Engineering Photons for Quantum Networks

Temporal-Mode Control, Optimization and Post-Selection Demonstrated in Trapped Ions

Carl Thomas

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Boris Blinov, Chair

Sara Mouradian

Xiaodong Xu

Program Authorized to Offer Degree:
Department of Physics

University of Washington

Abstract

Engineering Photons for Quantum Networks
Temporal-Mode Control, Optimization and Post-Selection Demonstrated in Trapped Ions

Carl Thomas

Chair of the Supervisory Committee:
Title of Chair Boris Blinov
Department of Chair

Control of the temporal waveform and Fock basis statistics of photons produced during spontaneous emission provides a crucial tool in the establishment of hybrid systems, optimal state transfer, interferometric stability and optimization of entanglement generation protocols for quantum networks. We describe novel methods to generate photons of any temporal waveform from emitters of any lifetime. Our broadly applicable approach has only two requirements for a candidate qudit: (1) control of the phase-parity and (2) modulation of the amplitude of a field coupling a ground state to an excited manifold which produces a photon during relaxation. We describe an approach to find optimal excitation pulse shapes, both numerically and experimentally, by employing variational algorithms to feedback on atomic populations. Additionally, we develop a quantum trajectory theoretical approach to determine emission statistics and establish tools for optimal post-selection to ensure maximum fidelity of photon generation protocols. We situate our work in the context of other prior research on bespoke single photon sources and networking including post-emission pulse shaping, temporal gating and cavity-based methods. In comparison, our free-space process has greater flexibility in producing any waveform, requires less infrastructure and can be readily applied across a wide domain of emitters of any frequency or lifetime. We validate our approach in a singly trapped $^{174}Yb^{+}$ ion and provide experimental results demonstrating

photon temporal waveform control.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

The writing of this dissertation has provided ample opportunity to reflect on all the people who have supported me throughout my life. I feel exceptionally lucky to have a family who have always encouraged me to pursue my interests. My parents worked very hard to provide me every opportunity in life. For that and so much more, I am profoundly grateful to them. Beyond my wonderful parents, my brother has been a constant source of friendship. He has always been there when I needed someone to call. Our regular conversations have helped keep me grounded through gray Seattle winters and long evenings in the lab. Most of all, I could not have completed this journey without my partner, Sarah. Moving from the East Coast to Washington together, where we knew almost no one, has been a great adventure, and I will always cherish our time together here, no matter where we end up next.

Scientifically, I have benefited tremendously from the guidance and mentorship of many people. Professors Marty Baylor and, especially, Joel Weisberg inspired me to continue my career as a scientist at a time when I was not sure what I wanted to do with my life. Without the graciousness of Dr. Neal Solmeyer to let a 20-something chemistry undergrad join his lab for a summer and learn about neutral atoms, I would not have ended up where I am today. After undergrad, it was at IonQ, in the early days, where I truly learned how to be an independent researcher. I learned so much from the team there, who were always patient, even when I asked for the third time how to fiber couple a laser. To Jonathan Mizrahi, Ken Wright, Neal Pisenti, Jwo-Sy Chen, Kristi Beck, and David Wong-Campos: thank you. To Kai Hudek, I owe a particular debt of gratitude. Kai taught me so much, not just about ion trap systems and optical design, but about how to manage the many challenges that research throws at us. Upon returning to IonQ a few years later for an internship, Ryan Bowler and

James Siverns also provided great role models.

At the University of Washington, there are so many others I need to thank. In my eyes, Alex Kato and Mila Zhukas set the standard for what graduate students are capable of. I have sought to emulate their example in the years since. Later, I was very grateful to have Rebecca Munk join our project. Rebecca is a kind and diligent scientist, and I hope our paths cross again. Jane Gunnell has been a great friend. Our conversations have helped me to see the big picture when the mundane, and unavoidable, challenges of research reared their heads. I have also enjoyed working beside the other members of our group, and although Gorby the Goldfish has passed on, she will always be with the lab in spirit. As I worked toward the completion of my PhD, my committee members Profs. Andrea Coladangelo, Xiaodong Xu, Kai-Mei Fu, and Sara Mouradian provided an excellent repository of knowledge and feedback.

Most of all, I need to thank my advisor, Boris Blinov. Boris has been a great mentor and project leader. Beyond teaching me a tremendous amount about ion trapping, he has shwon me how to deal with the ups and downs of research. He has been supportive and kind. I will miss our chats over coffee, hours spent in the lab together, and beers at Big Time. He empowers his students to take the lead in their own work, but is always present when help is needed. I feel profoundly lucky to have landed in his lab.

Chapter 1

WHY QUANTUM COMPUTING AND WHERE WE ARE NOW

Prior to the mid 2010s, few people outside of technical professionals and grad students buried deep in basements had ever heard of quantum computing. Fewer still had any idea what the field was really all about. Now it is hard to avoid. We have quantum fiber internet, quantum multivitamins, and no shortage of startups promising to lead the next quantum revolution, often with questionable credentials. Capital markets and state actors continue to pour billions into the field, while comparisons to the space race, the AI winter, or even the Manhattan project abound. A decade on, quantum information scientists no longer get a blank stare when asked what we do. Despite or because of this newfound quantum-mania, misconceptions and chicanery lurk behind every corner. But if we look beyond all the salesmanship, something profound IS happening.

Periodically for the last half-millennium, the world has been shaken up by a breakthrough in physics. Suddenly, new areas of exploration open up before us with novel technologies that revolutionize human life. We can trace these cataclysms back at least to Kepler whose observations of the night sky led to new navigational techniques, supercharging the Age of Exploration. Next came Newton whose development of mechanics and his famous three laws provided the substrate for the entirety of modern engineering and the natural sciences. The discovery of the laws of thermodynamics in the 1800s harnessed the power first of steam and then coal which built industrial societies. A half century later, the formulation of electromagnetism led to the advent of radio and telegraph with a new epoch of globalization greeting the 20th century as science itself became more deeply integrated.

The 1927 Solvay Conference on Electrons and Photons (pictured in Fig. 1.1) marked the

zenith of both the era of scientific internationalism and the last great scientific upheaval the world has witnessed: the first quantum revolution. The development of quantum mechanics threw open the doors to the atomic and information ages, with nuclear powering and digital computing being just two of the most prominent applications of the ideas of Einstein, Bohr, Curie and the rest of their coterie.

Now, nearly 100 years later, we *may* be witnessing the beginning of the next great leap forward. Quantum computers, or at least their biggest boosters, promise to unlock new techniques in material and climate science, as well as optimization, pharmaceuticals, and cryptography by solving problems that prove intractable on 'classical' devices [1]. But unlike the digital transistor used to implement the familiar 'bit', there is no single approach to realize an analogous quantum bit or 'qubit'. Instead, there are many ways to build a quantum computer. We can perform quantum logic operations by manipulating the electronic levels of a single atom, via control the state of a superconducting circuit loop or the interference of photons. Each qubit implementation from solid-state spin systems to trapped ions excels in some metrics desired for a quantum computer while lagging in others.

This thesis aims to contribute in some small way to the unfolding revolution in quantum technologies. Specifically, the bulk of this thesis deals with how to control photons from single qubit emitters to advance quantum networking. Such work could improve the performance of both same and mixed qubit type architectures. By engineering photons for specific tasks, we can improve the rate of quantum information transfer by mitigating noise, maximizing qubit-photon interaction and enabling distinct qubit types to 'talk' to each other.

The rest of Chapter 1 serves to provide context for these experiments as we briefly cover the history and concepts of both classical and quantum computing, address criteria for quantum computers and survey the state of the field to more deeply motivate the need for photon control before closing with a road map for the rest of this thesis. This chapter will be more informal in tone and style as we work to establish why anyone should care about these problems at all and provide a high-level conceptual framework to orient ourselves in the field. There will be plenty of space for stultifying, dense and 'properly scientific' contents in the

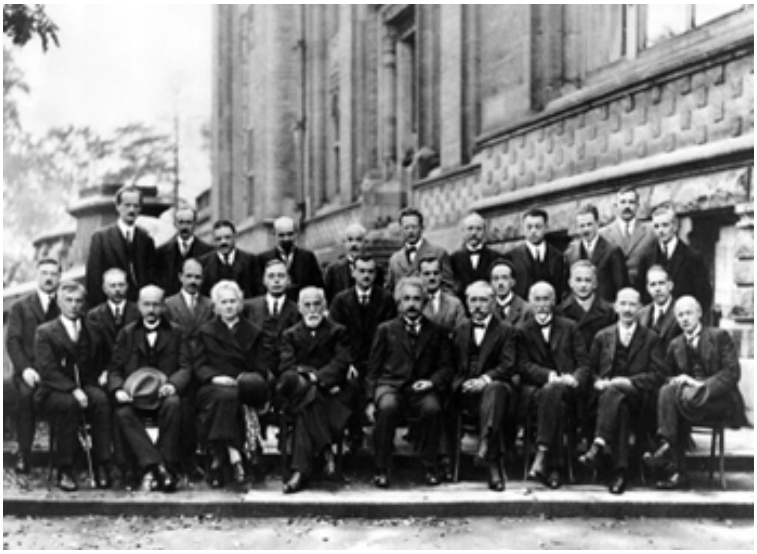

Figure 1.1: Attendees at the 1927 Solvay Conference on Electrons and Photons in Brussels, including Schrödinger, Pauli, Heisenberg, Debye, Bragg, Dirac, Compton, de Broglie, Born, Bohr, Planck, Curie and Einstein.

following chapters.

1.1 From classical to quantum computing

The problem statement for this thesis revolves optimal control over the properties of photons generated during spontaneous emission, and the application of these techniques to building quantum networks.

But before diving deeply into the world of networking, we should first establish what a quantum computer even is, how such a device differs from the classical machines we are more familiar with, and what the fundamental requirements to realize such a system are.

1.1.1 When classical computers fail

In order to motivate the need for quantum systems, we should first understand what a computer is, whether quantum or classical, at least in high-level terms. In its most basic

form, a computer is a device that reads a set of symbols, deterministically performs a set of operations based on these inputs and produces a new set of symbols representing the result of some computation. In modern digital computers, transistors acting as tiny switches, or bits, set to '0' or '1', assembled into a set of 'logic gates' perform these tasks by performing Boolean operations (that is, '0 AND 1 = 0', 'O OR 1 = 1'). These gates can be concatenated together in complex ways to solve many problems which would be totally impossible without digital computing.

Devices based on Boolean logic have proven incredibly useful for all types of applications; however, even the most powerful supercomputers struggle to predict the shape of folded proteins [2]. In fact, biochemical systems only represent one of the many areas where classical machines struggle. What makes certain tasks hard or slow while others take a fraction of second would be far too large of a question to address here. Indeed, an entire field, computational complexity, evolved to understand this question.

Still, we can at least get some intuition about which types of problems might be hard for a classical computer by recalling how these devices function. At its core, a bit is discrete: yes or no, up or down, 1 or 0. This logic works well when the when the system we want to model shares a similar structure. For example, an apple can be here or there, a baseball travels with a certain velocity, the market for pork futures has fixed inputs and outputs. However, this comfortable assumption that things have definite properties breaks down when we zoom in to the scale of atoms. Unlike our familiar macroscopic world, on the atomic scale systems things can be quite a bit stranger. A particle confined in a W-shaped potential energy environment may simultaneously exist in both the left and right hand sides. Similarly, atoms of a molecule vibrate back and forth like a series of balls connected by springs, but unlike their classical counterparts these molecules may vibrate with multiple 'speeds' at once. The non-discreteness of these systems means that modeling them on a traditional computer will be very inefficient because the quantum world operates on fundamentally different set of rules than the devices we employ to simulate them.

1.1.2 Introducing the qubit

The challenges of calculating atomic dynamics described above provided both the first inspiration to develop new computing systems and hints about how to build devices capable of the task. Feynman generally receives the credit for laying the conceptual foundations of the field and beginning to problematize a solution. We will take as starting point for our own survey of quantum computation his oft-quoted insight,

> *'Nature isn't classical, dammit, and if you want to make a simulation of Nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem because it doesn't look so easy.'*[3]

We now aim to build a light conceptual schema for how such a quantum simulation machine could be constructed while saving the more nuts-and-bolts implementation for later chapters.

Start by considering one of the simpler 'real' quantum systems: a single electron orbiting a hydrogen nucleus. The familiar form of the S, P, D and F orbitals describe the 'probability density function' of our electron. Already, here we encounter the non-discreteness of the quantum world. The probability density function represents the amplitude that, if observed, the electron will be found in any given region of space. In other words, until interaction with the environment, the electron does not exist at a single fixed location but rather is 'smeared out'; however, upon measurement the electron will always be located at just one point. The probability density function itself is derived from the electron's 'wave function'. At its core, the wave function just represents our method for describing the state of a system based on a few axioms of quantum mechanics. For example, the wave function for the electron orbiting the nucleus can be derived by assuming that the electron exhibits wave-like properties and then solving the time-independent Schrodinger equation with a Coulomb potential. The exact form of these solutions and the additional complexity brought upon by the inclusion of ideas like electron and nuclear spin are beyond the scope of this description; but, even this simplest example illustrates the key concepts at hand. Mathematically, the wave function is a normalized vector in complex Hilbert space. For our purposes, Hilbert space can be

taken as a vector space with an inner product, roughly meaning a distance metric exists between any two states. This space can either be finite or infinite dimensional depending on the quantum system in question. This description of a simplified hydrogen atom highlights two most fundamental axioms of quantum mechanics:

1. Superposition: a quantum system may occupy multiple states simultaneously

2. Wave function collapse: upon measurement, a quantum system in superposition will be forced to occupy only a single state.

Exactly what constitutes measurement has proven a very challenging problem to solve and continues to be one of the biggest open questions in physics; however, as a shorthand 'any thermodynamically irreversible interaction between a quantum system and its environment' usually suffices.

Returning to the example of the hydrogen atom, we can employ these central ideas of quantum mechanics to arrive at an ansatz about how to build a fundamentally new type of information processor. For the time being, we will ignore all orbitals except for the lowest energy S and P levels, and even here we will only consider the energy difference between the two orbitals and make no comment on angular momentum. In this simplified model, the electron exists entirely in one of the S and P orbitals or alternatively, the electron can instead exist in a superposition of both as represented in Fig. 1.2.

This allows us to introduce a simple shorthand for the state of our system:

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle; |0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}; |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tag{1.1}$$

where $|0\rangle$ and $|1\rangle$ correspond to the states where the electron occupies the S or P orbital respectively. We calculate the population of each state by taking $|\alpha|^2$ and $|\beta|^2$. The electron is always here or there, when we observe it so we can write the constraint $\alpha^2 + \beta^2 = 1$. Generalizing this to an n-level quantum systems, we have $|\Psi\rangle = \sum_{i=0}^{n-1} c_i |a_i\rangle$ with $\sum_{i=0}^{n} c_i^2 = 1$.

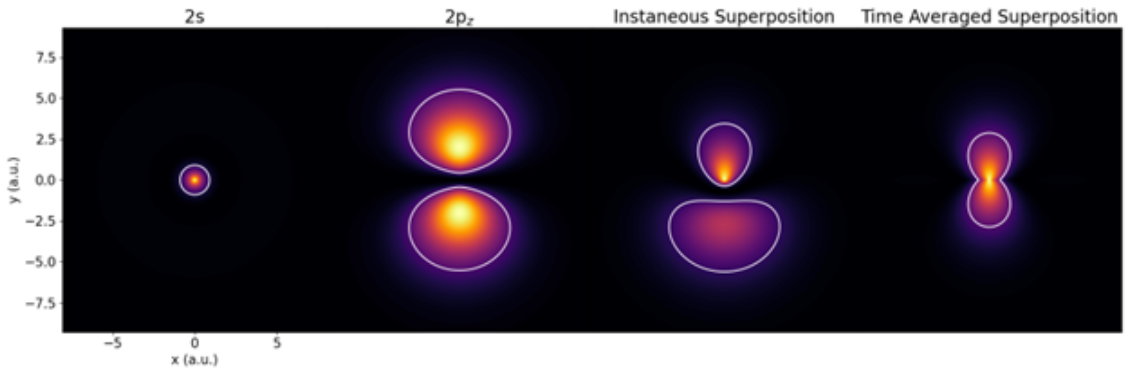

Figure 1.2: Visualization of superposition in atomic orbitals. The electron density of ^2S and ^2P_z orbitals are shown, as well as an instantaneous snap shot and time averaged picture of a superposition of the two orbitals. White contours bound the region within which, if measured, the electron will be found $\approx 80\%$ of the time. As the name 'wave function' implies, each orbital carries a distinct phase which oscillates at a characteristic frequency. In a superposition state, the beat note between the two orbitals leads to a time varying spatial density. The rate of oscillation of the superposition state also determines the frequency of the electromagnetic wave (usually a laser or microwave field) needed to couple the two atomic levels.

The matrix algebra used to describe quantum states and their interaction with physical forces also defines the conjugate transpose of $|\Psi\rangle$:

$$\langle\Psi| \equiv |\Psi\rangle^\dagger = (c_1^*,\ c_2^*,\ ..., c_n^*) \tag{1.2}$$

where * denotes the complex conjugate. Additionally, sometimes we will use density matrix formalism to denote a certain quantum states, known as pure states,

$$\rho = |\Psi\rangle\langle\Psi|\,. \tag{1.3}$$

We will define what instances constitute a pure state in Sec. 1.1.3. Refocusing on the simple two-level case, we realize that this system *almost* looks like a way to represent a classical bit,

except now rather than two discrete states, our system can take on any value between 0 and 1. If we can devise a way of controlling the populations of these levels, we could build a new type of quantum bit.

Before moving on, we need to address one other key way in which qubits differ from their classical counterparts. We said that quantum states exist as unit vectors in a *complex* Hilbert space. α and β are not constrained to be real values. Rather, we can incorporate a relative phase between $|0\rangle$ and $|1\rangle$. The most general form of a qubit in a pure state should instead be written as:

$$|\Psi\rangle = \alpha\,|0\rangle + e^{i\phi}\beta\,|1\rangle\,, \qquad \alpha, \beta \in \mathbb{R},\ \phi \in [0, 2\pi]. \tag{1.4}$$

The inclusion of this phase will be crucial for taking full advantage of the power of quantum gates and will factor heavily into further discussion in subsequent chapters. For now, we will just motivate the inclusion of phase by returning to the example of the hydrogen atom. In this instance, we recall that the S and P orbitals correspond to different energy levels with a frequency relationship given by Bohr's law:

$$E = \hbar\omega. \tag{1.5}$$

In this picture, the phase of the qubit advances at the transition frequency; however, it is standard to instead adopt a rotating frame of reference such that the phase of the qubit does not depend on time. Nonetheless, control of the relative phase between $|0\rangle$ and $|1\rangle$, either in the case of a single qubit or for two interacting qubits becomes crucial for implementation of non classical logic gates.

