which indicate the order in which the shapes involved would be written. Fig. B-1 shows two sequences of images representing the order in which the shapes involved in the grating pattern (for devices from Ch. 3) were written. Part a shows the sequence as these devices were originally implemented, and devices written with this order resulted in profiles like that plotted in the gray points of Fig. B-2. The asymmetry in the sidelobe profile was observed to be reproducible even in devices fabricated on a wholly different piece, indicating that the origin was systematic. In contrast, devices written with an order represented in Fig. B-1b resulted in the sidelobe profiles as presented in Ch. 3 and reproduced for comparison in the black trace of Fig. B-2. Clearly the more systematic write order results in more symmetric sidelobes, and at many distances also significantly lower intensities closer to the diffraction limit.

We can briefly note that in this case, orienting the devices vertically (as in Fig. B-1b) helped (and it was straightforward to simply expose pieces rotated accordingly), as did using a slightly older version of the output file format (.co7, which was followed by .co8) – the newer format resulted in splitting individual write fields (themselves




Figure B-1: Illustration of write orders employed for focusing grating devices of Ch. 3, with (a) showing six snapshots of the pattern during the course of writing using the order initially used and found to result in systematically asymmetric sidelobes; and (b) showing a sequence corresponding to a more symmetric write order as well as sidelobe profile.




Figure B-2: Comparison of sidelobe profiles obtained with different e-beam lithography write orders; the gray points represent the trace (as described in Ch. 3) reproducibly obtained with the write order in Fig. B-1a, with the black points the same as those presented in Fig. 3-12 and obtained in device written as in Fig. B-1b, replotted here for comparison.

500 μm -square areas) into smaller subfields. The older format considered each write field (600 μm -square) as a whole, and hence appeared to help in this. The devices in Ch. 4 were also written with the same orientation and format (and comparable write order), so the sidelobe intensities present there were not due to the kind of artifacts discussed here.

Certainly many other methods could be employed to tailor the write order, even in the newer file format. But, the LayoutBeamer software is often updated and doubtless the conventions for write order will change, so it would be of little value to dwell on the details that resulted in these write orders in this particular version – the main purpose of this section is simply to illustrate how the write order turned out to be significant in this case. Whatever software is used to define the e-beam write for future devices, when features at this level are of interest care should be taken to ensure the write order is logical. Of course for photolithographically defined devices these concerns would not arise.

Bibliography

- [AKB⁺07] L Aolita, K Kim, J Benhelm, CF Roos, and H Häffner. High-fidelity ion-trap quantum computing with hyperfine clock states. *Physical Review A*, 76(4):040303, 2007. [154](#)
- [AUW⁺10] Jason M Amini, H Uys, JH Wesenberg, S Seidelin, J Britton, JJ Bollinger, D Leibfried, C Ospelkaus, AP VanDevender, and DJ Wineland. Toward scalable ion traps for quantum information processing. *New journal of Physics*, 12(3):033031, 2010. [103](#)
- [AWB⁺10] Mustafa M Aslan, Nathan A Webster, Courtney L Byard, Marcelo B Pereira, Colin M Hayes, Rodrigo S Wiederkehr, and Sergio B Mendes. Low-loss optical waveguides for the near ultra-violet and visible spectral regions with al₂o₃ thin films from atomic layer deposition. *Thin Solid Films*, 518(17):4935–4940, 2010. [42](#), [49](#), [102](#), [150](#)
- [Bal14] CJ Ballance. *High-Fidelity Quantum Logic in Ca⁺*. PhD thesis, D. Phil. thesis, Univ. Oxford, 2014. [24](#), [34](#), [44](#), [157](#)
- [BAWP07] JDB Bradley, F Ay, K Wörhoff, and M Pollnau. Fabrication of low-loss channel waveguides in al₂o₃ and y₂o₃ layers by inductively coupled plasma reactive ion etching. *Applied Physics B*, 89(2-3):311–318, 2007. [49](#), [150](#)
- [BGL⁺12] R Bowler, J Gaebler, Y Lin, TR Tan, D Hanneke, JD Jost, JP Home, D Leibfried, and DJ Wineland. Coherent diabatic ion transport and separation in a multizone trap array. *Physical review letters*, 109(8):080502, 2012. [52](#)
- [BGT84] DA Bryan, Robert Gerson, and HE Tomaschke. Increased optical damage resistance in lithium niobate. *Applied Physics Letters*, 44(9):847–849, 1984. [148](#)
- [BHL⁺15] CJ Ballance, TP Harty, NM Linke, MA Sepiol, and DM Lucas. Laser-driven quantum logic gates with precision beyond the fault-tolerant threshold. *arXiv preprint arXiv:1512.04600*, 2015. [34](#)
- [BHL⁺16] CJ Ballance, TP Harty, NM Linke, MA Sepiol, and DM Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Physical Review Letters*, 117(6):060504, 2016. [15](#), [29](#), [153](#), [156](#)

- [BIS⁺16] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, John M Martinis, and Hartmut Neven. Characterizing quantum supremacy in near-term devices. *arXiv preprint arXiv:1608.00263*, 2016. [37](#)
- [BJO⁺08] Christopher Batten, Ajay Joshi, Jason Orcutt, Anatoly Khilo, Benjamin Moss, Charles Holzwarth, Milos Popovic, Hanqing Li, Henry I Smith, Judy Hoyt, et al. Building manycore processor-to-dram networks with monolithic silicon photonics. In *High Performance Interconnects, 2008. HOTI'08. 16th IEEE Symposium on*, pages 21–30. IEEE, 2008. [39](#)
- [BJPK⁺15] Tymon Barwicz, Alexander Janta-Polczynski, Marwan Khater, Yan Thibodeau, Robert Leidy, Jeffrey Maling, Stephan Martel, Sebastian Engelmann, Jason S Orcutt, Paul Fortier, et al. An o-band metamaterial converter interfacing standard optical fibers to silicon nanophotonic waveguides. In *Optical Fiber Communication Conference*, pages Th3F–3. Optical Society of America, 2015. [102](#), [152](#)
- [BKGN⁺16] Robin Blume-Kohout, John King Gamble, Erik Nielsen, Kenneth Rudinger, Jonathan Mizrahi, Kevin Fortier, and Peter Maunz. Certifying qubit operations below the fault tolerance threshold. *arXiv preprint arXiv:1605.07674*, 2016. [26](#)
- [BKM⁺13] R Barends, J Kelly, A Megrant, D Sank, E Jeffrey, Yu Chen, Y Yin, B Chiaro, J Mutus, C Neill, et al. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Physical review letters*, 111(8):080502, 2013. [33](#)
- [BKM⁺14] R Barends, J Kelly, A Megrant, A Veitia, D Sank, E Jeffrey, TC White, J Mutus, AG Fowler, B Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014. [33](#)
- [BKRB08a] J Benhelm, G Kirchmair, CF Roos, and R Blatt. Experimental quantum-information processing with c 43 a+ ions. *Physical Review A*, 77(6):062306, 2008. [25](#), [26](#)
- [BKRB08b] Jan Benhelm, Gerhard Kirchmair, Christian F Roos, and Rainer Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4(6):463–466, 2008. [24](#), [29](#), [34](#), [153](#)
- [BKRB14] M Brownnutt, M Kumph, P Rabl, and R Blatt. Ion-trap measurements of electric-field noise near surfaces. *arXiv preprint arXiv:1409.6572*, 2014. [30](#), [33](#)
- [BLWW15] S Burd, D Leibfried, AC Wilson, and DJ Wineland. Optically pumped semiconductor lasers for atomic and molecular physics. In *SPIE LASE*,