Bohr's law gives us another hint about how to implement a quantum computer. Not only does $E = \hbar\omega$ tell us about the relative frequency between the oscillations of the wave function of two orbitals, this tells us about how to couple these two states. If we start with our electron in $|0\rangle$ and expose this electron to an oscillating electromagnetic field whose switching frequency matches that given by Bohr's law, we will see the electron be resonantly driven

to the $|1\rangle$ state. Assuming a perfect two level system with no 'relaxation' (i.e. collapse of the higher energy state back to the lower level with the corresponding emission of a photon), we have now realized a conceptual picture for how to implement and control a single qubit. We can imagine turning on our EM field for enough time to fully rotate between $|0\rangle$ and $|1\rangle$, which we call a 'π' gate, for reasons which will soon be clear. Alternatively we could apply a $\frac{\pi}{2}$ gate, leaving our field on just long enough to prepare an equal superposition given by state $|\Psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

Taking this all together, we can finally introduce the most critical tool for thinking about single qubits: the Bloch sphere. This intuitive representation, derived from the Poincare sphere used to represent the polarization of light, allows us to represent any possible state of a pure two level quantum system as a point on its surface. The north and south poles are identified with the basis states $|0\rangle$ and $|1\rangle$, while all other points on the surface correspond to superpositions of the two. The polar angle encodes the populations of the basis states, and the azimuthal angle encodes their relative phase. Rotations of the state vector around the sphere then correspond to the action of single-qubit gates, providing an elegant geometric picture of qubit control. In this way, the Bloch sphere serves as the standard method for visualizing the full space of pure states available to a two-level quantum system. A representation of the Bloch sphere, along with the paths traced out by a state vector for several representative gates are shown in Fig. 1.3.

Employing Euler's formula, $e^{i\phi} = \cos(\phi) + i\sin(\phi)$ and mapping α and β to a unit circle (justified by the condition that total probability to find our qubit in one of the states to be 1) via $\alpha = \cos(\theta/2), \beta = \sin(\theta/2)$, we can rewrite Eq. (1.4) as

$$|\psi\rangle = \cos(\theta/2)\,|0\rangle + (\cos\phi + i\sin(\phi))\sin(\theta/2)\,|1\rangle\,. \tag{1.6}$$

We have already discussed a high level picture of how to implement rotations between $|0\rangle$ and $|1\rangle$, but what about rotations about $\hat{z}$, that is 'phase gates'? Well, as we know, the rate at which the phase between the two levels we use to encode our qubit evolves is determined by their relative energy spacing. We also know that the 'standard' approach is to project

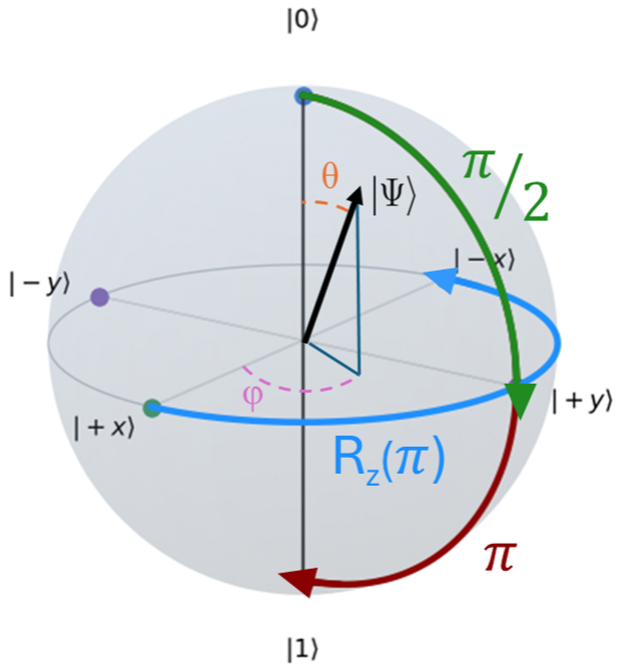

Figure 1.3: A representation of the Bloch sphere showing a quantum state, $|\Psi\rangle$, and its decomposition into polar and azimuthal angles, along with several exemplary rotations and phase gates.

into the rotating frame such that we can imagine our entire Bloch sphere spinning at this frequency so that for a single point on the surface of the sphere (a state vector), we can remove the time dependence of ϕ. A visualization of the rotating frame transformation is shown in Fig. 1.4. Thus, a phase gate can be implemented by temporarily exiting the rotating frame and allowing phase to accumulate for a fixed period state by the qubit frequency.

It turns out that we can also accomplish phase gates by modulation of our EM field, just as with the single qubit rotations described above. The difference is that instead of controlling the intensity and length of time our field is on to dial in an arbitrary superposition of $|0\rangle$ and $|1\rangle$, we instead change the phase of the EM wave [4]. In the rotating frame picture, changing the phase between the qubit levels and changing the phase of the field which couples those levels turns out to be completely equivalent! Finally, we have a full toolkit to prepare our quantum state on any point on the Bloch sphere. In other words, we have realized the full set of single qubit gates. In reality, the S and P orbitals of a hydrogen atom do not present an attractive qubit platform for many of reasons, too long to list now, but they still illustrate all the core requirements for such a system.

However, single qubit operations only represent part of the problem. If a computer could perform were flips between '0' and '1' without the ability to get these bits to interact, this would not be a very useful device. So our final task before leaving behind this conceptual survey of the basics of quantum computing is to clarify how two-qubits can be made to interact in interesting and non-classical ways. This would lay the groundwork for an entirely new set of gates based on quantum, rather than classical logic. For example, the use of conditional operations acting on qubits in superposition could allow bit-bit correlations that no classical circulate can replicate. By accessing new 'primitives', quantum devices create the opportunity for entirely new classes of algorithms which may outperform the best classical alternatives leading to computational speed-ups. The key word here is 'entanglement'.

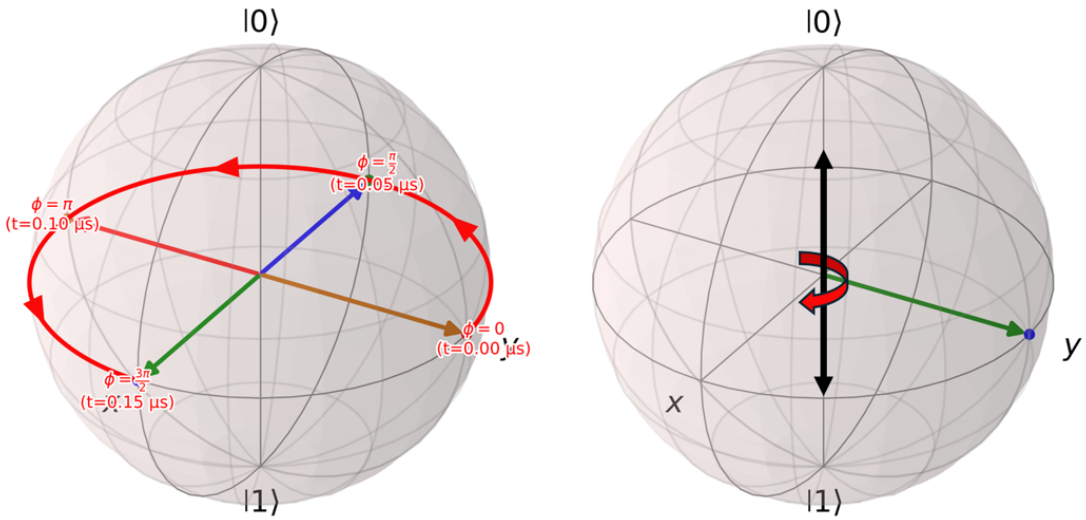

Figure 1.4: Projection into the rotating frame. Representation of the rotating frame used to remove time dependence for the phase between $|0\rangle$ and $|1\rangle$. In this example, the frequency splitting between the two levels is set to 5 MHz.

1.1.3 Understanding entanglement

Informally, entanglement captures correlations between subsystems that cannot be explained by any decomposition into independent local states. Formally a pure state of two particles, A and B, $|\Psi\rangle_{AB}$, is *entangled* if it is not separable, i.e., it cannot be written as $|\Psi\rangle_{AB} = |\psi\rangle_A \otimes |\phi\rangle_B$.

To build intuition, consider two-qubits prepared independently in equal superpositions, $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. The product state

$$|+\rangle \otimes |+\rangle = \tfrac{1}{2}\big(\,|00\rangle + |01\rangle + |10\rangle + |11\rangle\,\big)$$

yields four outcomes $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ with equal probability upon measurement in the computational basis. Correlations between states of the two-qubits are only the result of simple product rules.

By contrast, consider the state below:

$$|\Phi^+\rangle = \tfrac{1}{\sqrt{2}}(|00\rangle + |11\rangle)\,. \tag{1.7}$$

Unlike the product state, this wave function can not be factorized, or seperated into the state of two independent qubits. Measuring either qubit alone produces $|0\rangle$ or $|1\rangle$ with equal probability, but the *joint* outcomes are perfectly correlated: whenever the first qubit is found in $|0\rangle$ ($|1\rangle$), so is the second.

This is one example of what is called a 'Bell state' or a maximally entangled state for a two-qubit system. We will show in Ch.2 how entanglement can be used to generate interesting, non-classical gates. For the time being, we will simply provide example to understand how entanglement between quantum particles could actually be generated.

One context in which entanglement between particles routinely occurs can be found in the decay of electrons from excited orbitals back down to lower lying atomic states. Figure 1.5 illustrates such a process where in a single excited state, $|e\rangle$, may decay to two distinct lower energy states, $|0\rangle$ or $|1\rangle$. Each decay path leads to the emission of a photon with an energy corresponding to the difference between the ground and excited states. For example, the photon emitted when the electron collapses back to $|0\rangle$ may be in the blue range of the visible spectrum while that emitted when the electron ends in the $|1\rangle$ state could be red.

When a real decay process occurs, *both* of these decays path occur simultaneously with an amplitude determined by the relative strengths of each atomic transition. We should be careful to note that this process only generates a single photon, but this photon exists in a superposition of red and blue (at least until measured). But what does this mean about the state of the atom after the photon has been emitted? Well, because the production of a blue photon only occurs during a decay to $|0\rangle$ (and likewise for $|1\rangle$ with a red photon), the atom also exists in a superposition of two states. Because the state of the atom and photon are not separable and are actually causally connected (i.e.the emission of a blue photon causes the electron to collapse to $|0\rangle$), we now have a set of connected quantum particles whose state would be described by something looking like Eq. 1.7, albeit with the second '0' and

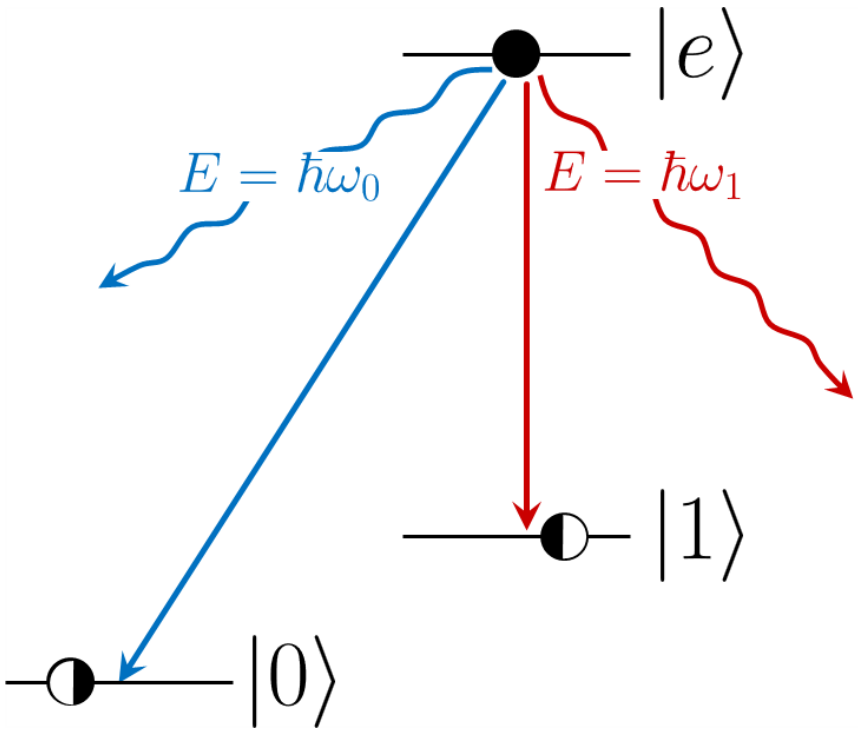

Figure 1.5: Level diagram for generating ion–photon entanglement in the frequency basis. An excited ionic state $|e\rangle$ can spontaneously decay to either ground-state qubit level $|0\rangle$ or $|1\rangle$, emitting a photon at frequency ω_0 (blue) or ω_1 (red), respectively. When a single photon is emitted during a decay process, it exists in a superposition of both frequencies, until measured. The frequency basis of the photon and logical basis of the atom are entangled in such a process.

'1' index swapped out for something like 'blue' and 'red.' In other words, the photon and atom are entangled! We can think of this entanglement as a wire which runs between our two particles. If we can figure out some clever ways of manipulating these entangled bits (after all, a photon can be a bit), then we can create new types of 2-qubit gates that not possible on a classical device. Then we could really unlock the power of quantum computing. But that will have to wait for a later chapter. Now, we will wrap up this crash course with some reflections on where the field of quantum computing stands in 2025.

1.2 Quantum computing, today

Ten years ago, virtually no large self described quantum computing companies existed (outside of D-wave, a heterodox type of computer based on quantum annealing). Some big corporations had small teams dedicated to the field, but outside of academia or government labs, the field was sparse. Now we have several major (> 100 employees) and *many* smaller players. But everyone seems to have a different idea about what the best way to build a quantum computer is. Unlike digital computers where the transistor reigns supreme, the field of quantum computing is characterized by the competition of many different architectures. Most qubit platforms have their fair share of promoters and detractors with each side eager to highlight why their platform of choice the best and why all the other platforms do not stand a chance. In recent years, we have also seen a greater interest in an alternative, the hybrid computer which has its own sets of challenges, addressed in later chapters.

Despite (or because of their differences), the field needs metrics to compare the performance across qubit types. Such metrics serve to transform an apples-to-oranges (or ions to dots) comparison to a universal language of 'how well does this thing do what I want it to?' Of course, there is a lot riding on 'what I want' so the question of which metrics are useful or even meaningful remains hotly debated. I will avoid some of the more contentious figures of merit promulgated in recent years (does anyone really understand what 'Quantum Volume' is?) and instead just touch on the most commonly employed criteria which (almost) everyone can agree to including on their spec sheets.

1.2.1 How to characterize a quantum computer

At a high level, the usefulness of a quantum computer to solve a particular problem encoded in a certain quantum cicruit (ie, series of gates) comes down to three major factors:

1. How 'good' are your qubits?

2. How many qubits do you have?

3. How are your qubits connected to each other?

Each of these questions hide within it many other questions. For example, when we want to quantify the 'goodness' of a qubit we could include metrics like how long our qubit will stay in a prepared state (coherence time), how fast we can perform an operation on one or two-qubits (single and multi qubit gate speed), how accurate are those gates-meaning how much control do we have over where the qubit state ends up (gate fidelity). Even for each of these criteria, we can dig much further. Coherence time, for example, refers to how fast your qubit state gets scrambled. This time scale typically ranges from μs to hours. Importantly, it can be both intrinsic to the qubit type (called T1, and set by how long it takes for one qubit level to collapse to the other in an atomic qubit, for example) or instead a measure of external control sources or environmental noise (or T2) where in phase information between qubit levels is effectively lost.

All of this is to say, there are *many* ways to characterize a quantum computer and providing an exhaustive list would be beyond the scope of this high level overview. Instead I will simply provide a short list of some of the most common figures of merit with brief definitions:

Coherence time ($\mathbf{T_1}$, $\mathbf{T_2}$): How long a qubit can store quantum information before it decoheres due to environmental noise or intrinsic relaxation. Typically expressed as

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \tag{1.8}$$

where T_2 represents the how long a definite phase between the qubit levels is maintained. T_1 represents the excited qubit state lifetime and T_ϕ quantifies the rate at which environmental noise (or experimental imperfections) contribute to effective dephasing (or loss of information about phase, from the experimenters perspective).

Two-qubit gate speed: How fast entangling operations can be applied. Gate speed, in tandem with coherence time, limit the depth of a quantum circuit (meaning how many gates can be applied to each qubit) as coherence must be maintained across the circuit. This leads to the condition

$$N_{gate}\tau << T_2 \tag{1.9}$$

where τ is the time required to run a gate and N_{gate} is the number of two-qubit entangling gates that can be performed. The speed of gates depends on factors such as the intrinsic state coupling strength and available power to stimulate transitions between qubit levels.

Two-qubit gate fidelity: The probability the intended gate is performed correctly. In other words, if the the experimentalist looks at the state of the two-qubits after performing a gate (or more realistically, performs the gate 1000 times, records the outcome of measurements and reconstructs the density matrix), what will be the distance in Hilbert space between the target and resultant quantum state? We typically quantify the fidelity as

$$\mathcal{F} = Tr(\rho_{ideal}\rho_{actual}) \tag{1.10}$$

with ρ_{ideal} and ρ_{actual} denoting the target and measured state density matrix. We can then write the corresponding error as by $\epsilon = 1 - \mathcal{F}$.

Qubit number: How many qubits are available for a quantum circuit on a given device. The qubit number determines the maximum width (given by how many parallel 2 qubit gates can be run simultaneously) of a circuit and is important for assessing what actual problems can be solved by a device. Qubit number alone does not provide a deep insight into the power of a device as a few 'good' (high fidelity, long coherence time) qubits may be able to perform a computation that many more bad qubits cannot. For this reason, there has been

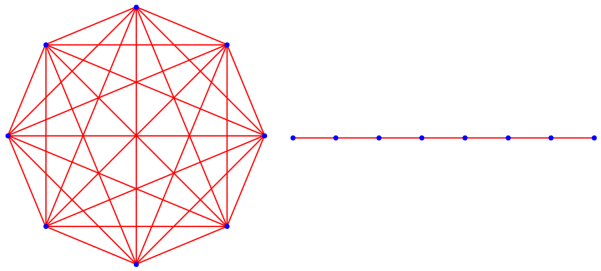

Figure 1.6: Schematic qubit connectivity graphs. Blue nodes denote qubits; red edges indicate native two-qubit couplings. **Left:** fully connected (all-to-all) topology typical of small trapped-ion registers, enabling direct two-qubit gates between any pair. **Right:** linear nearest-neighbor chain, representative of simple 1-D arrays of superconducting or neutral-atom qubits, where distant interactions require entanglement swapping, impacting circuit depth and fidelity. Most real nearest neighbor architectures will be represented by 2-D graphs with a variable number of neighbors based on physical implementation. For example, we could consider a honey comb like hexagonal array or a simpler square array.

an increasing push to define more platform agnostic measurements of effective computational power such as 'algorithmic' or 'logical' qubits.

Connectivity: How qubits are topologically connected in terms of ability to generate entangling gates. In other words, if each qubit is represented by a node on a graph, edges are drawn between pairs of qubits which can be directly entangled. The density of nodes and edges represents the relative ease or difficulty of connecting any two-qubits. Denser architectures tend to require less overhead to run a given circuit (i.e.no need for piece wise entanglement swapping). Representations of all-to-all and nearest-neighbor qubit conectivities are shown in Fig. 1.6.

To this preliminary list of qubit platform metrics, we could add things such as remote entanglement generation rate, state preparation and measurement errors and many more. But for now, these metrics provide enough granularity to get a high level overview of the state of different architectures.

1.2.2 Is one qubit platform 'the best'?

The answer to this section header is easy: No! At time of writing, no one qubit platform unambiguously out competes all the other platforms at any of the metrics we care about. Some typical performance metrics for a subset of architecture types is displayed in Tab. 1.1. The particular reasons which one platform excels at one metric while lagging behind in another ultimately comes down to either fundamental physics or maturity of the technology. Unpacking the explanations behind these strengths and weaknesses of each architecture is well beyond the scope of this thesis. Suffice it to say, so long as no qubit stands above the rest, we have arrived at our motivation for the hybrid quantum system.

	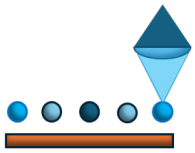		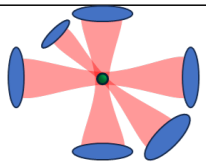	
	Trapped Ions	**Superconductors**	**Neutral atoms**	**Solid State**
Coherence Time (s)	10–10^3	10^{-3}	10^{-1}–10	10^{-6}–10
2 Q Gate Speed (μs)	10^{-3*}–10^3	10^{-3}–10^{-1}	10^{-3}–10	10
2 Q Fidelity	0.99–0.9999	0.9–0.99	0.9–0.99	0.9–0.99
Qubit # (physical)	10–10^2	10–10^3	10–10^4	10
Local connectivity	All to all	Nearest neighbor	Rydberg limited	Lattice limited

*Typical 2-qubit gate speeds for trapped ions are on the order of tens to hundreds of microseconds; ultrafast gates have also been demonstrated [5].

Table 1.1: Comparison of physical qubit platforms. Top row shows schematic icons for each platform, followed by performance metrics.

1.2.3 Quantum networks, hybrid computers and beyond

We have established in short why one might want to build a hybrid quantum computer. Now the operative question becomes: 'how could one build a hybrid quantum computer'? We can imagine device of the future: fast gates realized on solid state qubits in a form of Quantum Processing Unit while the results of these computations stored in ions with long coherence times like a memory bank. A conceptual representation for such a device is shown in Fig. 1.7. To build such a device, we must determine how to connect the constituent units.

To date, the most common approach to interfacing distinct qubit registers have aimed to use photons. Once again, a more in depth explanation of exactly how photons can be used to remotely entanglement diverse qubit types will be saved for a future chapter. In order to motivate the problems this thesis seeks to address, we will provide only a very brief sketch

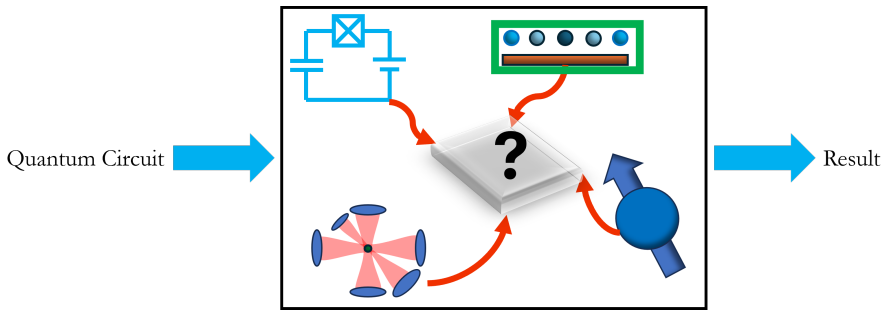

Figure 1.7: An agnostic quantum network. Future universal quantum computers will likely leverage the strengths of multiple qubit platforms. How to interconnect distinct architectures remains an active area of research.

now.

Returning to the description of entanglement given in Sec. 1.1.3, let us consider what would happen if we could build a network consisting of two such atomic nodes (call them Alice and Bob), each of which is capable of producing photons as envisioned in Fig. 1.5. We can collect photons from each node using standard optics like lenses and fiber optic cables. Importantly, the photons emitted from each node are perfectly identical. That means that if we look at the spectra of both photons on a diffraction grating, measure their polarizations or look at the amplitude of each photon's electric field in time, it would be impossible to say if any given photon came from Alice or Bob.

Now imagine that both nodes can be made to produce a photon at the exact same time. We collect both photons, mix them up on a 50:50 beam splitter and then split the red and blue components of each output port by using dichroic mirrors (an optic which reflects certain wavelengths while transmitting others). The output of these dichroics send these photons to a set of detectors which count coincidences between of different photon color pair combinations.

Because our photons and atoms are entangled, we know that if we detect two blue photons, we can write the state of our atoms as $|\Psi\rangle = |0_{\text{atom }0}; 0_{\text{atom }1}\rangle$. Similarly, if both photons are red, we can write $|\Psi\rangle = |1_{\text{atom }0}; 1_{\text{atom }1}\rangle$. But what happens if we detect one red photon and

one blue photon? We cannot tell whether the red photon or blue photon came from either Alice or Bob. In fact, if we have prepared our interferometer correctly, the reality is that neither photon came from either atom alone. Instead, both photons were produced by a superposition of both atoms. But, because we produced one photon of each color, we know that if we look at the state of the atoms, we must find either $|\Psi\rangle = |0;1\rangle$ or $|\Psi\rangle = |1;0\rangle$. In other words, because we have independently maintained entanglement between each atom and each photon and because we have interfered photons from both atoms such that at after mixing our photons on a beam splitter we cannot say which photon came from which atom, we have managed to swap entanglement from the two sets of atom-photon pairs such that the two atoms have become entangled.

This basic experiment represents the key to building a quantum network, distributed quantum computing and hybrid devices. Provided the photons from each node are identical, it does not matter if we have a trapped ion at each node, a solid state qubit or one of each. Thus if we can coax any two-qubits to produce identical photons, we can, in principle, entangle them.

The work described in the rest of this thesis seeks to establish new techniques for networking including controlling the properties of individual photons to generate indistinguishable photons. Specifically, I deal with how to modulate the temporal profile of photons from any quantum emitter, how to calculate emission statistics for photon generation process and how to select events which optimize various photon network protocols. It turns out these techniques may prove useful not just for hybrid system but also for same-type qubit networks as photons of different shapes may be ideally suited for different applications or to mitigate various error sources. For example, the protocols I describe in subsequent chapters can be used to engineer photons which are optimally robust to phase drifts across networks or which have the best chance of being absorbed at a second node in direct quantum state transfer experiments. We also describe the novel application of established methods coming from Quantum Trajectory Theory to calculate photon statistics during atomic excitation processes for experimental optimization in quantum networks.

1.3 The structure of this thesis

Chapter 2: Trapped ion quantum computing is a 'building block' chapter where in we cover the core criteria to realize a quantum computer in greater detail, discuss how trapped ion systems can satisfy these criteria, describe the theory behind ion trapping and establish how a system is actually built.

Chapter 3: The interaction of single atoms with fields of light is another 'building block' chapter. Here we cover the theory critical to understand the photon control experiments which represent the novel work of this thesis.

Chapter 4: Optimal control of single photons expands on the previous chapter to describe how we can actually apply numerical techniques to generate photons of any temporal profile from any emitter and how we can characterize photon statistics from an excitation process and use them to develop protocols for optimal post selection of experimental shots. We also discuss the application of these techniques to hybrid and same-qubit type system in more depth and situate our work in the context of other approaches to photon control. This chapter provides a variety of key simulational results.

Chapter 5: Generating shaped photons from an ion node: describes how our theoretical and numerical work can be brought to fruition in an ion trap. We detail how to actually implement photon control in a qubit platform and present results demonstrating control of the photon temporal waveform, the first such work in a free space quantum emitter.

Chapter 2

TRAPPED ION QUANTUM COMPUTING

Our goal in this chapter is to establish the criteria for *any* quantum computer, regardless of qubit platform, describe how these criteria can be implemented for trapped ion devices, provide a brief foray into theory relevant for ion trapping and finally detail the experimental realization of a system suitable for our photon control experiments.

2.1 DiVincenzo's criteria

As the field of quantum information began to pass from an era of purely conceptual discussion to the first realizations of quantum gates in the mid 1990's [6], the need to establish a consistent set of criteria for the construction of a physical quantum computer became apparent. DiVincenzo first articulated a comprehensive list of such requirements in 1996 [7]. These five original criteria (plus two added later and related to the networking of devices) have become the standard list for assessing the viability of a qubit platform. The criteria are as follows:

1. A scalable physical system with well-characterized qubit

2. The ability to initialize the state of the qubits to a simple fiducial state

3. Long relevant quantum coherence times

4. A universal set of quantum gates

5. A qubit-specific measurement capability

6. The ability to interconvert stationary and flying qubits

7. The ability to faithfully transmit flying qubits between specified locations.

These seven items are relevant because, taken together, they specify a device that has the logical structure necessary to encode (1), instantiate (2), control (4) and read out (5) qubits which are the basic elements necessary to run quantum circuits. Long coherence times (3) mean that such a device will both have the ability to store quantum information for later access in a circuit and the rate at which fidelities degrade due to intrinsic and extrinsic sources will not prevent running of deep circuits, or ultimately error correction. Finally, scalability (1 again) alongside transduction to (6) and distribution of (7) photons allow the only clear path forward so that small prototypes can grow into large, modular systems and be networked. Each criterion is necessary but not sufficient on its own; dropping any one prevents either universal control, reliable error-corrected operation, or system-scale growth [8].

We will now move onto concretizing how these criteria can satisfied using our qubit platform of choice, the trapped ion.

2.2 The trapped ion qubit

Trapped ions earned their reputation as a leading qubit type in large part because essentially every DiVincenzo requirement was either first demonstrated, or first demonstrated at high fidelity, in this platform. Laser cooling and optical pumping provide reliable initialization. State-dependent fluorescence allows near-ideal readout. Long-lived internal states deliver record coherence; and multiple entangling-gate families (Cirac–Zoller; Mølmer–Sørensen) establish universal control. Beyond computation, ions were among the first systems to interconvert stationary and flying qubits and to distribute entanglement over optical links, needed for the two 'network' criteria. Table 2.1 summarizes the key results for ions in the context of DiVincenzo's criteria. Subsequent subsections provide *some* greater depth on the realization of each criteria in ions.

Criterion	Solution in Trapped Ions	Notes
Scalable, well-characterized qubits	Qubits are encoded in the stable electronic states of trapped ions (hyperfine, Zeeman or optical levels). Ions are confined in RF Paul or surface electrode traps. Scaling pursued via linear chains, 2D arrays, and photonic interconnects between separate traps.	Ions are identical by nature. Modular architectures use ion shuttling (i.e. QCCD) [9] and/or photonic networking. Qubit registers with 10s of physical qubits have been demonstrated [10].
Qubit initialization	Optical pumping prepares ions in a ground hyperfine/Zeeman sublevel with near 100% fidelity.	Initialization is fast (< 100 μs). Typically performed with resonant lasers (e.g., 369 nm for Yb^+).
Long coherence times	Coherence times of 10–10^3s [11, 12], orders of magnitude longer than gate times [13].	Clock states (magnetically insensitive) yield the best coherence.
Universal gate set	Achieved with laser-driven single-qubit rotations and multi-qubit entangling gates mediated by shared motional modes. Work with magnetic field gradient gates have pushed fidelity to >0.99999 [14].	2-qubit gate speeds tend to be in μs regime.
Qubit readout	State-dependent fluorescence gives readout by detecting photons scattered only from one logical basis state with errors as low as a few parts per million [15].	Fast readout with low measurement errors require high numerical aperture (NA) optics.
Ion-photon entanglement	Ion qubits mapped to photons via spontaneous emission. Photons encode qubit state in polarization, frequency, or temporal bases.	Ion photon entanglement has been demonstrated in various species including: Cd^+ [16], Yb^+ [17], Ca^+ [18], Sr^+ [19] and Ba^+ [20].

Photon networking	Ion-emitted photons coupled into single-mode fibers for transmission. Frequency conversion bridges to telecom bands. Remote entanglement via photon interference.	Entanglement demonstrated between ions separated by meters to kilometers. Fastest remote entanglement rate demonstrated at 250 Hz [21].

Table 2.1: DiVincenzo's criteria for a quantum computer and how they are addressed in trapped ion systems.

2.2.1 Optical-Metastable-Ground: one ion, many qubits

Logical basis states can be encoded into trapped ion energy levels in a variety of ways. Two orbitals coupled by a quadrupole transition, for instance, can be used to encode a so-called 'optical' qubit. Alternatively, different hyperfine levels of a single ground orbital (like $^2\mathrm{S}_{1/2}, \mathrm{F} = 0$ and $^2\mathrm{S}_{1/2}, \mathrm{F} = 1$ in $^{171}\mathrm{Yb}^+$) constitute a 'ground' state qubit. Similar hyperfine or Zeeman levels in a metastable state also form a good basis for logic states. In recent years, significant attention has been paid to how these three distinct ways of writing a qubit into the electronic levels of an ion can be leveraged depending on the task at hand in 'Optical-Metastable-Ground' (or OMG) architectures [22]. A OMG encoding scheme for $^{137}\mathrm{Ba}^+$, a popular qubit choice, is shown in Fig. 2.1. For example, ground state qubits typically have the longest coherence times with gates realized using microwave horns (hyperfine splittings tend to be in the 100s of MHz to 10s of GHz regime), two photon Raman transitions [23] or magnetic field gradients [24].

Meanwhile, optical qubit encodings are usually ideal for ion-photon transduction or networking protocols. Metastable qubits, on the other hand, present a convenient way to store the state of one qubit while gate operations are performed on other qubits in a register.

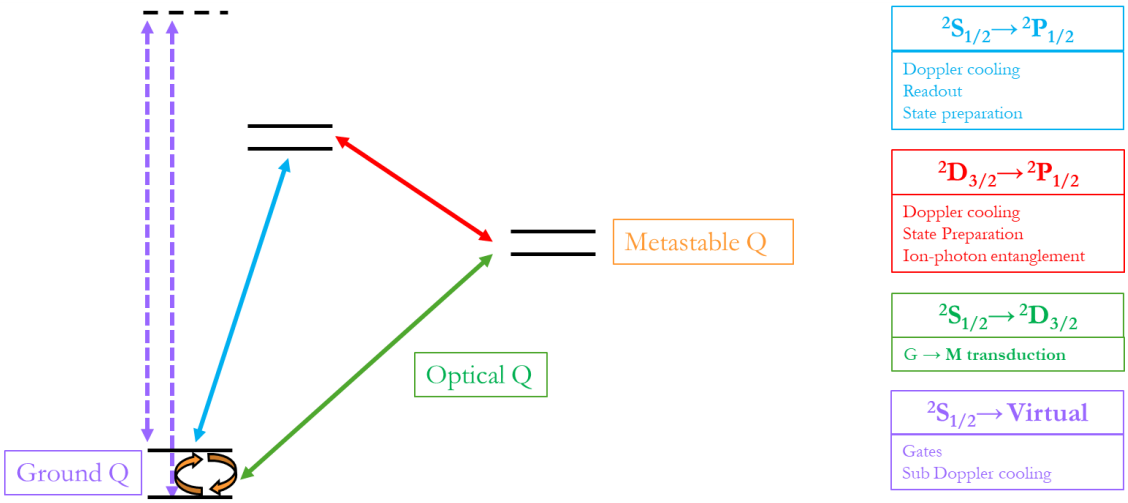

Figure 2.1: A simplified representation of OMG qubit encoding for the case of a $^{137}Ba^+$ ion

2.2.2 State initialization

In trapped ion qubits, state initialization can readily be achieved via leveraging state-dependent forces and a set of lasers with well defined frequencies and polarizations to drive an electron into a single state. Taking the example of $^{171}Yb^+$, a common choice of ion for gates, the application of a laser resonant with the transition from the $^2S_{1/2}, F = 1$ to $^2P_{1/2}, F = 1$ will result in a qubit prepared in the $^2S_{1/2}, F = 0$ state as shown in Fig. 2.2. Similar protocols exist for a wide variety of other ions.

2.2.3 Long-lived qubits

As we have discussed, trapped ion qubits have the longest coherence times demonstrated of any qubit platform. The fundamental limit on these times is derived from the relaxation rate between the two orbitals which can be on the order of years for strongly forbidden transitions. However, environmental dephasing stemming from magnetic field noise or the linewidth of lasers used to implement gates tends to practically limit coherence times. Still, the appropriate selection of qubits encoded in first order field-insensitive levels readily allow

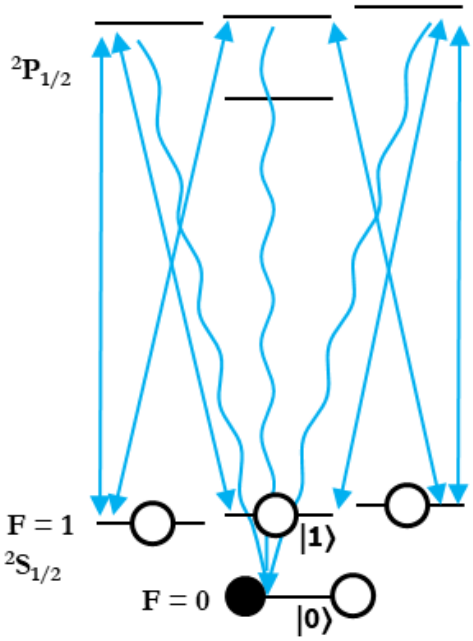

Figure 2.2: State initialization scheme in $^{171}Yb^+$. Open circles represent levels with some non-zero population before initialization begins. The solid black circle represents the state of the ion after initialization.

coherence time of seconds or longer.