- pages 93490P–93490P. International Society for Optics and Photonics, 2015. [44](#)
- [BMR⁺16] JH Béjanin, TG McConkey, JR Rinehart, CT Earnest, CRH McRae, D Shiri, JD Bateman, Y Rohanizadegan, B Penava, P Breul, et al. Three-dimensional wiring for extensible quantum computing: The quantum socket. *Physical Review Applied*, 6(4):044010, 2016. [33](#)
- [BPW⁺15] T Brecht, W Pfaff, C Wang, Y Chu, L Frunzio, MH Devoret, and RJ Schoelkopf. Multilayer microwave integrated quantum circuits for scalable quantum computing. *arXiv preprint arXiv:1509.01127*, 2015. [33](#)
- [BWC⁺11] KR Brown, AC Wilson, Y Colombe, C Ospelkaus, AM Meier, E Knill, D Leibfried, and DJ Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Physical Review A*, 84(3):030303, 2011. [82](#)
- [CBB⁺05] J Chiaverini, RB Blakestad, Joe Britton, JD Jost, Chris Langer, Dietrich Leibfried, R Ozeri, and David J Wineland. Surface-electrode architecture for ion-trap quantum information processing. *arXiv preprint quant-ph/0501147*, 2005. [17](#), [29](#)
- [CGL⁺96] S Cova, M Ghioni, A Lacaita, C Samori, and F Zappa. Avalanche photodiodes and quenching circuits for single-photon detection. *Applied optics*, 35(12):1956–1976, 1996. [117](#), [125](#), [126](#)
- [CLJ08] J Chiaverini and WE Lybarger Jr. Laserless trapped-ion quantum simulations without spontaneous scattering using microtrap arrays. *Physical Review A*, 77(2):022324, 2008. [37](#), [81](#)
- [CMBK14] S Crain, E Mount, S Baek, and Jungsang Kim. Individual addressing of trapped 171yb+ ion qubits using a microelectromechanical systems-based beam steering system. *Applied Physics Letters*, 105(18):181115, 2014. [31](#), [41](#), [48](#), [81](#), [95](#)
- [CMQ⁺10] WC Campbell, J Mizrahi, Q Quraishi, C Senko, D Hayes, D Hucul, DN Matsukevich, P Maunz, and C Monroe. Ultrafast gates for single atomic qubits. *Physical review letters*, 105(9):090502, 2010. [26](#)
- [CPI⁺09] WA Challener, Chubing Peng, AV Itagi, D Karns, Wei Peng, Yingguo Peng, XiaoMin Yang, Xiaobin Zhu, NJ Gokemeijer, Y-T Hsia, et al. Heat-assisted magnetic recording by a near-field transducer with efficient optical energy transfer. *Nature Photonics*, 3(4):220–224, 2009. [58](#)
- [CS14] J Chiaverini and JM Sage. Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range. *Physical Review A*, 89(1):012318, 2014. [30](#), [33](#), [110](#), [157](#)

- [CSW⁺14] Yves Colombe, Daniel H Slichter, Andrew C Wilson, Dietrich Leibfried, and David J Wineland. Single-mode optical fiber for high-power, low-loss uv transmission. *Optics express*, 22(16):19783–19793, 2014. [102](#), [151](#)
- [CXWR14] Li Chen, Qiang Xu, Michael G Wood, and Ronald M Reano. Hybrid silicon and lithium niobate electro-optical ring modulator. *Optica*, 1(2):112–118, 2014. [128](#), [129](#)
- [CZ95] Juan I Cirac and Peter Zoller. Quantum computations with cold trapped ions. *Physical review letters*, 74(20):4091, 1995. [22](#)
- [CZ00] Juan Ignacio Cirac and Peter Zoller. A scalable quantum computer with ions in an array of microtraps. *Nature*, 404(6778):579–581, 2000. [158](#)
- [DBSR⁺11] Gian-Franco Dalla Betta, David Stoppa, Justin Richardson, Lucio Panzeri, and Robert Henderson. *Avalanche photodiodes in submicron CMOS technologies for high-sensitivity imaging*. Citeseer, 2011. [113](#)
- [DG98] Eric R Dufresne and David G Grier. Optical tweezer arrays and optical substrates created with diffractive optics. *Review of Scientific Instruments*, 69(5):1974–1977, 1998. [58](#)
- [DLF⁺16] S Debnath, NM Linke, C Figgatt, KA Landsman, K Wright, and C Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, 2016. [29](#), [35](#), [58](#), [74](#), [97](#), [156](#), [157](#)
- [dLM⁺16] L. E. de Clercq, H.-Y. Lo, M. Marinelli, D. Nadlinger, R. Oswald, V. Negnevitsky, D. Kienzler, B. Keitch, and J. P. Home. Parallel Transport Quantum Logic Gates with Trapped Ions. *Physical Review Letters*, 116(8):080502, February 2016. [101](#), [102](#), [127](#), [153](#)
- [DMR⁺04] Nicola Daldosso, Mirko Melchiorri, Francesco Riboli, Manuel Girardini, Georg Pucker, Michele Crivellari, Pierluigi Bellutti, Alberto Lui, and Lorenzo Pavesi. Comparison among various Si₃N₄ waveguide geometries grown within a CMOS fabrication pilot line. *Lightwave Technology, Journal of*, 22(7):1734–1740, 2004. [78](#), [88](#), [116](#)
- [Elt15] Amira Madeleine Eltony. *Scalable trap technology for quantum computing with ions*. PhD thesis, Massachusetts Institute of Technology, 2015. [111](#)
- [EWA⁺13] Amira M Eltony, Shannon X Wang, Gleb M Akselrod, Peter F Herskind, and Isaac L Chuang. Transparent ion trap with integrated photodetector. *Applied Physics Letters*, 102(5):054106, 2013. [30](#), [31](#), [101](#)

- [FHE06] Hod Finkelstein, Mark J Hsu, and Sadik C Esener. Sti-bounded single-photon avalanche diode in a deep-submicrometer cmos technology. *IEEE Electron Device Letters*, 27(11):887–889, 2006. [113](#), [123](#)
- [FLC⁺10] Ryan M Field, Jenifer Lary, John Cohn, Liam Paninski, and Kenneth L Shepard. A low-noise, single-photon avalanche diode in standard 0.13 μm complementary metal-oxide-semiconductor process. *Applied Physics Letters*, 97(21):211111, 2010. [52](#), [113](#), [115](#), [116](#), [117](#)
- [FMCM12] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012. [37](#), [51](#)
- [FSG09] Austin G Fowler, Ashley M Stephens, and Peter Groszkowski. High-threshold universal quantum computation on the surface code. *Physical Review A*, 80(5):052312, 2009. [34](#)
- [FWH12] Austin G Fowler, Adam C Whiteside, and Lloyd CL Hollenberg. Towards practical classical processing for the surface code. *Physical review letters*, 108(18):180501, 2012. [155](#)
- [GFEG04] MD Groner, FH Fabreguette, JW Elam, and SM George. Low-temperature al₂o₃ atomic layer deposition. *Chemistry of Materials*, 16(4):639–645, 2004. [150](#)
- [Gho95] Pradip K Ghosh. Ion traps. 1995. [16](#)
- [GHSN⁺13] Simon Gröblacher, Jeff T Hill, Amir H Safavi-Naeini, Jasper Chan, and Oskar Painter. Highly efficient coupling from an optical fiber to a nanoscale silicon optomechanical cavity. *Applied Physics Letters*, 103(18):181104, 2013. [99](#), [101](#), [102](#)
- [GHY⁺14] A Goban, C-L Hung, S-P Yu, JD Hood, JA Muniz, JH Lee, MJ Martin, AC McClung, KS Choi, DE Chang, et al. Atom–light interactions in photonic crystals. *Nature communications*, 5, 2014. [36](#), [149](#)
- [GJG⁺08] A Gorin, A Jaouad, E Grondin, V Aimez, and P Charette. Fabrication of silicon nitride waveguides for visible-light using pecvd: a study of the effect of plasma frequency on optical properties. *Optics Express*, 16(18):13509–13516, 2008. [49](#), [83](#), [98](#)
- [GPR⁺07] Andrea Guarino, Gorazd Poberaj, Daniele Rezzonico, Riccardo Degl’Innocenti, and Peter Günter. Electro–optically tunable microring resonators in lithium niobate. *Nature Photonics*, 1(7):407–410, 2007. [128](#)
- [GRSV07] William MJ Green, Michael J Rooks, Lidija Sekaric, and Yurii A Vlasov. Ultra-compact, low rf power, 10 gb/s silicon mach-zehnder modulator. *Optics express*, 15(25):17106–17113, 2007. [127](#)

- [GRZC05] Juan J García-Ripoll, Peter Zoller, and J Ignacio Cirac. Coherent control of trapped ions using off-resonant lasers. *Physical Review A*, 71(6):062309, 2005. [153](#)
- [GTL⁺16] JP Gaebler, TR Tan, Y Lin, Y Wan, R Bowler, AC Keith, S Glancy, K Coakley, E Knill, D Leibfried, et al. High-fidelity universal gate set for ${}^9\text{Be}^+$ ion qubits. *Physical Review Letters*, 117(6):060505, 2016. [15](#), [29](#), [34](#), [153](#), [156](#)
- [HAB⁺14] TP Harty, DTC Allcock, CJ Ballance, L Guidoni, HA Janacek, NM Linke, DN Stacey, and DM Lucas. High-fidelity preparation, gates, memory and readout of a trapped-ion quantum bit. *arXiv preprint arXiv:1403.1524*, 2014. [26](#), [27](#), [30](#), [34](#), [48](#)
- [Hau84] Hermann A Haus. *Waves and fields in optoelectronics*. Prentice-Hall,, 1984. [61](#)
- [HBHB10] Maximilian Harlander, Michael Brownnutt, Wolfgang Hänsel, and Rainer Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New Journal of Physics*, 12(9):093035, 2010. [101](#)
- [HBP⁺06] CW Holzwarth, T Barwicz, MA Popović, PT Rakich, EP Ippen, FX Kärtner, and Henry I Smith. Accurate resonant frequency spacing of microring filters without postfabrication trimming. *Journal of Vacuum Science & Technology B*, 24(6):3244–3247, 2006. [67](#)
- [HCW⁺12] DA Hite, Y Colombe, AC Wilson, KR Brown, U Warring, R Jördens, JD Jost, KS McKay, DP Pappas, D Leibfried, et al. 100-fold reduction of electric-field noise in an ion trap cleaned with in situ argon-ion-beam bombardment. *Physical review letters*, 109(10):103001, 2012. [33](#)
- [HDA⁺09] Peter Fønss Herskind, Aurélien Dantan, Magnus Albert, Joan P Marler, and Michael Drewsen. Positioning of the rf potential minimum line of a linear paul trap with micrometer precision. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 42(15):154008, 2009. [155](#)
- [Hem14] Cornelius Hempel. *Digital quantum simulation, Schrödinger cat state spectroscopy and setting up a linear ion trap*. PhD thesis, University of Innsbruck, 2014. [29](#)
- [HKL⁺87] Charles H Henry, RF Kazarinov, HJ Lee, KJ Orlowsky, and LE Katz. Low loss si 3 n 4-sio 2 optical waveguides on si. *Applied optics*, 26(13):2621–2624, 1987. [139](#)
- [HMM⁺10] D Hayes, DN Matsukevich, P Maunz, D Hucul, Q Quraishi, S Olmschenk, W Campbell, J Mizrahi, C Senko, and C Monroe. Entanglement of atomic qubits using an optical frequency comb. *Physical review letters*, 104(14):140501, 2010. [26](#)