2.2.4 Single-qubit gates

We have already briefly touched on (in both Chapter 1 and Sec. 2.2.1) how single qubit gates can be implemented trapped ions. The basic approach is to identify a set of energy levels with appropriate metrics (i.e. coherence time, coupling strengths etc.) to assign a set of logical states. Adopting the Bloch sphere picture, rotations from $|0\rangle$ to $|1\rangle$ can then be performed via a number of means using constant wave resonant lasers, two-photon transitions with a beat note at the qubit frequency or microwave horns. Phase gates can be accomplished by either exiting the rotating frame set by the qubit splitting and allowing phase accumulation or through virtual gates by adding the desired phase to the local oscillator which couples the transition at future times.

2.2.5 Multi-qubit gates

The realization of two-qubit gates is sufficiently more involved than their single-qubit counterparts to merit a brief schematic treatment. Ions are confined in radio-frequency traps by ponderomotive forces from oscillating electric fields, the details of which we will address later. For gates, the key resource is that ions in a common potential share quantized normal modes of motion. These collective modes act as a quantum bus that can mediate entangling interactions between internal (qubit) states.

A simple visualization is shown in Fig. 2.4. Consider two ions with a common motional mode frequency ω_m. In order to generate entanglement between the logical basis states of two qubits without unintentionally changing the shared motional state of the ion chain, we must operate in the Lamb–Dicke regime. Practically, we can assess whether we operate in this regime by considering the inequality

$$\eta^2(2n+1) \ll 1, \qquad \eta = \sqrt{\frac{\hbar}{2m\omega_x}}\, k_x, \tag{2.1}$$

where ω_x is the trap frequency along the x-axis (discussed later in the section on ion-trap

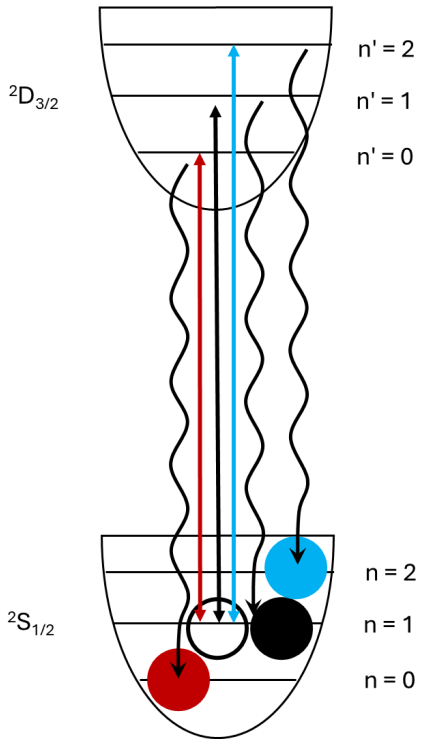

Figure 2.3: Representation of motional sideband transitions in the Lamb–Dicke regime for a narrow quadrupole transition. The open black circle indicates the initial state of the ion in motional level $n = 1$. By applying laser pulses on the first red or blue motional sideband (colored double sided arrows), the ion is excited to the upper electronic state while changing the motional quantum number by $\Delta n = \pm 1$. In the Lamb-Dicke regime, spontaneous decay from the excited state (curved arrows) occurs predominantly without changing the motional state ($n' = n$), so the ion returns to the original electronic level while its motional state has been changed in a controlled way. Repeating these cycles allows preparation of the motional ground state via resolved sideband cooling and provides a mechanism for generating entanglement between ions that share a common motional mode. To resolve the sidebands, the optical transition linewidth must be narrower than the trap frequency, so a quadrupole transition or an effective narrow two-photon Raman transition is typically used.

theory), n is the motional quantum number, m is the mass of the ion, and k_x is the projection of the laser wave vector on the ion in the x-direction. In other words, satisfying this inequality corresponds to the case where spontaneous decay occurs at the qubit carrier transition frequency and changes in the motional state of the ion by more than one quantum are suppressed. Thus, by tuning lasers to the red or blue motional sidebands of an optical transition, we can implement controlled changes of the ion's motional state. A representation of transitions in the Lamb–Dicke regime is shown in Fig. 2.3.

After ground-state cooling and in the Lamb–Dicke regime, we apply a sequence of sideband pulses (the Cirac–Zoller construction). Starting from $|0,0\rangle \otimes |n=0\rangle$, a blue-sideband $\pi/2$ pulse on ion 1 at angular frequency $\omega_0 + \omega_m$ (ω_0 is the qubit splitting) prepares

$$\frac{1}{\sqrt{2}} \left(|0,0\rangle \otimes |n=0\rangle + |1,0\rangle \otimes |n=1\rangle \right).$$

A red-sideband π pulse on ion 2 at $\omega_0 - \omega_m$ transfers the phonon to the second ion's spin, yielding (up to a global phase)

$$\frac{1}{\sqrt{2}} \left(|0,0\rangle \otimes |n=0\rangle + |1,1\rangle \otimes |n=0\rangle \right) = |\Phi^+\rangle \otimes |n=0\rangle .$$

Thus the shared motional mode has served as a bus to entangle the two qubits while returning the motion to its ground state.

'Modern' two qubit gate implementations tend to be a bit more complex than the simple illustrative example given above. However, the use of shared motional states of the ion chain continues to provide the basis for two qubit gates.

2.2.6 State readout

Similar to state initialization, readout makes use of state-dependent forces that only couple a subset of the pertinent atomic energy levels for a particular qubit. Returning to the common choice of $^{171}\mathrm{Yb}^+$ for our qubit as shown in Fig. 2.5, a laser can be applied which only couples the $^2\mathrm{S}_{1/2}, \mathrm{F}=1$ to $^2\mathrm{S}_{1/2}, \mathrm{F}=0$ states in a cycling transition, so called because the states form a closed subspace from which the electron will never, or at least extremely rarely, be pumped

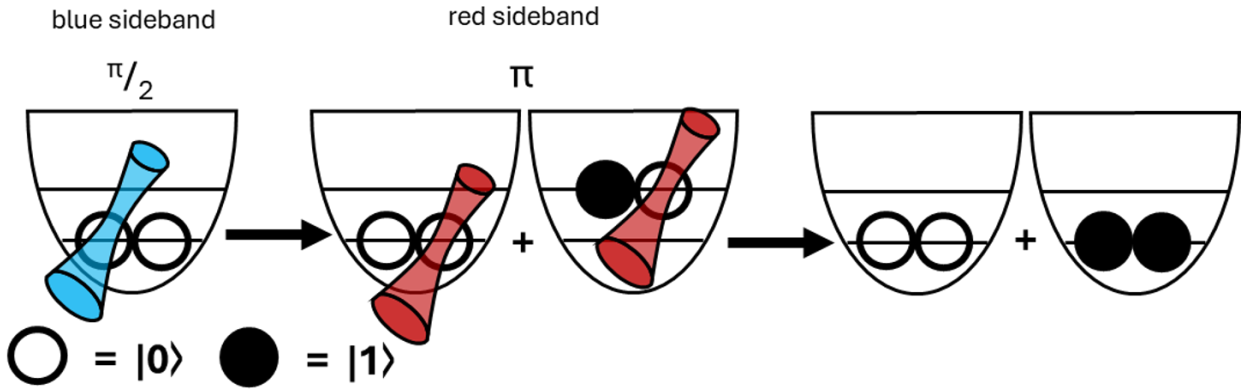

Figure 2.4: A schematic representation of 2-qubit entanglement implemented via a shared motional mode with red and blue sideband pulses. Adapted from [25].

out of. Thus the ion will only scatter light in the instance that it has some probability amplitude in the $|1\rangle$ state. By counting the number of photons detected in a certain period and making inferences based on the background signal and expected ionic scatter, the state of the qubit can be read out. The relevant levels for the case of $^{171}Yb^{+}$ and a conceptual photon detection histogram are shown in Fig. 2.5. In the instances that the qubit is not simply in either $|1\rangle$ or $|0\rangle$ but instead the state vector lies somewhere else on the surface of the Bloch sphere, repeated measurements along with appropriate rotations and phase gates can be used to reconstruct the full qubit density matrix in a process referred to as state tomography.

2.2.7 Ion-photon entanglement

The need for a candidate qubit to satisfy first five DiVincenzo criteria is relatively self explanatory. However, the motivation for criteria 6 and 7 are *slightly* less apparent. These two instead relate to the networking of multiple devices together. The fundamental rationale for these requirements can be understood in analogy to the need for interconnection between classical computers. While a single classical computer may prove useful for solving particular

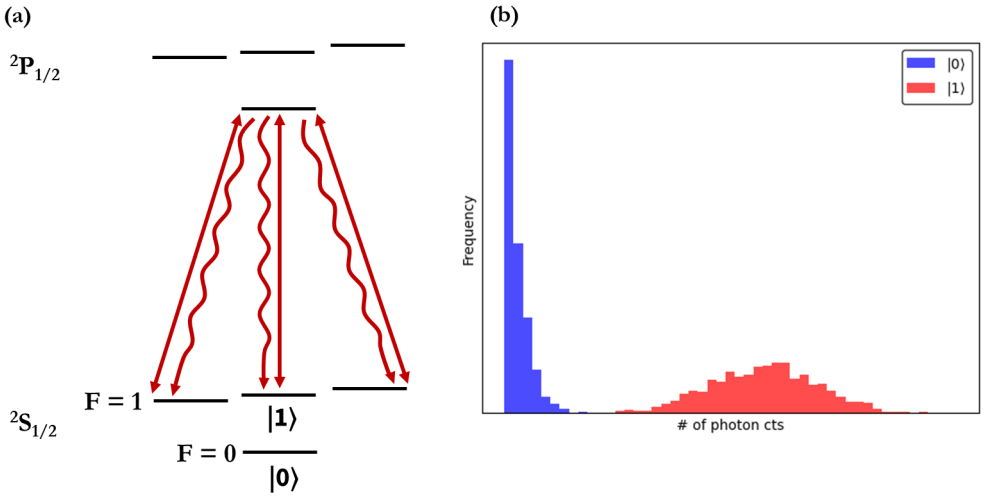

Figure 2.5: State discrimination in a $^{171}Yb^+$ ion. (a) A cycling transition only scatters photons if the qubit is in $|1\rangle$. (b) A histogram of photon counts, integrated over a certain time period, is used to determine which logical basis state the measured ion was most likely in.

problems, the real power of modern computing relies on the ability to network many devices together. In other words, just as many of the most important applications of today's conventional computers require access to a shared mechanism by which they can communicate, the quantum computer of the future will be one which has its own quantum network. In fact, due to the intrinsic limitations of many qubit platforms, such a quantum internet would not just be a convenience but is rather a necessity to solve problems beyond the scope of classical devices.

For example, a single trapped ion register relies on the ability to prepare a linear chain, or in some proposals a 2D crystal, of ions. However, we cannot simply load more and more ions into a single trap indefinitely. Architectural constraints put an upper limit on the number of qubits which a single device can support. The voltages necessary to maintain a line of nearly equally spaced ions on a chip trap grow as more and more ions are added to the chain, increasing the Columbic forces between qubits. Additionally, the need to cool many vibrational modes enforces constraints as the density and number of modes increases with each ion. Ultimately, it becomes infeasible, or impossible, to add more useful qubits to a single register.

Approaches based on a Quantum Charge Coupled Device (QCCD) architecture, wherein ions are shuttled between multiple registers and multiple operation-specific zones in a single trap have gained significant ground in the last 10 years. However, even here, practical considerations such as vacuum chamber volume and complexity of optical system infrastructure will prevent any single trap from scaling indefinitely.

Instead, it becomes requisite to identify and implement other mechanisms to allow distant qubits to interact. Thus far, the only viable approach identified is to use photons as a mediator to encode quantum information from a 'stationary' qubit. These photons can then be distributed using the mature technology of fiber optics to generate remote entanglement between devices or transfer quantum states between nodes.

The first step then in achieving such a quantum network relies on mapping the state of a stationary qubit onto its 'flying' (i.e. photonic) counterpart in a process known as

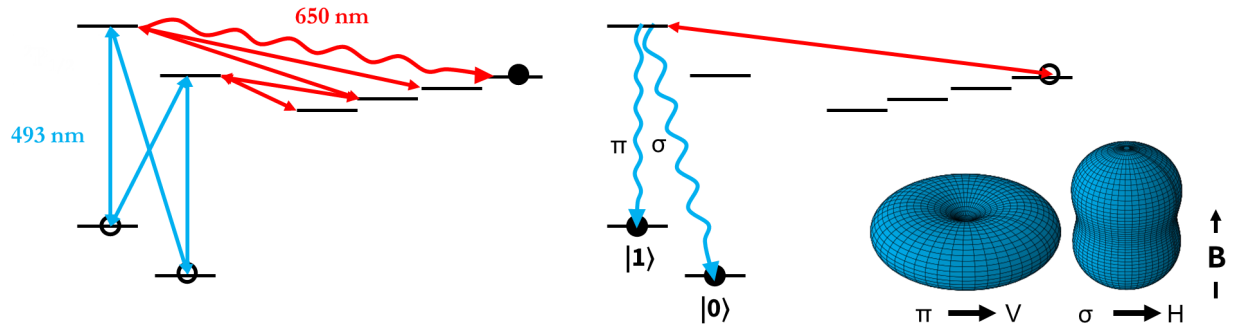

Figure 2.6: Simple polarization basis ion-photon entanglement scheme for $^{138}Ba^{+}$. **Left:** the ion is prepared in an edge state of $^{5}D_{3/2}$ manifold through application of appropriately polarized 493 nm 650 nm light. **Right:** An excitation pulse of σ^{-} polarized 650nm light drives to a single Zeeman level of the $^{6}P_{1/2}$ state which decays, emitting a photon in a superposition of polarization states entangled with the atom logical basis. Surfaces show the relative electric field amplitude for the two polarizations of photons and the mapping to lab frame states.

transduction. In section 1.1.3, we described one approach to generate entanglement between an atom and a photon, using the color, or frequency, of a photon to encode the state of the atom. We will describe the two other dominant paradigms for atom-photon entanglement.

Ion-photon entanglement, and by extension the first hybrid quantum system, was initially realized via entanglement of the polarization basis of a flying qubit with the logical state of an ion [16]. Figure 2.6 shows how polarization based entanglement can be generated for a $^{138}Ba^{+}$ qubit with logical states written into the Zeeman sublevels of the $^{2}S_{1/2}$ manifold. The application of appropriately polarized 493 nm and 650 nm light can be used to prepare the ion in an edge state of the $^{5}D_{3/2}$ manifold. A pulse of 650 nm light can then excite the electron to the $^{2}P_{1/2}, m_s = +1/2$ level. The electron will then decay back to a lower energy state, emitting either a 493 nm or 650 nm photon. In the instance that a red photon is measured, the shot is thrown out and the process is repeated. However, when a 493 nm photon is emitted, this photon exists in a superposition of polarization states corresponding to a π (i.e. $\Delta m_s = 0$) or σ (i.e. $\Delta m_s = 1$) transition. The dipole emission pattern of these two photons are distinct, allowing the two decay paths to be translated into lab frame vertical and horizontal polarizations, respectively, with relative amplitudes determined by the atomic coupling strengths, calculated by the corresponding Clebsch-Gordon coefficients. Thus, the polarization state of the photon and logical basis state of the ion are entangled. Appropriate rotations via the use of polarizing optics and microwave gates on the ion can be used to verify entanglement by measuring correlations between the state of the photon and the ion and comparing the results to the limits imposed by Bell's inequalities. The process of ion-photon entanglement via the polarization method for $^{138}Ba^{+}$ is shown in Fig. 2.6.

The third major method of ion-photon entanglement instead maps the qubit state to the time bin of photon detection [26]. In such experiments, the qubit is first prepared in an equal superposition of logical states. A strong excitation coupling only one of the qubit levels to an excited state is applied. A π pulse then rotates the qubit state followed by another excitation pulse. The time bin corresponding to when the photon was emitted is then entangled with the qubit and the use of delay lines allows reconstruction of the full ion-photon state.

2.2.8 Photon networking

The final of DiVincenzo's criteria takes the need for ion-photon entanglement a step further to establish how such entangled photons can be used for entanglement swapping, that is generating remote entanglement between two distant ions. Here we describe one approach for to this criteria.

We sketch a standard two-photon, polarization-based entanglement-swapping protocol. Two identical ion nodes are operated synchronously. Each node prepares ion–photon entanglement of the form

$$|\Psi\rangle_A = \tfrac{1}{\sqrt{2}}(|0\rangle_A |H\rangle_A + |1\rangle_A |V\rangle_A), \quad |\Psi\rangle_B = \tfrac{1}{\sqrt{2}}(|0\rangle_B |H\rangle_B + |1\rangle_B |V\rangle_B),$$

with photons launched into single-mode, polarization-maintaining fibers. The optical paths from both nodes are mode-matched and interfered on a 50:50 non-polarizing beam splitter; each output feeds a polarizing beam splitter and single-photon detectors (Fig. 2.7). When the two photons are *indistinguishable* in all other degrees of freedom (arrival time/temporal mode, spectrum, spatial mode, polarization), interference erases which-path information. Certain two-fold detection patterns then herald a maximally entangled Bell state of the remote ions. In the case of coincident clicks in (H, V) on *different* output ports project the ions onto $|\Psi^-\rangle = (|0,1\rangle - |1,0\rangle)/\sqrt{2}$ up to a relative phase. Other analyzer choices project onto the Bell states $|\Psi^\pm\rangle$ and $|\Phi^\pm\rangle$, defined in the computational basis as

$$|\Phi^\pm\rangle = \frac{1}{\sqrt{2}}\big(|0,0\rangle \pm |1,1\rangle\big), \tag{2.2}$$

$$|\Psi^\pm\rangle = \frac{1}{\sqrt{2}}\big(|0,1\rangle \pm |1,0\rangle\big). \tag{2.3}$$

Events that do not correspond to the Bell state signatures are discarded.