- [HO81] Detlef Heitmann and Carmen Ortiz. Calculation and experimental verification of two-dimensional focusing grating couplers. *IEEE Journal of Quantum Electronics*, 17(7):1257–1263, 1981. 58
- [HOS⁺06] WK Hensinger, S Olmschenk, D Stick, D Hucul, M Yeo, M Acton, L Deslauriers, C Monroe, and J Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation. *Applied Physics Letters*, 88(3):034101, 2006. 52
- [HOT⁺10] Florian Haupt, Sumant SR Oemrawsingh, Susanna M Thon, Hyochul Kim, Dustin Kleckner, Dapeng Ding, Donald J Suntrup III, Pierre M Petroff, and Dirk Bouwmeester. Fiber-connectorized micropillar cavities. *Applied Physics Letters*, 97(13):131113, 2010. 102
- [HR06] Serge Haroche and Jean-Michel Raimond. *Exploring the quantum: atoms, cavities, and photons*. Oxford university press, 2006. 21, 91
- [HR13] Serge Haroche and Jean-Michel Raimond. *Exploring the quantum: atoms, cavities, and photons (oxford graduate texts)*. Oxford University Press, USA, 2013. 20
- [HRB08] Hartmut Häffner, Christian F Roos, and Rainer Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, 2008. 23, 25
- [HS04] JP Home and Andrew M Steane. Electrode configurations for fast separation of trapped ions. *arXiv preprint quant-ph/0411102*, 2004. 35
- [HSA⁺16] TP Harty, MA Sepiol, DTC Allcock, CJ Ballance, JE Tarlton, and DM Lucas. High-fidelity trapped-ion quantum logic using near-field microwaves. *Physical Review Letters*, 117(14):140501, 2016. 154
- [HSM⁺15] Nicholas C Harris, Gregory R Steinbrecher, Jacob Mower, Yoav Lahini, Mihika Prabhu, Tom Baehr-Jones, Michael Hochberg, Seth Lloyd, and Dirk Englund. Bosonic transport simulations in a large-scale programmable nanophotonic processor. *arXiv preprint arXiv:1507.03406*, 2015. 133
- [ISAF08] Kazuhiro Ikeda, Robert E Saperstein, Nikola Alic, and Yeshaiahu Fainman. Thermal and kerr nonlinear properties of plasma-deposited silicon nitride/silicon dioxide waveguides. *Optics express*, 16(17):12987–12994, 2008. 50
- [Jam98] Daniel FV James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B: Lasers and Optics*, 66(2):181–190, 1998. 98

- [JLH⁺14] Petar Jurcevic, Ben P Lanyon, Philipp Hauke, Cornelius Hempel, Peter Zoller, Rainer Blatt, and Christian F Roos. Quasiparticle engineering and entanglement propagation in a quantum many-body system. *Nature*, 511(7508):202–205, 2014. 47
- [JSH⁺03] Stefan Jakschik, Uwe Schroeder, Thomas Hecht, Martin Gutsche, Harald Seidl, and Johann W Bartha. Crystallization behavior of thin al-d-al 2 o 3 films. *Thin Solid Films*, 425(1):216–220, 2003. 150
- [KAS⁺15] Clemens J Krückel, Peter A Andrekson, Daryl T Spencer, Jared F Bauters, Martijn JR Heck, John E Bowers, et al. Continuous wave-pumped wavelength conversion in low-loss silicon nitride waveguides. *Optics letters*, 40(6):875–878, 2015. 49
- [KAZ⁺16] Mark Keil, Omer Amit, Shuyu Zhou, David Groswasser, Yonathan Japha, and Ron Folman. Fifteen years of cold matter on the atom chip: promise, realizations, and prospects. *Journal of Modern Optics*, pages 1–46, 2016. 58
- [KBB11] Muir Kumph, Michael Brownnutt, and Rainer Blatt. Two-dimensional arrays of radio-frequency ion traps with addressable interactions. *New Journal of Physics*, 13(7):073043, 2011. 35
- [KHC11] Tony Hyun Kim, Peter F Herskind, and Isaac L Chuang. Surface-electrode ion trap with integrated light source. *Applied Physics Letters*, 98(21):214103, 2011. 30, 31, 81, 155
- [KHL⁺16] Muir Kumph, Philip Holz, Kirsten Langer, Martin Meraner, Michael Niedermayr, Michael Brownnutt, and Rainer Blatt. Operation of a planar-electrode ion-trap array with adjustable rf electrodes. *New Journal of Physics*, 18(2):023047, 2016. 35
- [KK09] Jungsang Kim and Changsoon Kim. Integrated optical approach to trapped ion quantum computation. *Quantum Information & Computation*, 9(3):181–202, 2009. 30, 31, 40, 41, 81, 82
- [KKY⁺07] Hitoshi Kiuchi, Tetsuya Kawanishi, Masumi Yamada, Takahide Sakamoto, Masahiro Tsuchiya, Jun Amagai, and Masayuki Izutsu. High extinction ratio mach–zehnder modulator applied to a highly stable optical signal generator. *IEEE Transactions on Microwave Theory and Techniques*, 55(9):1964–1972, 2007. 145
- [KMW02] David Kielpinski, Chris Monroe, and David J Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, 2002. 15, 35, 37, 81

- [KNI⁺04] Kenji Kintaka, Junji Nishii, Yoshitaka Imaoka, Junpei Ohmori, Shogo Ura, Ryohei Satoh, and Hiroshi Nishihara. A guided-mode-selective focusing grating coupler. *IEEE Photonics Technology Letters*, 16(2):512–514, 2004. 58
- [KPM⁺05] J Kim, S Pau, Z Ma, HR McLellan, JV Gates, A Kornblit, Richard E Slusher, Robert M Jopson, I Kang, and M Dinu. System design for large-scale ion trap quantum information processor. *Quantum Information & Computation*, 5(7):515–537, 2005. 52
- [KWM⁺98] BE King, CS Wood, CJ Myatt, QA Turchette, D Leibfried, WM Itano, C Monroe, and DJ Wineland. Cooling the collective motion of trapped ions to initialize a quantum register. *Physical Review Letters*, 81(7):1525, 1998. 157
- [KZI⁺10] C Knoernschild, XL Zhang, L Isenhower, AT Gill, FP Lu, M Saffman, and J Kim. Independent individual addressing of multiple neutral atom qubits with a micromirror-based beam steering system. *Applied Physics Letters*, 97(13):134101, 2010. 95
- [lak] Lakeshore thermal conductivity values: http://www.lakeshore.com/Documents/LSTC_appendixI_1.pdf. 82
- [LBD⁺05] PJ Lee, KA Brickman, L Deslauriers, PC Haljan, LM Duan, and C Monroe. Phase control of trapped ion quantum gates. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S371, 2005. 24
- [LDM⁺03] Dietrich Leibfried, Brian DeMarco, Volker Meyer, David Lucas, Murray Barrett, Joe Britton, WM Itano, B Jelenković, Chris Langer, Till Rosenband, et al. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(6930):412–415, 2003. 24
- [LGA⁺08] Jaroslaw Labaziewicz, Yufei Ge, Paul Antohi, David Leibbrandt, Kenneth R Brown, and Isaac L Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. *Physical review letters*, 100(1):013001, 2008. 30, 33
- [LKOW07] D Leibfried, E Knill, C Ospelkaus, and DJ Wineland. Transport quantum logic gates for trapped ions. *Physical Review A*, 76(3):032324, 2007. 51, 101, 102, 153
- [LLC16] Kuan-Yu Lin, Guang Hao Low, and Issac L Chuang. Effects of electrode surface roughness on motional heating of trapped ions. *arXiv preprint arXiv:1605.02808*, 2016. 110
- [LOJ⁺05] C Langer, R Ozeri, John D Jost, J Chiaverini, B DeMarco, A Ben-Kish, RB Blakestad, J Britton, DB Hume, WM Itano, et al. Long-lived qubit

- memory using atomic ions. *Physical review letters*, 95(6):060502, 2005. [26](#)
- [LRC12] Myung-Jae Lee, Holger Rucker, and Woo-Young Choi. Effects of guard-ring structures on the performance of silicon avalanche photodetectors fabricated with standard cmos technology. *IEEE Electron Device Letters*, 33(1):80–82, 2012. [113](#)
- [Mau16] Peter Lukas Wilhelm Maunz. High optical access trap 2.0. Technical report, Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States), 2016. [41](#)
- [MBB⁺13] Emily Mount, So-Young Baek, Matthew Blain, Daniel Stick, Daniel Gaultney, Stephen Crain, Rachel Noek, Taehyun Kim, Peter Maunz, and Jungsang Kim. Single qubit manipulation in a microfabricated surface electrode ion trap. *New Journal of Physics*, 15(9):093018, 2013. [25](#), [48](#)
- [MBM⁺16] Karan K Mehta, Colin D Bruzewicz, Robert McConnell, Rajeev J Ram, Jeremy M Sage, and John Chiaverini. Integrated optical addressing of an ion qubit. *Nature Nanotechnology*, 11:1066–1070, 2016. [58](#), [64](#), [74](#), [80](#), [81](#)
- [MEB⁺14] Karan K Mehta, AM Eltony, CD Bruzewicz, IL Chuang, RJ Ram, JM Sage, and J Chiaverini. Ion traps fabricated in a CMOS foundry. *Applied Physics Letters*, 105(4):044103, 2014. [32](#), [78](#), [101](#), [103](#), [111](#)
- [Meh12] Karan Kartik Mehta. Resonant photonic crystal photodetectors for the infrared in silicon. Master’s thesis, Massachusetts Institute of Technology, 2012. [138](#)
- [MGM⁺11] Attila Mekis, Steffen Gloeckner, Gianlorenzo Masini, Adithyaram Narasimha, Thierry Pinguet, Subal Sahni, and Peter De Dobbelaere. A grating-coupler-enabled CMOS photonics platform. *IEEE Journal of Selected Topics in Quantum Electronics*, 17(3):597–608, 2011. [78](#), [80](#)
- [MIH06] Koji Matsumoto, Akihiro Inomata, and Shin-ya Hasegawa. Thermally assisted magnetic recording. *Fujitsu Sci. Tech. J.*, 42(1):158–167, 2006. [58](#)
- [MK13] Christopher Monroe and Jungsang Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013. [15](#), [30](#), [31](#), [33](#)
- [MMK⁺95] Chris Monroe, DM Meekhof, BE King, Wayne M Itano, and David J Wineland. Demonstration of a fundamental quantum logic gate. *Physical review letters*, 75(25):4714, 1995. [15](#)

- [MMO⁺07] DL Moehring, P Maunz, S Olmschenk, KC Younge, DN Matsukevich, L-M Duan, and C Monroe. Entanglement of single-atom quantum bits at a distance. *Nature*, 449(7158):68–71, 2007. 35
- [MMS⁺16] Esteban A. Martinez, Christine A. Muschik, Philipp Schindler, Daniel Nigg, Alexander Erhard, Markus Heyl, Philipp Hauke, Marcello Dalmonde, Thomas Monz, Peter Zoller, and Rainer Blatt. Real-time dynamics of lattice gauge theories with a few-qubit quantum computer. *Nature*, 534(7608):516–519, 06 2016. 29
- [MNM⁺15] Thomas Monz, Daniel Nigg, Esteban A Martinez, Matthias F Brandl, Philipp Schindler, Richard Rines, Shannon X Wang, Isaac L Chuang, and Rainer Blatt. Realization of a scalable shor algorithm. *arXiv preprint arXiv:1507.08852*, 2015. 29, 33
- [MOR13] Karan K Mehta, Jason S Orcutt, and Rajeev J Ram. Fano line shapes in transmission spectra of silicon photonic crystal resonators. *Applied Physics Letters*, 102(8):081109, 2013. 133, 148
- [MOTZ⁺14] Karan K Mehta, Jason S Orcutt, Ofer Tehar-Zahav, Zvi Sternberg, Reha Bafrali, Roy Meade, and Rajeev J Ram. High-Q CMOS-integrated photonic crystal microcavity devices. *Scientific reports*, 4, 2014. 52, 78
- [MPSO09] Jonathan CF Matthews, Alberto Politi, André Stefanov, and Jeremy L O’Brien. Manipulation of multiphoton entanglement in waveguide quantum circuits. *Nature Photonics*, 3(6):346–350, 2009. 47
- [MR16] Karan K Mehta and Rajeev J Ram. Precise and diffraction-limited waveguide-to-free-space focusing gratings. *arXiv preprint arXiv:1607.00107*, 2016. 57
- [MRR⁺14] C Monroe, R Raussendorf, A Ruthven, KR Brown, P Maunz, L-M Duan, and J Kim. Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, 89(2):022317, 2014. 35, 37
- [MSB⁺11] Thomas Monz, Philipp Schindler, Julio T Barreiro, Michael Chwalla, Daniel Nigg, William A Coish, Maximilian Harlander, Wolfgang Hänsel, Markus Hennrich, and Rainer Blatt. 14-qubit entanglement: Creation and coherence. *Physical Review Letters*, 106(13):130506, 2011. 29
- [MSW⁺08] AH Myerson, DJ Szwer, SC Webster, DTC Allcock, MJ Curtis, G Imreh, JA Sherman, DN Stacey, AM Steane, and DM Lucas. High-fidelity readout of trapped-ion qubits. *Physical review letters*, 100(20):200502, 2008. 27, 54
- [NC10] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010. 15, 20