In practice, only a small fraction of emitted photons are collected and transmitted with collection efficiency, fiber/optic loss, and detector quantum efficiency bound the heralding rate. Nevertheless, the protocol is robust: successful two-fold detections *herald* remote entanglement independent of local photon state-detection efficiencies. Variants replace polarization with

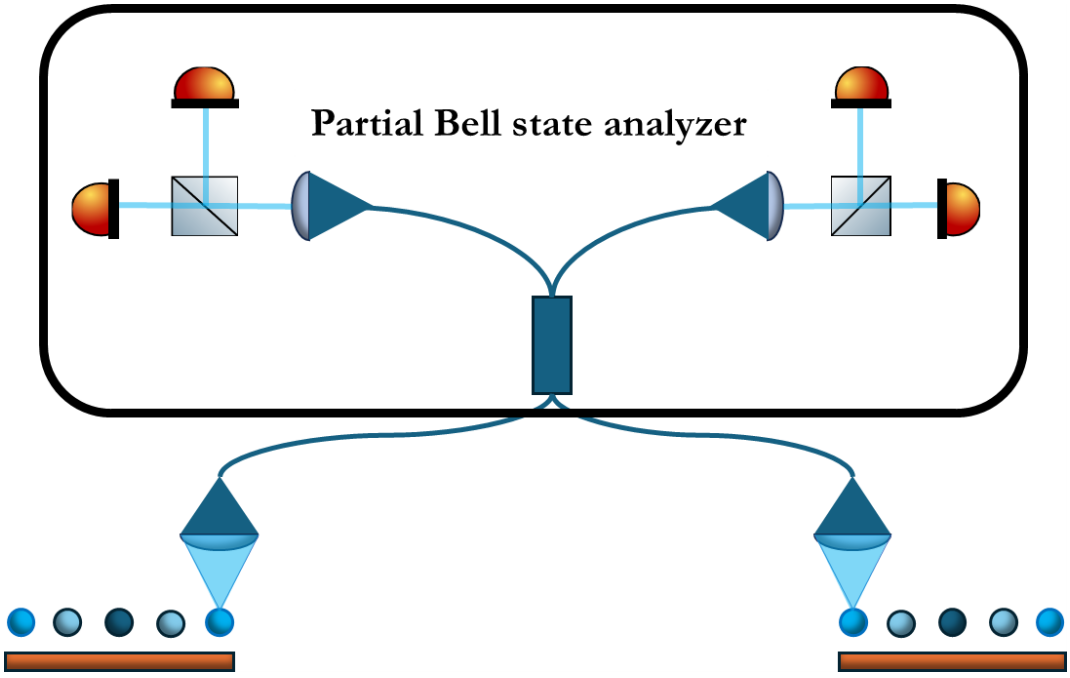

Figure 2.7: A simplified representation of ion-ion entanglement between two chip ion nodes. Ions at the end of each trap are used as communication qubits. Appropriate excitation protocols lead to the emission of a photon from each node, entangled with the logical state of the ion qubit. Photons are combined at a 50:50 non-polarizing fiberized beam splitter. Provided perfectly synchronized excitation pulses, identical ideal collection optics and no other sources of interferometric error such as path length fluctuations, the resultant photons will be indistinguishable. The output mode of the splitter has equal amplitude contributions from both ions, projective measurements of photon polarization on a set of polarizing beam splitters is used to herald the generation of a Bell state of the communication ions.

frequency-bin or time-bin encodings while retaining the same two-photon interference logic. A complementary 'single-photon' scheme (weak excitation and a single detector click) also heralds remote entanglement but requires phase-stable interferometry between the nodes.

Once a communication ion at each node is entangled, local entangling gates swap that entanglement onto target memory qubits in the registers, enabling arbitrary inter-node connectivity. With fiber switches and quantum repeaters, the same building block extends to larger-scale quantum networks [27].

2.3 Ion trapping theory

Thus far we have established in broad strokes how a quantum computer, and quantum network, can be built using trapped ions as building blocks. However, we have not yet addressed how an ion trap itself can be created and how the additional infrastructure necessary to do interesting experiments can be built. Of course, most of the work of an experimentalist deals with these two tasks and as such this thesis could not seriously be considered complete without an introduction to the basic aspects of ion trapping.

A central advantage of ions, relative to many other qubit platforms, is that they are *charged* and therefore couple strongly to electric fields. These fields can be used to control and contain our qubit of choice. However, purely static fields cannot produce a stable three-dimensional trap for a charged particle for reasons we will soon discuss. Practical ion traps therefore use either (i) static electric fields together with a strong static magnetic field (Penning traps), or (ii) time-varying radio-frequency electric fields to generate an effective confining potential (RF traps). Our group, and most trapped ion quantum information groups, rely exclusively on RF traps. Still, Penning traps continue to be an active area of research and noting the similarities and differences between the two approaches, at least in brief, is worthwhile.

Penning traps use static electric quadrupole fields together with a strong, uniform magnetic field B to confine charged particles. Axial motion is harmonic from the DC potential, the same as in RF traps. However, radial confinement arises from the $E \times B$ term, yielding cyclotron motion. In comparison to RF traps, this method of confinement does not have to

deal with the intrinsic forces produced by a time varying field which leads to an intrinsic drive of charged particles referred to as micromotion. This enables very large, stable 2D ion crystals via a rotating-wall potential useful for analog quantum simulation and precision spectroscopy. Disadvantages include high magnetic field strengths and 2D crystals which rotate in the lab frame [28, 29]. By contrast, RF traps avoid large magnetic fields, offer flexible segmented-electrode designs and fast entangling gates, but must manage micromotion and and higher heating rates compared to their Magnetic field counterparts.

2.3.1 Earnshaw's theorem

Before introducing the RF Paul trap and the language of pseudopotentials, it is useful to make precise what cannot work. A single, charged particle cannot be confined in three dimensions by purely static, curl-free electric fields in free space. The relevant result is Earnshaw's theorem. In short, if the charge density in a given region is 0 (not including the charge of the ion itself), stable minima, or maxima cannot be formed. An intuitive 'proof' for this statement relies on Gauss's Law. Recalling that

$$E = -\nabla\Phi, \tag{2.4}$$

we see that a potential minimum in some volume would require lines to terminate at that point. In other words, we could enclose the minimum in a spherical shell with field lines flowing in but not out. Such a situation would present a contradiction because the very presence of net field flux through the surface implies the existence of charge within.

However, free space can support saddle points in the electric field which we will find is the next best thing. Electric field lines can be focused along one axis, but must be divergent in the transverse direction. Figure 2.8 shows one example of an electric field with such saddle points. In this snap shot, the black surfaces are electrodes held at ground, while the red surface is an electrode to which a potential oscillating at radio frequencies is applied. At the moment shown, the RF electrode has a negative voltage applied. We see the positively charged ion will be pushed towards the red electrode. The 2-D slice represented in the above

figure only shows part of the picture. We can complete our trap in a variety of ways. The first RF traps consisted of a radially symmetric 'ring' electrode which could be generated by rotating the shown slice by π radians about the y axis. This yields a trap where the RF null corresponds to a single point located at the origin. Alternatively, by extruding the shown plane out of the page, we can realize a linear RF trap with a null that consists of a line. This picture is the basis for the standard 4-rod trap design and its descendants, which we rely on for our photon shaping experiments. As such electrode cannot extend infinitely, two additional DC electrodes, usually referred to as end caps, are oriented along the z-axis with the distance between the electrodes determining the axial trap dimension. Similarly, if $\vec{z}$ sets the axis dictated by the direction between the two end caps of a 4-rod trap, by projecting into a plane parallel to $\vec{z}$, we can arrive at the geometry of a chip trap.

2.3.2 An RF solution

Without the ability to trap ions with static fields, our approach will be based on switching the polarity of the saddle point in time. A quadrupole potential that is not confining in all directions at any instant can still produce stable effective potential when the rate of switching is fast compared to the particle's response. This is the basis for the Paul trap.

As our experiment relies on the use of a linear 4-rod trap, we will describe the potential for this trap. Assigning the coordinate system as shown in Fig. 2.9, we can write the RF and DC potentials at the center of linear trap as [30]

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

and

$$\Phi_{DC}(x, y, z) = \frac{\alpha U_{DC}}{2z_0^2}(2z^2 - x^2 - y^2) \tag{2.6}$$

Here, r_0 and z_0 are the characteristic trap radial and axial dimensions given by ≈ 0.3 and 1.5 mm respectively. Ω is the RF frequency. V_{RF} and U_{DC} are the alternating and static voltages applied to the rods and end caps accordingly. κ and α are geometric factors dependent on the particular shape of electrodes. We see that the potential points in opposite directions along x

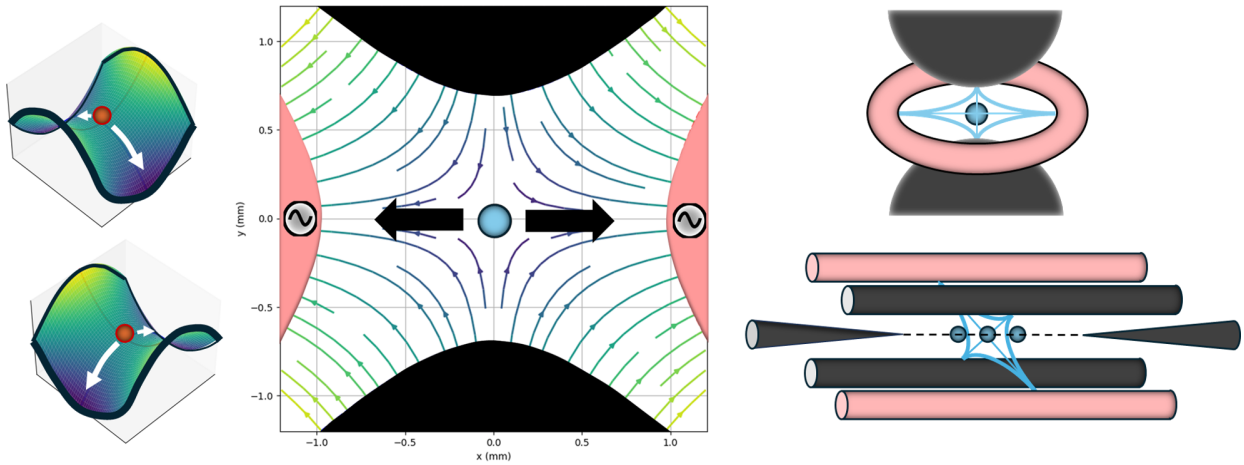

Figure 2.8: Representations of the potential and geometry of a quadrupole RF ion trap. **Left:** Instantaneous RF quadrupole potentials at opposite phases ($+V_{\mathrm{RF}}$, $-V_{\mathrm{RF}}$) form orthogonal saddle surfaces: one focuses in x and defocuses in y; a half-cycle later the roles reverse. **Center:** Time-dependent field lines in the radial $x-y$ plane between the RF electrodes (pink) and DC electrodes (black). The rapid alternation produces a time-averaged ponderomotive pseudopotential that confines the ion near the center, yielding slow secular motion superposed with drive-synchronous micromotion. Arrows show field lines. **Right:** Two possibly geometries of the RF trap. Top shows a 'ring' trap where the RF null is located at a single point. Bottom shows a standard 4-rod trap, which forms a RF null along the trap axis.

and y, but switches with the trap frequency. In such a device, relevant for many macro and chip traps, axial confinement is provided by endcap electrodes held at a static voltage.

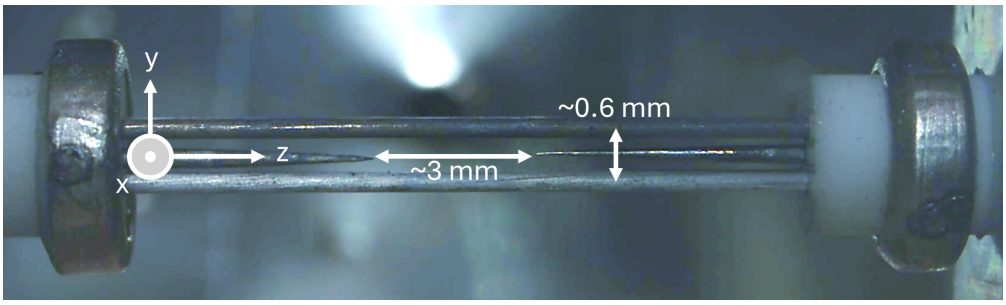

Figure 2.9: The 4-rod trap used in photon control experiments described in later chapters. Rods at two opposite vertices have a potential applied which oscillates at a radio frequency will the remaining two rods are grounded. Two electro-etched needles, centered along the axis of the RF null, provide axial confinement

We can derive the equations of motion and pertinent trap parameters for the 4-rod trap. Using $\vec{F} = q\vec{E} = -q\nabla\Phi$, with q being the ion charge, we arrive at:

$$m\ddot{x} + qx[\frac{\kappa}{R_0^2}(V_{RF}\cos(\Omega t)) + \frac{\alpha}{z_0^2}U_{DC}] = 0, \tag{2.7}$$

$$m\ddot{y} - qy[\frac{k}{R_0^2}(V_{RF}\cos(\Omega t)) - \frac{\alpha}{z_0^2}U_{DC}] = 0 \tag{2.8}$$

and

$$m\ddot{z} + qz\frac{\alpha}{z_0^2}U_{DC} = 0 \tag{2.9}$$

for an ion of mass m.

Recalling the form of a simple harmonic oscillator, we see that the equation of motion for the z coordinate can be rewritten in the familiar way, $\vec{F} = ma = -k\vec{z}$ with the oscillator frequency given by

$$\omega_z = \sqrt{\frac{k}{m}} = \sqrt{\frac{2\alpha q U_{DC}}{mz_0^2}}. \tag{2.10}$$

Thus, along the trap axis, we have harmonic motion.

The situation is somewhat more complex in the radial direction where our equation of motion depends on the RF frequency, Ω. The radial motion of ions is described by the Mathieu equations whose derivation is beyond the scope of this section. In the case of a perfectly quadrupole potential formed by hyperbolic electrodes of the type employed in the first RF traps constructed by Wolfgang Paul [31], these are the exact solutions. For the case of linear rod trap like ours, they instead provide a good first order approximation. We will simply restate the first order solution for the case of radial motion which can be found in many resources on ion trapping such as 'Ion Traps' by P. Ghosh [32]:

$$x(t) = x_0 \cos(\sqrt{a_x + \frac{q_x^2}{2}}\Omega t)[1 - \frac{q_x}{2}\cos(\Omega t)]. \tag{2.11}$$

Here q_x and a_x are the Matheiu parameters given by

$$q_x = \frac{2qkV_{RF}}{mr_0^2\Omega^2} \tag{2.12}$$

and

$$a_x = \frac{4q\alpha U_{DC}}{mz_0^2\Omega^2}. \tag{2.13}$$

A similar equation will describe motion along y with the condition $q_y = -q_x$ and $\alpha_y = \alpha_x$. We see that our motion along x and y coordinates is the product of two distinct frequencies. So-called micromotion at the RF frequency arises from the presence of the $\cos(\Omega t)$ term. The remaining 'secular' motion is slower as both α_x and q_x^2 are much less than 1.

The amplitude of micromotion is reduced by the $q_x/2$ term, hence the name. Ignoring the micromotion allows us to treat the remaining motion as though it were harmonic and we can then write the corresponding radial frequency as

$$\omega_x = \frac{\Omega}{2}\sqrt{a_x + \frac{q_x^2}{2}}. \tag{2.14}$$

Experimentally, it is often easier to determine trap frequencies based on the spacing of ions in a linear crystal, which can be used to calculate axial frequencies by comparing the repulsive of ion-ion interaction to the restoring force needed to keep ions at the measured distance. Radial frequencies can then be calibrated by increasing axial confinement (or reducing radial

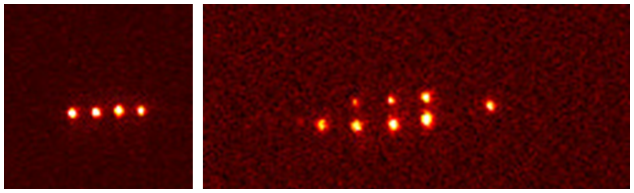

Figure 2.10: Ion crystals trapped in our 4-rod system. Left: a linear chain of 4 ions with ∼15 micron spacing, suitable for measuring axial mode frequencies. Right: a 2D crystal forms when the ion chain buckles as axial confinement is increased.

trap strength) until the chain buckles. Two ion crystals trapped in our system and suitable for such a calculation are shown in Fig. 2.10.

Finally, we should note that real trap potentials are rarely quite so clean. Besides corrections to form of potentials based on electrode geometries, stray fields often displace the ion equilibrium position. Such effects may arise due to patch potentials on electrode surface induced by atomic deposition or UV charging [33]. Displacement from the RF null leads to a number of deleterious effects as the ions are pushed up against the side of the trap psuedopotential. Beyond reducing the effective trap depth, displacement also leads to 'excess micromotion' which in turn can create problems for chain cooling as ion motion leads to an instantaneous Doppler shift whose magnitude may lead to inefficient cooling as the applied laser field effectively becomes chirped in the frame of the ion. A conceptual representation of intrinsic and excess micromotion, as well as secular motion, is shown in Fig. 2.11. Compensation electrodes can be used to cancel stray fields by generating an equal and opposite additional potential term. Some of the consequences of excess micromotion will be addressed in Ch 5.