- [NLK⁺14] Michael Niedermayr, Kirill Lakhmanskiy, Muir Kumph, Stefan Partel, Johannes Edlinger, Michael Brownutt, and Rainer Blatt. Cryogenic surface ion trap based on intrinsic silicon. *New Journal of Physics*, 16(11):113068, 2014. [70](#), [109](#)
- [NLR⁺99] H Ch NaÈgerl, D Leibfried, H Rohde, G Thalhammer, J Eschner, F Schmidt-Kaler, and R Blatt. Laser addressing of individual ions in a linear ion trap. *Physical Review A*, 60(1):145, 1999. [47](#), [95](#)
- [NPW⁺16] Jelena Notaros, Fabio Pavanello, Mark T Wade, Cale Gentry, Amir Atabaki, Luca Alloatti, Rajeev J Ram, and Milos Popovic. Ultra-efficient CMOS fiber-to-chip grating couplers. In *Optical Fiber Communication Conference*, pages M2I–5. Optical Society of America, 2016. [64](#)
- [OC12] Antony Orth and Kenneth Crozier. Microscopy with microlens arrays: high throughput, high resolution and light-field imaging. *Optics Express*, 20(12):13522–13531, 2012. [58](#)
- [OIB⁺07] R Ozeri, WM Itano, RB Blakestad, J Britton, J Chiaverini, JD Jost, C Langer, D Leibfried, R Reichle, S Seidelin, et al. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A*, 75(4):042329, 2007. [26](#), [44](#)
- [OLA⁺08] Christian Ospelkaus, Christopher E Langer, Jason M Amini, Kenton R Brown, Dietrich Leibfried, and David J Wineland. Trapped-ion quantum logic gates based on oscillating magnetic fields. *Physical review letters*, 101(9):090502, 2008. [26](#)
- [OMS⁺12] Jason S Orcutt, Benjamin Moss, Chen Sun, Jonathan Leu, Michael Georgas, Jeffrey Shainline, Eugen Zraggen, Hanqing Li, Jie Sun, Matthew Weaver, et al. Open foundry platform for high-performance electronic-photonic integration. *Optics express*, 20(11):12222–12232, 2012. [52](#), [78](#), [101](#), [111](#)
- [OR10] Jason S Orcutt and Rajeev J Ram. Photonic device layout within the foundry cmos design environment. Institute of Electrical and Electronics Engineers, 2010. [67](#)
- [Orc12] Jason Scott Orcutt. *Monolithic electronic-photonic integration in state-of-the-art CMOS processes*. PhD thesis, Massachusetts Institute of Technology, 2012. [107](#)
- [ORI⁺10] Ardavan F Oskooi, David Roundy, Mihai Ibanescu, Peter Bermel, John D Joannopoulos, and Steven G Johnson. Meep: A flexible free-software package for electromagnetic simulations by the fdtd method. *Computer Physics Communications*, 181(3):687–702, 2010. [63](#)

- [OWC⁺11] C Ospelkaus, U Warring, Y Colombe, KR Brown, JM Amini, D Leibfried, and DJ Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476(7359):181–184, 2011. 26
- [OYM⁺07] S Olmschenk, KC Younge, DL Moehring, DN Matsukevich, P Maunz, and C Monroe. Manipulation and detection of a trapped yb⁺ hyperfine qubit. *Physical Review A*, 76(5):052314, 2007. 26
- [Pal98] Edward D Palik. *Handbook of optical constants of solids*, volume 3. Academic press, 1998. 42, 149
- [PCR⁺08] Alberto Politi, Martin J Cryan, John G Rarity, Siyuan Yu, and Jeremy L O’Brien. Silica-on-silicon waveguide quantum circuits. *Science*, 320(5876):646–649, 2008. 47, 82
- [PG61] M Pollak and TH Geballe. Low-frequency conductivity due to hopping processes in silicon. *Physical Review*, 122(6):1742, 1961. 109
- [PHSG12] Gorazd Poberaj, Hui Hu, Wolfgang Sohler, and Peter Guenter. Lithium niobate on insulator (lnoi) for micro-photonics devices. *Laser & Photonics Reviews*, 6(4):488–503, 2012. 51, 128
- [PLM⁺09] Nishant Patil, Albert Lin, Edward R Myers, Kounghmin Ryu, Alexander Badmaev, Chongwu Zhou, H-S Philip Wong, and Subhasish Mitra. Wafer-scale growth and transfer of aligned single-walled carbon nanotubes. *IEEE Transactions on Nanotechnology*, 8(4):498–504, 2009. 83
- [PMGL⁺16] M Palmero, S Martínez-Garaot, D Leibfried, DJ Wineland, and JG Muga. Fast phase gates with trapped ions. *arXiv preprint arXiv:1609.01892*, 2016. 35, 153
- [PRH13] Adam M Packer, Botond Roska, and Michael Häusser. Targeting neurons and photons for optogenetics. *Nature Neuroscience*, 16(7):805–815, 2013. 58
- [RGF⁺03] A Rochas, M Gani, B Furrer, PA Besse, RS Popovic, G Ribordy, and N Gisin. Single photon detector fabricated in a complementary metal–oxide–semiconductor high-voltage technology. *Review of Scientific Instruments*, 74(7):3263–3270, 2003. 113
- [RGMZ⁺13] Sebastian Romero-García, Florian Merget, Frank Zhong, Hod Finkelshtein, and Jeremy Witzens. Visible wavelength silicon nitride focusing grating coupler with AlCu/TiN reflector. *Optics Letters*, 38(14):2521–2523, 2013. 99
- [RLA⁺07] Ivan Rech, Ivan Labanca, Giacomo Armellini, Angelo Gulinatti, Massimo Ghioni, and Sergio Cova. Operation of silicon single photon avalanche diodes at cryogenic temperature. *The Review of scientific instruments*, 78(6):063105–063105, 2007. 55, 123

- [RMK⁺13] Payam Rabiei, Jichi Ma, Saeed Khan, Jeff Chiles, and Sasan Fathpour. Heterogeneous lithium niobate photonics on silicon substrates. *Optics express*, 21(21):25573–25581, 2013. 128, 129
- [RPC⁺15] Ashutosh Rao, Aniket Patil, Jeff Chiles, Marcin Malinowski, Spencer Novak, Kathleen Richardson, Payam Rabiei, and Sasan Fathpour. Heterogeneous microring and mach-zehnder modulators based on lithium niobate and chalcogenide glasses on silicon. *Optics express*, 23(17):22746–22752, 2015. 128, 148
- [SB87] Richard A Soref and Brian R Bennett. Electrooptical effects in silicon. *IEEE Journal of Quantum Electronics*, 23:123–129, 1987. 127
- [SBO⁺14] Joshua W Silverstone, Damien Bonneau, Kazuya Ohira, Nob Suzuki, Haruhiko Yoshida, Norio Iizuka, Mizunori Ezaki, Chandra M Nataraajan, Michael G Tanner, Robert H Hadfield, et al. On-chip quantum interference between silicon photon-pair sources. *Nature Photonics*, 8(2):104–108, 2014. 47
- [SEHS10] Ch Schneider, Martin Enderlein, Thomas Huber, and Tobias Schätz. Optical trapping of an ion. *Nature Photonics*, 4(11):772–775, 2010. 16
- [SGG99] Stephen M Schultz, Elias N Glytsis, and Thomas K Gaylord. Volume grating preferential-order focusing waveguide coupler. *Optics Letters*, 24(23):1708–1710, 1999. 58
- [SGO⁺15] Chen Sun, Michael Georgas, Jason Orcutt, Benjamin Moss, Yu-Hsin Chen, Jeffrey Shainline, Mark Wade, Karan Mehta, Kareem Nammari, Erman Timurdogan, et al. A monolithically-integrated chip-to-chip optical link in bulk CMOS. *Solid-State Circuits, IEEE Journal of*, 50(4):828–844, 2015. 32, 52
- [Sho97] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM journal on computing*, 26(5):1484–1509, 1997. 15
- [SKC12] Jeremy M Sage, Andrew J Kerman, and John Chiaverini. Loading of a surface-electrode ion trap from a remote, precooled source. *Physical Review A*, 86(1):013417, 2012. 54, 86, 90
- [SKHR⁺03] Ferdinand Schmidt-Kaler, Hartmut Häffner, Mark Riebe, Stephan Gulde, Gavin PT Lancaster, Thomas Deuschle, Christoph Becher, Christian F Roos, Jürgen Eschner, and Rainer Blatt. Realization of the cirac-zoller controlled-not quantum gate. *Nature*, 422(6930):408–411, 2003. 23
- [SKK⁺00] CA Sackett, David Kielpinski, BE King, C Langer, V Meyer, CJ Myatt, M Rowe, QA Turchette, WM Itano, DJ Wineland, et al. Experimental entanglement of four particles. *Nature*, 404(6775):256–259, 2000. 24

- [SLY⁺97] Stephen Sheard, Tingdi Liao, Guoguang Yang, Philip Prewett, and Jianguo Zhu. Focusing waveguide grating coupler using a diffractive doublet. *Applied Optics*, 36(19):4349–4353, 1997. 58
- [SM99] Anders Sørensen and Klaus Mølmer. Quantum computation with ions in thermal motion. *Physical review letters*, 82(9):1971, 1999. 23
- [SM00] Anders Sørensen and Klaus Mølmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62(2):022311, 2000. 23, 24, 47
- [SN06] Simon M Sze and Kwok K Ng. *Physics of semiconductor devices*. John wiley & sons, 2006. 113
- [SNBT86] Th Sauter, W Neuhauser, R Blatt, and PE Toschek. Observation of quantum jumps. *Physical Review Letters*, 57(14):1696, 1986. 95
- [SNJ⁺11] Erik W Streed, Benjamin G Norton, Andreas Jechow, Till J Weinhold, and David Kielpinski. Imaging of trapped ions with a microfabricated optic for quantum information processing. *Physical review letters*, 106(1):010502, 2011. 30, 31, 81
- [SNM⁺13] Philipp Schindler, Daniel Nigg, Thomas Monz, Julio T Barreiro, Esteban Martinez, Shannon X Wang, Stephan Quint, Matthias F Brandl, Volckmar Nebendahl, Christian F Roos, et al. A quantum information processor with trapped ions. *New Journal of Physics*, 15(12):123012, 2013. 29, 47, 58, 74, 82, 95, 97, 154, 156
- [SOW⁺13] Jeffrey M Shainline, Jason S Orcutt, Mark T Wade, Kareem Namari, Benjamin Moss, Michael Georgas, Chen Sun, Rajeev J Ram, Vladimir Stojanović, and Miloš A Popović. Depletion-mode carrier-plasma optical modulator in zero-change advanced cmos. *Optics letters*, 38(15):2657–2659, 2013. 127
- [Ste04] Andrew M Steane. How to build a 300 bit, 1 giga-operation quantum computer. *arXiv preprint quant-ph/0412165*, 2004. 52
- [STY⁺13] Jie Sun, Erman Timurdogan, Ami Yaacobi, Ehsan Shah Hosseini, and Michael R Watts. Large-scale nanophotonic phased array. *Nature*, 493(7431):195–199, 2013. 39, 47, 81
- [SVL⁺16] DH Slichter, VB Verma, D Leibfried, RP Mirin, SW Nam, and DJ Wineland. Uv-sensitive superconducting nanowire single photon detectors for integration in an ion trap. *arXiv preprint arXiv:1611.09949*, 2016. 126, 155
- [TBB⁺02] Dirk Taillaert, Wim Bogaerts, Peter Bienstman, Thomas F Krauss, Peter Van Daele, Ingrid Moerman, Steven Verstuyft, Kurt De Mesel, and

- Roel Baets. An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers. *Quantum Electronics, IEEE Journal of*, 38(7):949–955, 2002. 99, 101
- [TBB04] Dirk Taillaert, Peter Bienstman, and Roel Baets. Compact efficient broadband grating coupler for silicon-on-insulator waveguides. *Optics letters*, 29(23):2749–2751, 2004. 64
- [TISF10] DTH Tan, K Ikeda, PC Sun, and Y Fainman. Group velocity dispersion and self phase modulation in silicon nitride waveguides. *Applied Physics Letters*, 96(6):061101, 2010. 50
- [TNT+15] TG Tiecke, KP Nayak, JD Thompson, T Peyronel, NP De Leon, V Vuletić, and MD Lukin. Efficient fiber-optical interface for nanophotonic devices. *Optica*, 2(2):70–75, 2015. 102
- [TSAS+14] Erman Timurdogan, Cheryl M Sorace-Agaskar, Jie Sun, Ehsan Shah Hosseini, Aleksandr Biberman, and Michael R Watts. An ultralow power athermal silicon modulator. *Nature communications*, 5, 2014. 127
- [TTdL+14] TG Tiecke, Jeffrey Douglas Thompson, Nathalie Pulmones de Leon, LR Liu, V Vuletić, and Mikhail D Lukin. Nanophotonic quantum phase switch with a single atom. *Nature*, 508(7495):241–244, 2014. 36, 149
- [TTZ+13] Jeffrey Douglas Thompson, TG Tiecke, Alexander S Zibrov, V Vuletić, and Mikhail D Lukin. Coherence and Raman sideband cooling of a single atom in an optical tweezer. *Physical Review Letters*, 110(13):133001, 2013. 76
- [Tur66] EH Turner. High-frequency electro-optic coefficients of lithium niobate. *Applied Physics Letters*, 8:303–304, 1966. 132, 145, 146
- [USNK86] Shogo Ura, Toshiaki Suhara, Hiroshi Nishihara, and Jiro Koyama. An integrated-optic disk pickup device. *IEEE Journal of Lightwave Technology*, 4(7):913–918, 1986. 58
- [VCA+10] AP VanDevender, Y Colombe, J Amini, D Leibfried, and DJ Wineland. Efficient fiber optic detection of trapped ion fluorescence. *Physical review letters*, 105(2):023001, 2010. 30, 31, 81
- [Ver89] Joseph Thomas Verdeyen. Laser electronics. 1989. 61
- [VRMK16] Andre Van Rynbach, Peter Maunz, and Jungsang Kim. An integrated mirror and surface ion trap with a tunable trap location. *Applied Physics Letters*, 109(22):221108, 2016. 155