2.3.3 Laser cooling

We have now seen how charged particles can be confined in space with a ponderomotive potential. However, when ions are first loaded into a trap (a process described in Sec. 2.4.4),

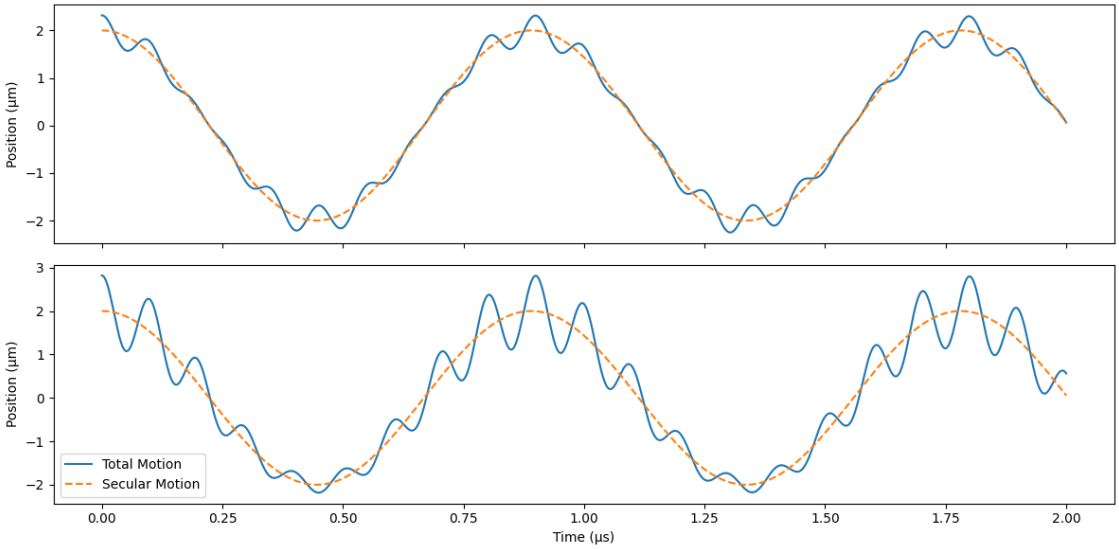

Figure 2.11: A conceptual representation of secular motion and micromotion of a trapped ion for cases of **Top:** no excess micromotion and **Bottom:** excess micromotion present. The fast ripple at the RF drive (10.845 MHz) is micromotion.

ions tend to be quite hot with thermal distributions corresponding to velocities on the order of hundreds of meters per second. At these temperatures ions are poorly localized and adding more ions yields a diffuse cloud rather than a crystal. For quantum information, we need individual optical addressing and state readout, and ultimately chains near the motional ground state for high-fidelity gates. As such, we need a way to cool our ions down. Ion cooling is typically divided into two regimes based on the minimum temperatures (or motional quanta) achievable, namely Doppler and sub-Doppler cooling.

The first approach, used in virtually every trapped ion quantum information experiment, is Doppler cooling [34]. This form of laser cooling relies on a velocity dependent Doppler shift which an ion 'sees' when interacting with a light field. The basic idea, described in greater depth below, is that a laser red detuned from an atomic transition is applied to a hot ion. As the ion moves around the trap, the ion is more likely to absorb a photon when its direction of motion has some component antiparallel to the k-vector of the cooling beam as this leads

to a beam closer to resonance in the ion frame. Each absorption imparts a momentum kick opposite the ion's motion. The ion then emits a photon, receiving another kick opposite the direction of photon emission to conserve momentum. Averaging over many emissions, the direction of flight of the resultant photons is isotropic. In other words, we see that the average of photon frequencies as measured in the lab frame in this process matches the non-detuned transition frequency. What this means is that the photons emitted over many cycles will be slightly 'bluer' than the photons absorbed. In the lab frame the ion is losing energy in the only way it can, through a reduction of its thermal motion. Cooling continues until the kicks from absorption and emission are equal. For efficient cooling, care must be taken to ensure the cooling beam has a projection along all the principal trap axes to ensure cooling in all directions.

To understand the limits of Doppler cooling, we first need a brief introduction to some atomic physics. The essential ingredients for Doppler cooling are:

1. a strong, nearly closed cycling transition with natural linewidth Γ

2. laser light red-detuned from the atomic resonance

3. projection of the laser wavevector $\vec{k}$ onto each princiap axis

If we take the case of a two-level atom driven by monochromatic light, we first define the saturation parameter, $I_{sat} = \frac{\pi\Gamma\ hc}{3\lambda^3}$ where Γ, h, c and λ are the transition linewidth, Plank constant, speed of light and photon wavelength respectively . The saturation intensity can be understood by reference to the excited-state lifetime, $\tau = 1/\Gamma$. The lifetime quantifies the average time an electron remains in the excited state before decaying and emitting a photon via spontaneous emission[1]. Because an electron can only spend at most half its time in the

[1]Actually, although the use of a photon counter would give the impression that such decay is discrete, with the probability for a photon to be emitted in some time interval given by a Poisson distribution, the real picture is a bit more complex. If the atom is prepared in the excited state at $t=0$, the excited-state population decays as $P_e(t) = e^{-t/\tau}$, while the emitted *field amplitude* has an envelope $|E(t)| \propto e^{-t/(2\tau)}$ (so the photon *intensity* decays as $e^{-t/\tau}$). The corresponding spectrum is Lorentzian with full width at half

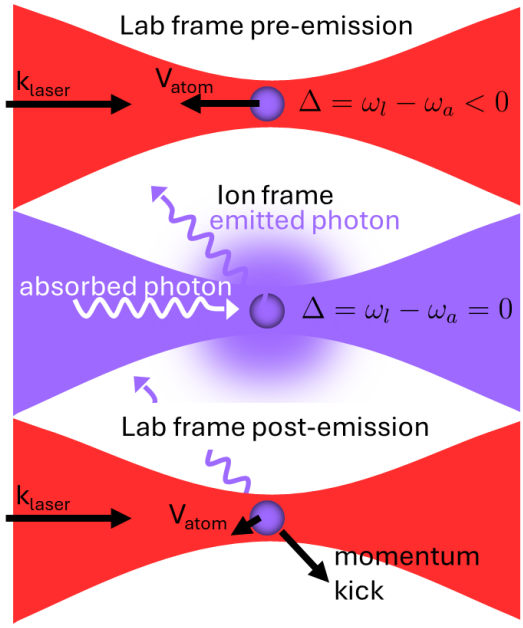

Figure 2.12: Illustration of Doppler cooling. **Top (lab frame, pre-emission):** an ion moving with velocity $\vec{v}$ encounters a red-detuned laser ($\Delta = \omega_L - \omega_0 < 0$); for ions moving *toward* the beam, the Doppler shift brings the light closer to resonance. **Middle (ion frame at absorption):** the light is resonant ($\Delta \approx 0$), the ion absorbs a photon and receives a momentum kick $-\hbar\vec{k}$ opposing its motion; subsequent spontaneous emission is nearly isotropic so its average momentum is ≈ 0. **Bottom (lab frame, post-emission):** the net effect per cycle is a damping force that reduces the ion's speed. Repeated cycles cool the motion toward the Doppler limit $k_B T_D = \hbar\Gamma/2$; sub-Doppler techniques are required to go below this limit.

excited state, the maximum rate at which photons can be generated via spontaneous for a transition is $\Gamma/2$. This occurs when the transition is full saturated.

The scattering rate can be calculated with

$$R_{\rm sc}(\Delta, s_0) = \frac{\Gamma}{2}\,\frac{s_0}{1+s_0+\left(\frac{2\Delta}{\Gamma}\right)^2}\,. \tag{2.15}$$

where Δ is the detuning from the transition frequency and s_0 is the saturation parameter given by $I/Isat$. The corresponding scattering force is then

$$\vec{F_{sc}} = \hbar\vec{k}R_{sc} \tag{2.16}$$

where $\hbar\vec{k}$ is the momentum of each photon.

Spontaneous emission randomizes the photon recoil direction. The randomized recoil sets the lowest temperature, T_D, which can be achieved via Doppler cooling given by [35]

$$k_B T_D = \frac{1}{2}\hbar\Gamma \tag{2.17}$$

where k_B is the Boltzmann constant. In a harmonic trap with secular frequency ω, this corresponds to a mean phonon number

$$\bar{n}_D \;\approx\; \frac{k_B T_D}{\hbar\omega} - \frac{1}{2} = \frac{\Gamma}{4\omega} - \frac{1}{2}\,, \tag{2.18}$$

which is generally $\bar{n}_D \gg 1$ for typical $\Gamma \gg \omega$.

When discussing cooling with trapped ions it is useful to quantify how strongly photon recoil couples to motion. Combining our earlier discussion of the Lamb-Dicke parameter with the Doppler cooling limit, we can estimate whether our system satisfies the conditions to operate in the Lamb-Dicke regime. For a $^{174}Yb^{+}$ ion, assuming a radial trap frequency of ≈ 600 kHz, and with a 369 nm transition of linewidth $\Gamma \approx 123$ MHz, this yields:

$$\eta\sqrt{2\bar{n}+1} = \frac{2\pi}{\lambda}\sqrt{\frac{\hbar}{2m\omega_{sec}}}\sqrt{2\left(\frac{\Gamma}{4\omega_{sec}} - \frac{1}{2}\right)+1} \approx 0.687. \tag{2.19}$$

---

maximum $\Gamma = 1/\tau$ (the natural linewidth). If photons did not have a temporal waveform, we could not observe interference. These concepts will be treated in greater depth in subsequent chapters.

The corresponding temperature and mean phonon number are about 0.47mK and 16 respectively. Axial frequencies in our trap are about one third of the radial frequency. In other words, for typical operation, we sit close to our slightly above the Lamb-Dicke regime. For the purposes of the photon control experiments, we limit ourselves to Doppler cooling. However, to perform actual circuits with an ion chain, systems must be able to prepare a crystal with a very low (ideally no) additional motional quanta number. Some prominent sub-Doppler techniques include resolved sideband cooling [36, 37], polarization gradient cooling [38] and electromagnetically induced transparency cooling [39, 40]. Additionally, we could push trap frequencies higher by increasing voltages on axial and radial electrodes to suppress changes in motional state during photon generation experiments.

2.4 Building an ion trap

In the prior sections of this chapter, we have established how to trap an ion, *in theory*. Now we will outline the construction of an actual ion trap suitable for the photon control experiments we aim to carry out.

2.4.1 Trap design

It turns out that many electrode geometries, not just the perfectly hyperbolic surfaces used in the first ion traps, can form good trap potentials, albeit at the cost of a field that is no longer purely quadrupolar. Here we describe some of the practical considerations in designing a trap and provide a light overview of how trap fields can actually be modeled.

When developing a new trap or selecting an existing design for a given quantum information application, some factors to take into account include:

- **Trap frequencies:** Typical quantum information applications tend to target traps with radial frequencies on the order of 1 MHz and axial frequencies of several hundred kHz although exact numbers are problem specific.

- **Trap depth:** Pseudopotential depth should comfortably exceed the thermal energy of

ions during loading and the potential generated by stray fields. Chip traps often have depths on the order of 0.1 eV while macro traps can achieve 10s or 100s of eV.

- **Stray field compensation:** The presence of background fields pushes ions off the RF null, inducing excess micromotion and reducing effective trap depth. Inclusion of additional electrodes to cancel the effects of stray fields can prove valuable.

- **Trap displacement from electrode surfaces:** Studies of anomalous heating have shown surfaced induced chain heating scales as $1/d^4$ where d is the distance from the ion to electrode [41]. As such, traps which position ions farther from the surface tend to be preferred, all else held equal.

- **Material selection:** Non-magnetic materials with low out assing rates and good electrical properties are ideal for trap electrodes.

Ensuring that a particular electrode geometry will actually form a suitable trapping potential is often as simple as a qualitative drawing of the electric field lines. If field lines drawn from (+) polarity RF electrodes can be expected to terminate on ground DC electrodes such that a test charge would be squeezed in one direction while being pushed in the other, such a geometry should form a trap. This can be seen by flipping the polarity of the RF electrodes and imagining the time averaged potential formed. However, it is advisable to actually model expected field behavior for the construction of a new device. Below we sketch the high level, conceptual approach to numerically predicting trap environments.

In the absence of free charge in the trapping region, the electrostatic potential satisfies Laplace's equation $\nabla^2\Phi(\vec{r}) = 0$ with boundary conditions set by the electrode voltages. In practice most geometries are too complicated for analytics, so we solve Laplace's equation with a finite-element method. Practically, this is done with a physics modeling tool like COMSOL where a given trap geometry is discretized into small cells. Potentials are applied to electrode surfaces and electric fields for each volume element are solved through iterative relaxations corresponding to averaging neighboring cells while enforcing the boundary conditions set by

electrode potentials. Eventually, iterative solutions begin to converge and we can gain a good numerical approximation of the electric field magnitude for the trap.

An approximation of the trap restoring force in directions constrained by the RF potential can be estimated by averaging out effects of driven motion at the RF frequency (micromotion). This allows us to write the trap psuedopotential at position $\vec{r}$ as [42]

$$\Psi(\vec{r}) = \frac{e^2}{4m\Omega_{\mathrm{RF}}^2} |\tilde{\vec{E}}(\vec{r})|^2 \tag{2.20}$$

where e, m, Ω_{RF} and $|\tilde{\vec{E}}(\vec{r})|$ are the charge of an electron, ion mass, trap angular frequency and magnitude of the RF electric field respectively. The total effective trapping potential is $U_{\mathrm{eff}}(\vec{r}) = e\,\Phi_{\mathrm{DC}}(\vec{r}) + \Psi(\vec{r})$ where Φ_{DC} is the potential generated from the DC electrodes.

2.4.2 The vacuum system

Reliable ion trapping requires an ultra-high vacuum (UHV) environment to suppress background-gas collisions that heat or eject the ion.

Our chamber is evacuated and kept at UHV by a combination of a titanium-sublimation (TiSub) pump and an ion pump, while an ion gauge monitors pressures. The ion pump (typically tens of L/s) provides continuous pumping by ionizing residual gas which are then accelerated by the electric field generated by the pump anode-cathode and become adsorbed into the pump surface. However, not all molecules in the background gas are efficiently ionized by the pump. The TiSub helps here by flashing a Ti filament to deposit a fresh, getter film on nearby surfaces. The ion gauge measures pressure down to the $\sim 10^{-11}$ Torr range. Proper bake out results in baseline pressures of $\sim 1 \times 10^{-10}$ Torr.

All trap components and chamber elements are assembled using UHV-compatible materials (tungsten rods for trap electrodes, stainless steel for chamber itself, etc.). The entire vacuum system is baked to desorb water and hydrocarbons which would otherwise prevent us from reaching UHV. We use a three stage pumping system consisting of a diaphragm/rough pump ($\sim 10^{-1}$ torr) followed by a turbomolecular pump to achieve high vacuum ($\sim 1x10^{-6}$ torr) before turning on the ion pump. We ramp the oven temperature up to $\sim$ 500K over the

course of several days and monitor pressure using an ion pump and a residual gas analyzer (RGA). Cool down is begun when baseline pressure begins to plateau and the composition of gas read on the RGA shows little contribution of atmospheric components and H_2, expected from stainless steel out gassing. After cool down, we run the TiSub repeatedly until pressures in at least the low 10^{-9} Torr regime are reached. If we are not able to pump down to these pressures, leak checks with acetone and He, monitored on the RGA, can help identify under (or over) tightened flanges or other faulty components.

2.4.3 Supplying trap voltages

To generate the several hundred volt at 10s of MHz required for a RF trap while driving from a 50 Ω source, we couple the RF through a shielded helical resonator. The resonator behaves as a high-Q quarter-wave resonator with effective inductance, L and capacitance, C, and resistance, R, determined by the helix geometry and load (i.e. the trap, feedthroughs and wires). Its resonance is

$$\omega_0 = \frac{1}{\sqrt{L\,C}}, \qquad f_0 = \frac{\omega_0}{2\pi}, \tag{2.21}$$

with quality factor $Q \simeq \omega_0 L/R$ [43]. The quality factor of the resonator can easily be measured via the use of a bidirectional coupler or standing wave ratio meter. The loaded resonator used in our system has a Q of ∼200 and a frequency of 10.845 MHz.

An antenna mounted inside the coil shield broadcasts the RF frequency which is picked up by the resonator. Assuming perfect coupling between the antenna and resonator, we can estimate the peak voltage at the trap as

$$V_{\mathrm{pk}} \;\approx\; \sqrt{\frac{2\,Q\,P_{\mathrm{in}}}{\omega_0\,C_{\mathrm{load}}}}. \tag{2.22}$$

Equation 2.22 makes clear the tradeoffs: higher Q, more input power, and smaller C_{load} increase the trap RF amplitude, while higher drive frequency reduces it for fixed P_{in}. Practically, we tune f_0 by adjusting the number of loops of the helical coil. Coupling is optimized by adjusting the position and shape of the antenna. Finally, the quality factor

is measured from the 3 dB bandwidth ($Q = f_0/\Delta f$) by inspecting the ratio of forward to reflected power with the standing wave ratio meter.

Additional voltages are applied for the end cap electrodes and the two DC rods. To reduce inductive coupling from the RF rods, π filters are placed right before the vacuum feedthroughs for these electrodes and capacitors are used as an RF shunt between the power supply signal and ground outputs. One end of the helical resonator is typically grounded so that the DC bias of the RF electrodes is zero. However, an additional DC offset can be added by instead connecting the resonator to a power supply with the use of another π filter.

2.4.4 Atomic ovens

The final element of the vacuum chamber infrastructure is the atomic source used to generate ions. The first step in produce ions is creating a flux of neutral atoms. This is accomplished through the use of a resistive oven composed of a stainless steel tube, crimped at one end, and loaded with flakes of ytterbium. The tube is then spot welded at the bottom and top to two vacuum feedthrough leads. Applying a current of a few amps to the oven produces sufficient neutral atom flux to load ions on the time scale of minutes if other system parameters are set correctly. It is important to prevent excess atomic deposition from coating various view ports and trap surfaces as this occludes optical access and can lead to the formation of troublesome patch potentials [44]. Therefore, a shield with an aperture is mounted inside the chamber with a $\sim$ 800 micron hole drilled through to roughly collimate the atom flux. The oven used in our chamber, and its alignment relative to other trap elements is shown in Fig. 2.13.

Alignment of the oven is performed by sending a collimated laser through the view port opposite the oven and ensuring the beam can pass through the trap center, aperture and top of the oven simultaneously. The oven orientation can be verified by observing the deposition of Yb on the opposite viewports. Representations of the alignment process and atomic deposition are shown in Fig. 2.14.

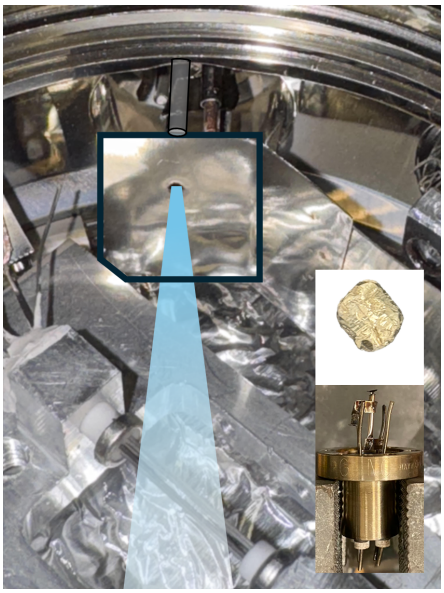

Figure 2.13: Resistively heated atomic oven used to generate a neutral ^{174}Yb beam for ion loading. **Main image:** view inside the vacuum chamber; the blue cone indicates the intended atom-flux path from the oven to the trap center. The aperture is outlined in a thick black line and a representation of the oven is added (towards the top center) for clarity. **Inset (top right):** atomic Ytterbium. Chips are cut off and loaded into the oven. **Inset (bottom right):** stainless-steel tube oven spot-welded to feedthroughs leads. A fixed shield with an $\sim$800 μm aperture (magnified box) collimates the flux and limits deposition on viewports and electrodes. Typical operation uses a few amperes of current, yielding sufficient flux to load ions within minutes.