- [Wan12] Shannon Xuanyue Wang. *Quantum gates, sensors, and systems with trapped ions*. PhD thesis, Massachusetts Institute of Technology, 2012. [25](#), [43](#)
- [WBB⁺03] David J Wineland, M Barrett, J Britton, J Chiaverini, B DeMarco, WM Itano, B Jelenković, C Langer, D Leibfried, V Meyer, et al. Quantum information processing with trapped ions. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 361(1808):1349–1361, 2003. [43](#)
- [WBL⁺14] Cheng Wang, Michael J Burek, Zin Lin, Haig A Atikian, Vivek Venkataraman, I-Chun Huang, Peter Stark, and Marko Lončar. Integrated high quality factor lithium niobate microdisk resonators. *Optics express*, 22(25):30924–30933, 2014. [128](#)
- [WBW⁺15] Jie Wang, Fang Bo, Shuai Wan, Wuxia Li, Feng Gao, Junjie Li, Guoquan Zhang, and Jingjun Xu. High-q lithium niobate microdisk resonators on a chip for efficient electro-optic modulation. *Optics express*, 23(18):23072–23078, 2015. [128](#), [148](#)
- [WCZ⁺10] Jigang Wu, Xiquan Cui, Guoan Zheng, Ying Min Wang, Lap Man Lee, and Changhuei Yang. Wide field-of-view microscope based on holographic focus grid illumination. *Optics Letters*, 35(13):2188–2190, 2010. [58](#)
- [WGH12] Eric AG Webster, Lindsay A Grant, and Robert K Henderson. A high-performance single-photon avalanche diode in 130-nm cmos imaging technology. *IEEE Electron Device Letters*, 33(11):1589–1591, 2012. [52](#), [113](#), [157](#)
- [WKYY⁺00] Ed L Wooten, Karl M Kissa, Alfredo Yi-Yan, Edmond J Murphy, Donald A Lafaw, Peter F Hallemeier, David Maack, Daniel V Attanasio, Daniel J Fritz, Gregory J McBrien, et al. A review of lithium niobate modulators for fiber-optic communications systems. *IEEE Journal of selected topics in Quantum Electronics*, 6(1):69–82, 2000. [127](#)
- [WLG⁺09] Shannon X Wang, Jaroslaw Labaziewicz, Yufei Ge, Ruth Shewmon, and Isaac L Chuang. Individual addressing of ions using magnetic field gradients in a surface-electrode ion trap. *Applied Physics Letters*, 94(9):094103, 2009. [48](#)
- [WMI⁺97] David J Wineland, C Monroe, WM Itano, D Leibfried, BE King, and DM Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *arXiv preprint quant-ph/9710025*, 1997. [18](#), [26](#), [27](#), [28](#), [47](#)

- [WOC⁺13] U Warring, C Ospelkaus, Y Colombe, R Jördens, D Leibfried, and DJ Wineland. Individual-ion addressing with microwave field gradients. *Physical review letters*, 110(17):173002, 2013. 26, 48, 95, 97
- [Won02] Ka-Kha Wong. *Properties of lithium niobate*. Number 28. IET, 2002. 42
- [WPK⁺15] Mark T Wade, Fabio Pavanello, Rajesh Kumar, Cale M Gentry, Amir Atabaki, Rajeev Ram, Vladimir Stojanović, and Miloš A Popović. 75% efficient wide bandwidth grating couplers in a 45 nm microelectronics cmos process. In *2015 IEEE Optical Interconnects Conference (OI)*, pages 46–47. IEEE, 2015. 99
- [WRGH11] Eric AG Webster, Justin A Richardson, Lindsay A Grant, and Robert K Henderson. Single-photon avalanche diodes in 90 nm cmos imaging technology with sub-1 hz median dark count rate. In *Int. Image Sensor Workshop (IISW), Onuma, Hokkaido*, 2011. 113
- [WSD⁺16] Peter O. Weigel, Marc Savanier, Christopher T. DeRose, Andrew T. Pomerene, Andrew L. Starbuck, Anthony L. Lentine, Vincent Stenger, and Shayan Mookherjea. Lightwave circuits in lithium niobate through hybrid waveguides with silicon photonics. *Scientific Reports*, 6:22301 EP –, 03 2016. 128
- [WVAA⁺16] Jeremy D Witmer, Joseph A Valery, Patricio Arrangoiz-Arriola, Christopher J Sarabalis, Jeff T Hill, and Amir H Safavi-Naeini. High-q photonic resonators and electro-optic coupling using silicon-on-lithium-niobate. *arXiv preprint arXiv:1612.02421*, 2016. 128
- [WXA⁺16] Cheng Wang, Xiao Xiong, Nicolas Andrade, Vivek Venkataraman, Xi-Feng Ren, Guang-Can Guo, and Marko Lončar. Second harmonic generation in nano-structured thin-film lithium niobate waveguides. *arXiv preprint arXiv:1610.04197*, 2016. 148
- [XPR⁺11] Chi Xiong, Wolfram Pernice, Kevin K Ryu, Carsten Schuck, King Y Fong, Tomas Palacios, and Hong X Tang. Integrated gan photonic circuits on silicon (100) for second harmonic generation. *Optics express*, 19(11):10462–10470, 2011. 128
- [XPS⁺12] Chi Xiong, Wolfram HP Pernice, Xiankai Sun, Carsten Schuck, King Y Fong, and Hong X Tang. Aluminum nitride as a new material for chip-scale optomechanics and nonlinear optics. *New Journal of Physics*, 14(9):095014, 2012. 51
- [XPT12] Chi Xiong, Wolfram HP Pernice, and Hong X Tang. Low-loss, silicon integrated, aluminum nitride photonic circuits and their use for electro-optic signal processing. *Nano Letters*, 12(7):3562–3568, 2012. 101

- [XSPL05] Qianfan Xu, Bradley Schmidt, Sameer Pradhan, and Michal Lipson. Micrometre-scale silicon electro-optic modulator. *nature*, 435(7040):325–327, 2005. [51](#), [127](#)
- [YB07] Joel KW Yang and Karl K Berggren. Using high-contrast salty development of hydrogen silsesquioxane for sub-10-nm half-pitch lithography. *Journal of Vacuum Science & Technology B*, 25(6):2025–2029, 2007. [67](#)
- [YHM⁺14] S-P Yu, JD Hood, JA Muniz, MJ Martin, Richard Norte, C-L Hung, Seán M Meenehan, Justin D Cohen, Oskar Painter, and HJ Kimble. Nanowire photonic crystal waveguides for single-atom trapping and strong light-matter interactions. *Applied Physics Letters*, 104(11):111103, 2014. [102](#)
- [ZJL⁺07] L Zhang, HC Jiang, C Liu, JW Dong, and P Chow. Annealing of al₂o₃ thin films prepared by atomic layer deposition. *Journal of Physics D: Applied Physics*, 40(12):3707, 2007. [150](#)
- [ZPM⁺16] Philip Zupancic, Philipp M. Preiss, Ruichao Ma, Alexander Lukin, M. Eric Tai, Matthew Rispoli, Rajibul Islam, and Markus Greiner. Ultra-precise holographic beam shaping for microscopic quantum control. *Optics Express*, 24(13):13881–13893, Jun 2016. [78](#)