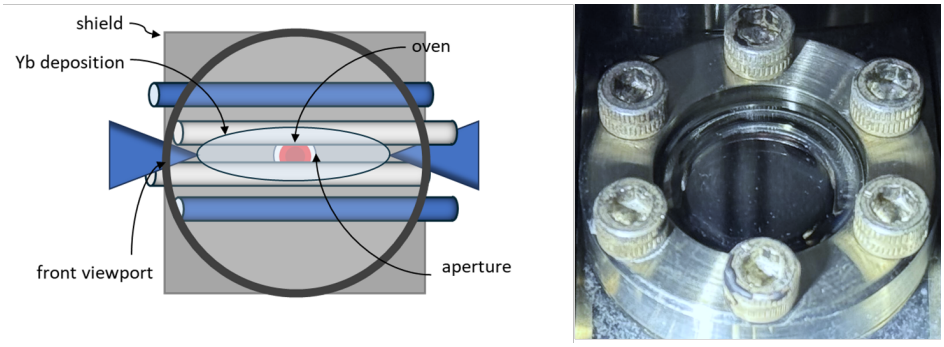

Figure 2.14: Schematic representation of oven alignment relative to the trap and shield aperture. **Left:** A resistively heated Yb oven (bright red circle) emits a neutral-atom plume (light blue oval) that is collimated by a fixed shield with an $\sim$800 μm aperture. The plume is centered on the trap axis (between the four electrodes), minimizing Yb deposition on electrode surfaces. Alignment is performed by sending a focused beam through the front viewport (black circle) so that it passes sequentially through the trap center, the aperture, and the oven opening. **Right:** Alternatively, alignment can be verified by observing deposition formed on the view port opposite the the oven. The silhouette of rods is seen in the areas where there is minimal deposition (silvery gray). This level of deposition indicates a unnecessarily large collimation angle. Ideal collimation would lead to minimal atomic flux outside of the trap center to prevent deleterious deposition

2.4.5 Lasers

In order to ionize, Doppler cool and perform photon control experiments with $^{174}Yb^{+}$, three lasers are needed. First, we must ionize the neutral atom flux generated by our ovens. Figure 2.15 shows a simplified level structure for both neutral and singly ionized Yb. To achieve efficient ionization, we resonantly excited on the strong $^1S_0 \rightarrow {}^1P_1$ transition near 399 nm (natural linewidth $\Gamma_{399}/2\pi \approx 29$ MHz). A second photon provided by the 369 nm laser used for Doppler cooling promotes the atom to the continuum. Tens of microwatts of 399 nm light tuned within ∼50 MHz of the nominal transition frequency and focused to a ∼75 micron waist at the trap center proves sufficient to load ions in minutes, given sufficient 369 nm power and oven current. A grating stabilized external-cavity diode laser (ECDL) provides sufficiently narrow linewidth source with mode-hop-free tuning over several hundred MHz via a combination of diode current, temperature and piezo-controlled grating angle. By observing neutral atom florescence (NAF), we scan across the resonance and optimize ionization rate. For day-to-day operation we reference the 399 nm frequency to a high-accuracy wavelength meter. In our trap geometry, the neutral atom flux and 399 nm k-vector are nearly orthogonal limiting the effects of first order Doppler shift to the collimation angle of the oven-aperture set up. When run at low currents, the range of Doppler shifts generated by the thermal distribution of neutral atoms leads to a distribution much narrower in frequency than the isotope shift between 174 and other isotopes present in the oven. This allows good isotope selectivity when loading ions. By scanning the 399 nm frequency across the atomic transition, we can also get a rough measurement of the transition linewidth. This is accomplished by looking for neutral atom florescence and recording the rate of photon scatter over a range of several hundred MHz. The results are shown in Fig. 2.16. NAF can also be used to infer the position of the laser beam relative to the trap for proper alignment as shown in Fig. 2.17

Doppler cooling, state detection, and fluorescence imaging use the $^2S_{1/2} \leftrightarrow {}^2P_{1/2}$ cycling transition of $^{174}Yb^{+}$ at 369.5 nm with $\Gamma/2\pi \approx 19.6$ MHz (natural lifetime ∼8.1 ns). To realize efficient cooling and bright fluorescence, the cooling laser's instantaneous linewidth should

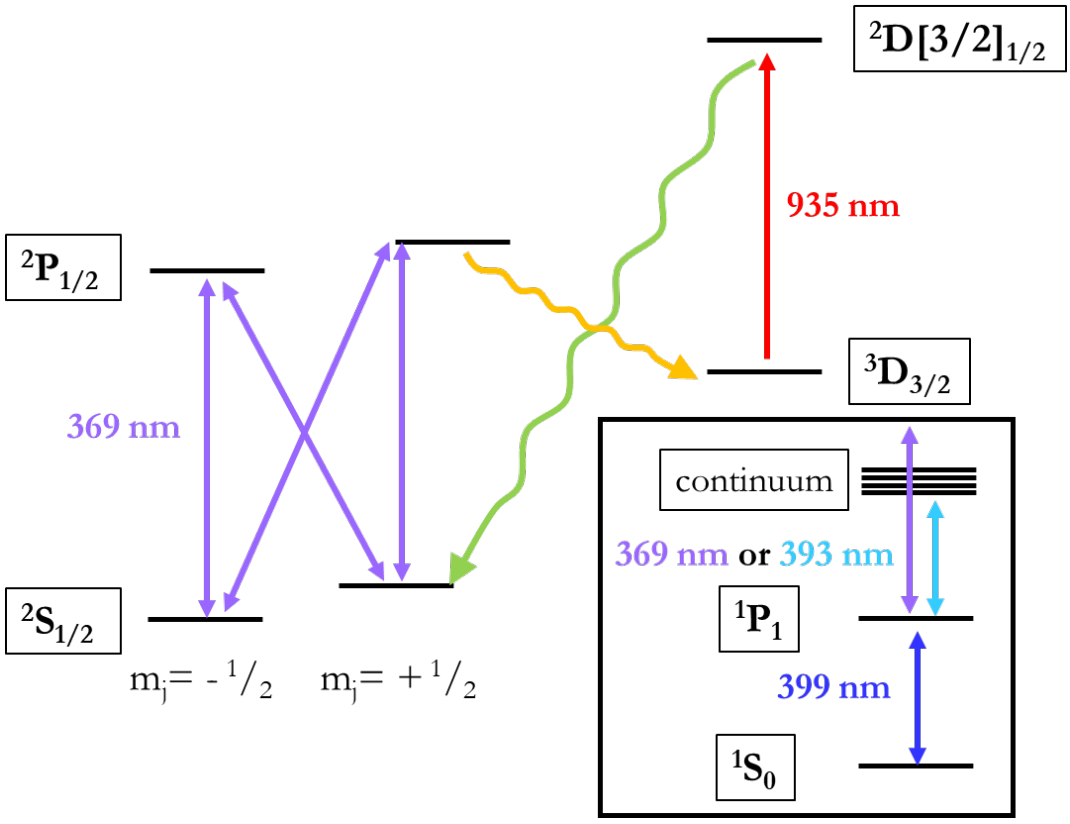

Figure 2.15: Relevant level structure for $^{174}Yb^+$. The Doppler-cooling cycling transition is $^2S_{1/2} \leftrightarrow {}^2P_{1/2}$ at 369 nm (purple). Population that leaks to $^2D_{3/2}$ is repumped via $^2D_{3/2} \leftrightarrow {}^2D[3/2]_{1/2}$ at 935 nm (red), which predominantly decays back to $^2S_{1/2}$ (green). The inset shows the neutral-ytterbium photoionization scheme: $^1S_0 \rightarrow {}^1P_1$ at 399 nm (blue), followed by a second photon near 369 or 393 nm to the continuum.

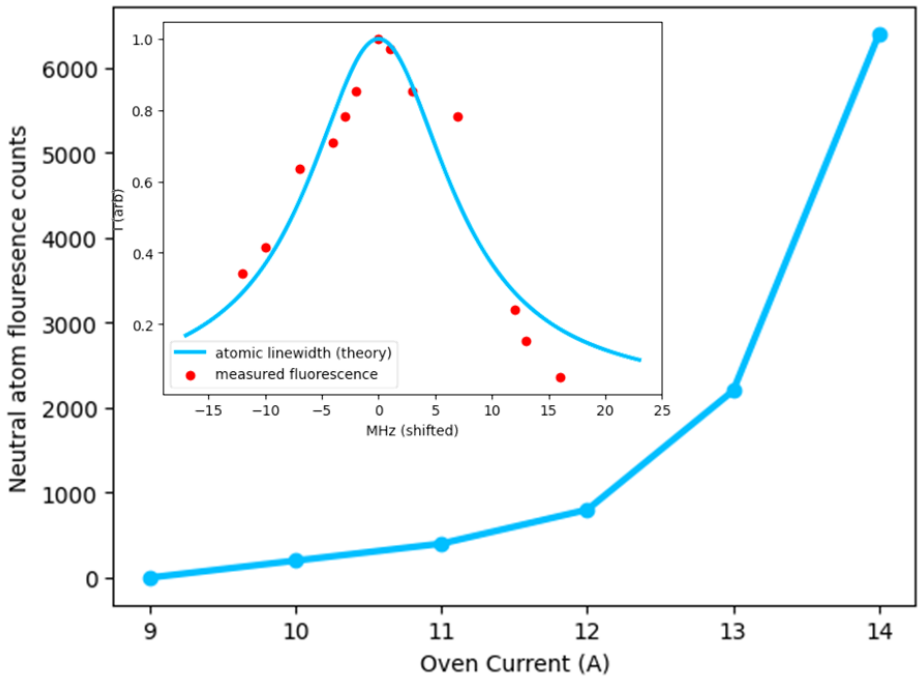

Figure 2.16: Indirect measurements of relative neutral Yb flux at various oven currents derived by counting scattered 399 nm photons as oven power is increased. **Inset:** Approximate measurement of $^1S_0 \Rightarrow ^1 P_0$ transition linewidth for ^{174}Yb measured by scanning 399 nm frequency. Although data is noisy, we see no discernible 1st order Doppler broadening as flux direction and k-vector are nearly orthogonal, limited by collimation angle of oven.

be well below Γ and the frequency should be stable at the few-MHz level over experimental timescales. A UV ECDL satisfies these requirements, typically achieving <100 kHz.

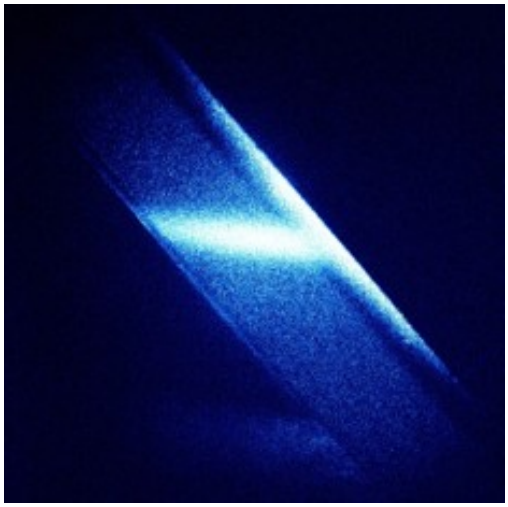

Figure 2.17: Resonance fluorescence from neutral Yb flux illuminated by the 399 nm $^1S_0 \rightarrow {}^1P_1$ beam. The bright stripe marks the laser–atom overlap, while the sharp dark boundaries are trap electrodes. Because the atomic plume is wider than the laser, NAF can be used to measure the 399 nm beam position and waist in situ relative to the electrode edges which proves useful for aligning photoionization and subsequent cooling beams.

Magnetic field coils are used to set the direction of the quantization axis and break the degeneracy of the Zeeman sublevel. Orienting the magnetic field perpendicular to the propagation of lasers ensures that a single, linearly polarized beam of 369 nm will cool the transition. For photon shaping experiments, we will replace a single linear polarization with separate, balanced left and right circularly polarized beams. Population that leaks from $^2P_{1/2}$ into the metastable $^2D_{3/2}$ manifold (with a branching ratio of $\sim$ 1:200 relative to the almost cycling transition at 369 nm) is returned to the cooling cycle via the $^2D_{3/2} \rightarrow {}^3[3/2]_{1/2}$ transition at 935.2 nm, which rapidly decays to $^2S_{1/2}$. A third ECDL at 935 nm readily meets this specification with modest optical power (tens to hundreds of μW at the ion). Figure 2.18 and Tab. 2.2 show and list all of the 'laser-side' elements in our optical system.

Some additional optics are located after the output of the optical fibers on the 'trap-side' and will be addressed in a later section.

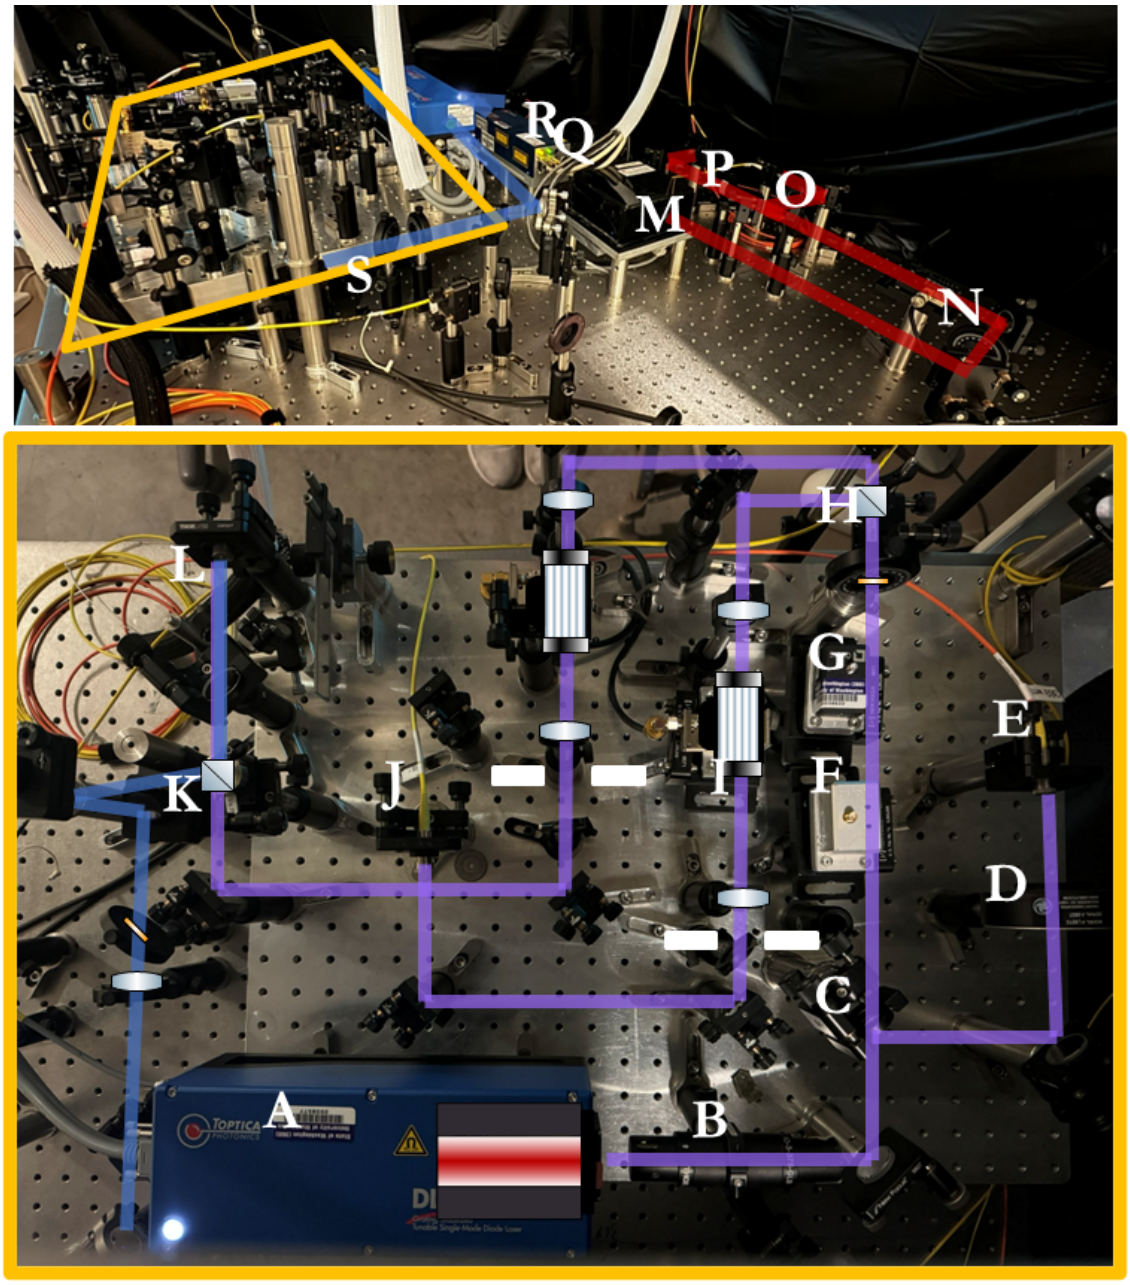

Figure 2.18: The three ECDLs used for our experiment with 'laser' side optics necessary for beam shuttering, amplitude modulation and polarization control. 369 nm path shown in purple, 399 nm show in blue and 935 nm shown in red. Labeled elements are recorded in Tab. 2.2.

Element	Name	Function
A	369 ECDL	Provides cooling, state detection and photon generation light
B	Optical isolator	Prevents etalon effects destabilizing laser
C	10% pickoff	Sends 369 to WM
D	Shutter	Prevents fiber switch from burning when WM not used
E	WM multimode fiber/coupler	Monitors the 369 nm frequency
F	12.6 GHz EOM	Bridges S hyperfine splitting
G	2.1 GHz EOM	Bridges P hyperfine splitting
H	HWP & 355 laser line PBS	Splits σ^+ and σ^- paths
I	Pulse synthesis apparatus	Includes AOM, lenses and iris for cooling and excitation
J	σ^+ path single mode fiber coupler	Acts as a spatial mode filter and sends light to trap
K	369 & 399 merge PBS	Combines ionization light with σ^+ path
L	σ^- single mode fiber coupler	Acts as a spatial mode filter and sends light to trap
M	935 ECDL	Provides repumping light
N	Repump AOM	Shutters repump light to monitor ion fluorescence
O	Repump fiber coupler	Sends 935 light to trap
P	935 WM path	Monitors 935 nm frequency
Q	399 ECDL	Provides light for first stage photo-ionization
R	399 WM path	Monitors 399 nm frequency
S	Periscope	Sends 399 light to σ^+ path

Table 2.2: 'Laser side' optics and their applications for trapping and photon control experiments in $^{174}Yb^+$.

2.4.6 The imaging system

Fluorescence from a single ion is collected by a 20× long–working-distance objective (numerical aperture NA = 0.40). Although designed for infinite conjugal use with a lens tube, we instead displace the objective from the trap center at a distance slightly longer than its nominal working distance. Thus, rather than forming an image of the ion at $z = \infty$, we use the objective as a slow lens. The ion at distance s_1 from the objective's principal plane is imaged to an intermediate plane at distance s_1' given by the thin–lens relation with magnification of this first stage, m_1,

$$\frac{1}{f_{\text{obj}}} = \frac{1}{s_1} + \frac{1}{s_1'}, \qquad m_1 \equiv -\frac{s_1'}{s_1}. \tag{2.23}$$

An adjustable circular aperture is placed at (or very near) the intermediate image plane. By locating an aperture at s_1', we can form a tight spatial mask, reducing the effective field of view at the image plane to a dimension close to the ion spot or crystal size. In this set up, the objective front housing still acts as the entrance pupil, setting the solid angle light cone collected from the ion, while the aperture diameter controls the exit pupil. Therefore, if the aperture is placed appropriately close to the intermediate image plane (typically the tolerance here is set by the depth of focus of the objective), we can minimize background scatter without decreasing ion photons collected by the imaging system.

A second lens of focal length f_2 is positioned a distance u_2 from the aperture. We deliberately set $u_2 \gtrsim f_2$ so that this final stage lens forms an image of the intermediate plane, and thus the ion, at a second image plane where we locate the sensor at a distance u_2'. A flip mirror allows us to send light either to an electric charge coupled device camera or a photomultiplier tube (PMT).

$$\frac{1}{f_2} = \frac{1}{u_2} + \frac{1}{u_2'}, \qquad m_2 \equiv -\frac{u_2'}{u_2}. \tag{2.24}$$

The overall magnification at the camera is then the product

$$M = m_1\, m_2 = \Big(-\frac{s_1'}{s_1}\Big)\Big(-\frac{u_2'}{u_2}\Big), \tag{2.25}$$

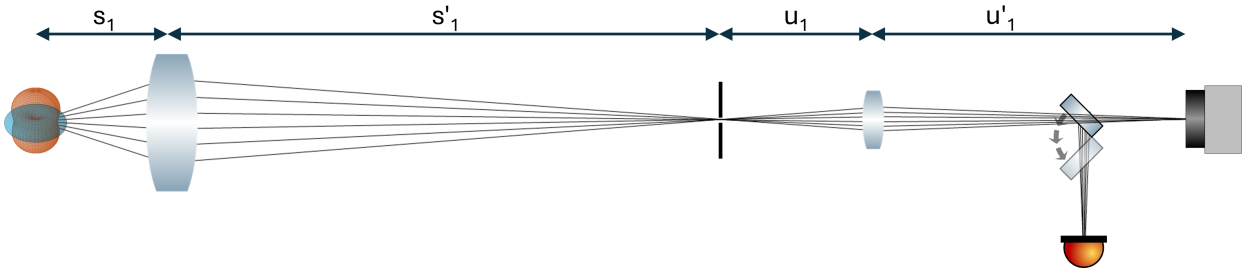

Figure 2.19: Schematic representation of the light collection system used to image ions and count photons.

with the sign indicating image parity. A schematic for our imaging system is shown in Fig. 2.19

Based on our optics, we can estimate the approximate rate of photons we can collect under ideal conditions (i.e., a strongly saturated transition with minimal instantaneous Doppler shift due to ion motion). The use of our 0.40 NA objective corresponds to half angle of $\theta = \arcsin(0.4) \approx 23.6$ degrees. We then calculate the corresponding solid angle, Ω with

$$\Omega = 2\pi(1 - \cos(\Theta)) \tag{2.26}$$

yielding about 0.52 steradians or $\sim 4\%$ for an isotropic emitter. Factoring in losses through the viewport, objective and additional optics, we estimate the total throughput for the system to be about $\sim 2\%$. The quantum efficiency for the PMT used in our experiments is ~ 0.3 at 369nm. This yields a total count rate of at most $0.3 \times 0.02 \times 1/(2\tau) \approx 3.7 \times 10^5$photons/second. Experimentally, we observe significantly fewer counts which we attribute largely to effects of excess micromotion, detailed more in Ch 5.

2.4.7 Putting it all together

Once the trap is assembled and on the optical table, electronics are in place, and laser and imaging optics are aligned, trapping can finally begin. An image of the full 'trap side' system

can be found in Fig. 2.20 with accompanying labeled components described in Tab. 2.3.

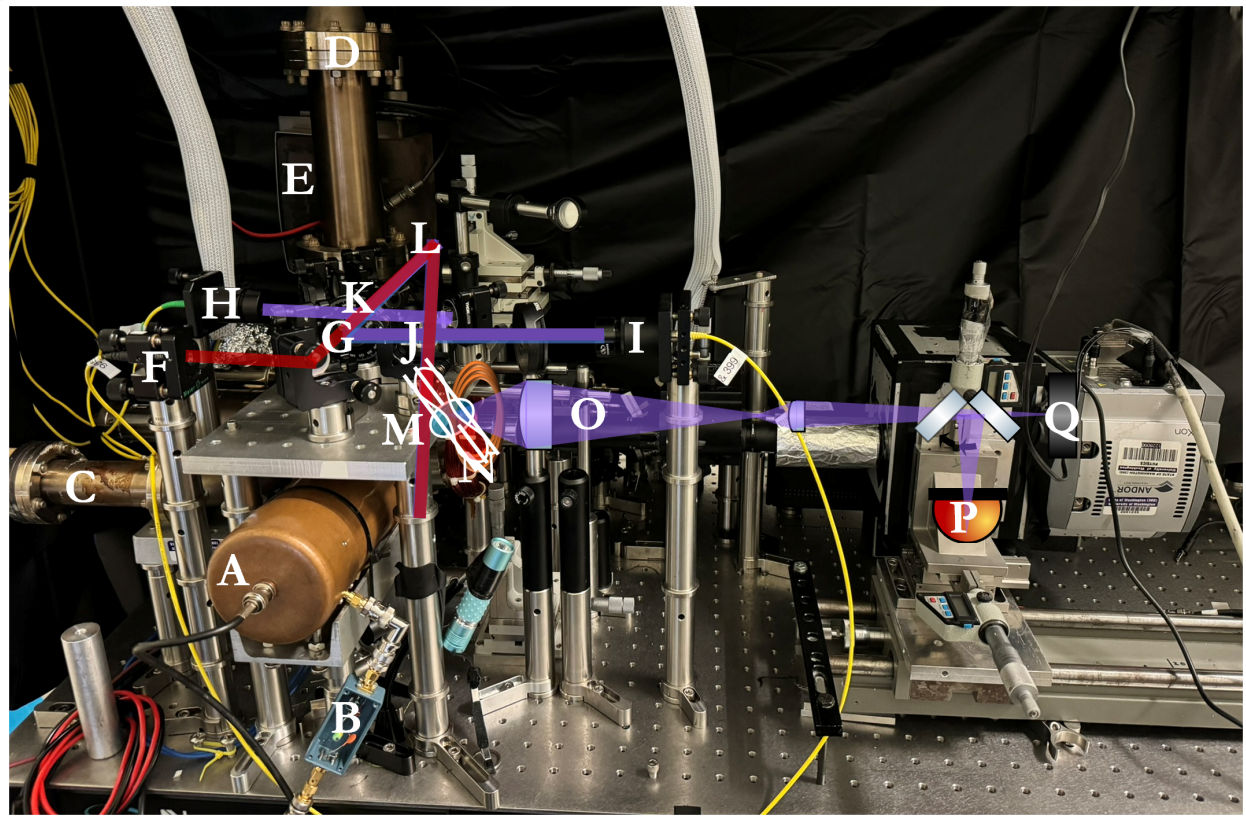

Figure 2.20: Full 'trap side' assembly for photon control experiments showing vacuum chamber, RF electronics, laser and imaging optics. Labeled components reflect entries in Tab. 2.3.

In the hopes of assisting future ion trappers in troubleshooting their systems, the rest of this section is dedicated to a brief overview of how initial trapping was established in the 4-rod trap used for our experiments.

First, the imaging system is aligned by adjusting the displacement of the microscope objective from the trap center until a clear silhouette of trap features (usually the conical end cap electrodes) can be seen either by eye or on the camera used for trapping. Then, the aperture and relay lens are attached to achieve the desired magnification. Minor adjustment

Element	Name	Function/Notes
A	Helical resonator	Supplies RF to trap
B	π circuit for RF bias	4 additional bias electrodes, not shown
C	Ion gauge	Monitors background gas pressure
D	Titanium sublimation pump	Getter pump for low ionizability gas
E	Ion pump	Maintains UHV
F	935 nm input	Used for repumping
G	Dichroic mirror	Combines 935 with 369, 399
H	369 (σ^+) input	Used for Doppler cooling and optical pumping
I	399, 369 (σ^-) input	Used for ionization, Doppler cooling and pulse excitation
J	Polarizing beam splitter	Combines σ^+ and σ^- paths
K	Quarter wave plate	Converts σ^+, σ^- beams to left, right circular polarizations
L	$f = 200$ mm lens in 3 axis stage	Reduces beam size, provides beam steering
M	4-rod trap	Generates potential to confine ions
N	Magnetic field coil	Sets quantization axis with 2 other coils, not shown
O	Ion imaging optics	Collects photon scattered by ion
P	photomultiplier tube	Counts photons
Q	EMCCD camera	Images ions

Table 2.3: Optical and trapping elements with their applications in trapping and photon control experiments with $^{174}Yb^+$.

of the objective position along the optical axis is required to form an image at the intermediate plane set by the aperture position. The final stage imaging lens must be positioned farther from the aperture than its nominal focal length in order to reimage the trapping region at the sensor plane.

The center of the trap is then located by identifying several trap features. In the case of the 4-rod trap used, the average position between the two end caps is used for initial alignment. As end caps may be slightly misaligned relative to the 4-rods which determine the quadrupole potential, it can be useful to identify the axial location of the trap center using end caps and then to calibrate the correct x and y coordinates through measurement of all 4-rod positions by moving the objective to center the bottom/top of rods in the camera field of view. The correct object plane can be found by measuring the position of the objective when the front and back rods are in focus and taking the average. Because the 0.40 NA objective sits within a recessed view port, it can not be translated sufficiently to image trap features. As such, for initial trapping, a 0.21 NA objective was used whose longer working distance puts it above the recess, allowing all trap features to be found.

Once we are confident the imaging system coordinates are established, we can then align the three lasers to the trap center. This is most easily done by first combining the IR and UV with a cold mirror and then walking the beams and checking that their spatial modes are well overlapped at a point directly after the cold mirror and a second point $\sim$5m away. The beams are then centered on a lens mounted on a 3-axis stage which is used to control beam pointing in the trap. This lens is adjusted until diffraction of the lasers on an identifiable trap feature (once again, usually one of the end caps) is observed in the beam exiting the chamber. The lens is then moved along the optical axis until the amplitude of diffraction fringes is maximized, indicating the beam waist is located at the trap feature.

Returning to the imaging system, laser scatter should now be clearly visible on the target feature. This process of adjustment may lead to the lasers no longer being well centered on the lens. Due to chromatic dispersion between the 369 nm and 399 nm and especially 935 nm beam, this can lead to poor overlap of the three lasers at the trap location. Therefore, small

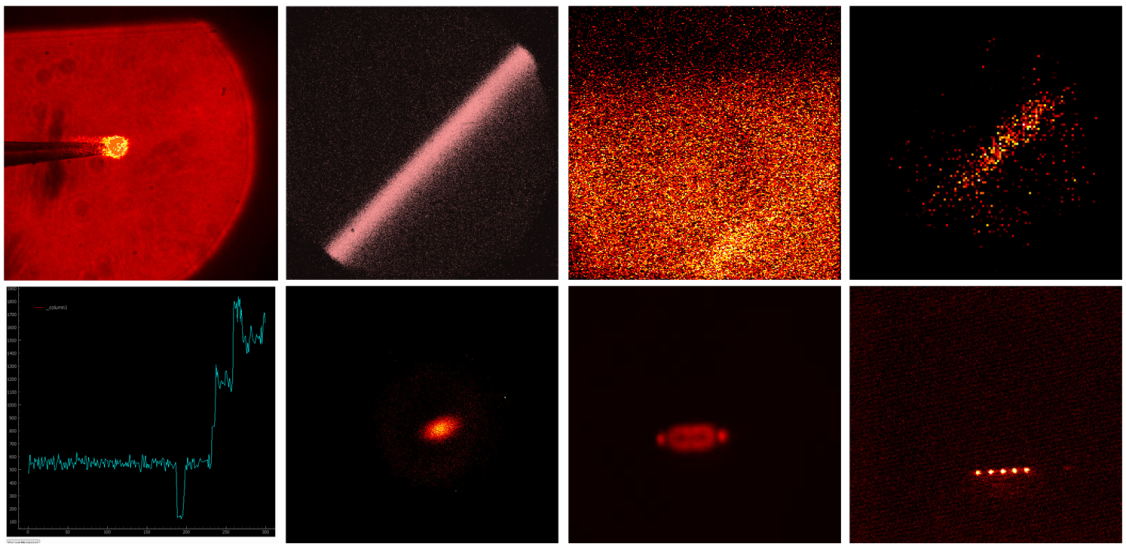

Figure 2.21: Images of the trapping process. **Clockwise from top left:**(i) One endcap is imaged at the sensor plane. The aperture is also in focus. A large field of view is initially used by selection of a relatively slow final stage imaging lens. This allows the top (pictured) or bottom rod to be viewed in focus simultaneously, useful for measuring an offsets of end caps from the center of the 4 rods. Laser scatter is seen illuminating the tip of the endcap. Checking all three wavelengths overlap at both end caps guarantees good beam alignment at the trap center. (ii) Objective is set to image the middle of the trap, oven is turned on so that laser path through trap can be inferred and lasers are adjusted until 399 nm beam is centered with respect to imaging system. (iii) A bandpass filter is then used to only collect 369 nm light and an ion cloud is visible. (iv) Trapping parameters (beam pointing, laser frequencies and trap voltages) are adjusted to maximize the brightness of the cloud to find the trap center. (v) The cloud is dumped and the oven is turned back on. Light is sent to the PMT allowing the number of ions trapped to be controlled. This is done by monitoring discrete jumps in photon counts indicate loading of one ion after another with ionization light shuttered when the desired number of ions are loaded. (vi) Ions are reimaged on the camera, but are still too hot to crystallize. (vii) Trapping parameters are once again adjusted until ions are sufficiently cool to form crystals. (viii) Reducing axial confinement leads to the formation of linear ion chains.

adjustments to the beam paths can be made until scatter from all three lasers is maximized at the same point of the trap. The beams should then be moved until the second end cap is similarly found. Once again, as beam pointing between the two trap locations of the three lasers may be slightly different, individual laser alignment may require iterative adjustments to ensure proper overlap at two trap features. Once this is done, the beams can be centered at the expected trap center and the attempt to trap can begin.

The following process to identify the trap center proved repeatable and deterministic and allowed realignment to the trap and successful retrapping. First, the oven was turned up to several amps, bounded by ensuring the ion gauge read less than 1×10^{-9} Torr. A narrow line 399 nm filter was added to the imaging system to allow the identification of neutral atom florescence (NAF). If lasers were well centered as described above, and the lasers were well tuned, some photon counts above the background were always seen. Using the final lens, the lasers were then moved until the NAF beam was centered on the camera at the expected trap null.

Once this was done, the 399 nm filter was replaced with a 369 nm filter to look for ion florescence. As ions may be initially hot and the trap center may be still somewhat displaced from the expected location based on optical alignment, I found the PMT to be a much more useful tool for initial trapping than the camera. The oven was run relatively hot to allow a large cloud of ions to fill the trapping region. At this stage, a small increase in florescence over the background was always observed. Ion florescence was distinguished from background scatter by turning off the 935 nm light which leads to the ions being pumped into the dark state.

After ion florescence was definitively identified, 935 nm and 369 nm laser frequencies were scanned to maximize initial cloud brightness. The true trap center was then found by moving lasers beams and imaging system in tandem until the brightest ion florescence was achieved. At this stage, the ion cloud should be discernible on the camera. Using a smaller number of bins, longer integration time, high gain and potentially subtracting background can be useful for finding the cloud with the camera. Because this cloud is so large (typically hundreds of

microns in each direction), it is possible to actually see the path of the beams through the cloud much like the image of NAF.

Species	**Transition**	f **(THz)**	**Note**
Neutral ^{174}Yb	$^1S_0 \rightarrow {}^1P_1$	751.52650	Resonant first step for photoionization
^{174}Yb$^+$	$^2S_{1/2} \rightarrow {}^2P_{1/2}$	811.29140	Doppler cooling and fluorescence detection
^{174}Yb$^+$	$^2D_{3/2} \rightarrow {}^3[3/2]_{1/2}$	320.57203	Repump from the metastable $^2D_{3/2}$ manifold
^{6}Li (reference)	$^2S_{1/2} \rightarrow {}^2P_{1/2}$ (D_1)	446.79964	Frequency reference (Li D_1 line) used for calibration

Table 2.4: Measured transition frequencies for ^{174}Yb$^+$ and a reference transition used for calibration of the wavemeter.

The next step is to go from loading a cloud to loading a small crystal, and ultimately single ions. To do this, the trap is dumped by turning off and on the RF. The oven current is then turned down ($\sim$3.4 A is sufficient for our device). The reduces the loading rate and allows us to count as ions are loaded into the trap by monitoring the PMT signal which shows a discrete increase as each additional ion is loaded. When a few ions have been loaded, the oven is turned off and we switch back to use of the camera. By iteratively adjusting laser frequencies and trap voltages, we are able to cool ions down until individual atoms are resolvable. At this stage, we switch to the use of the 0.40 NA objective. If the imaging system is well aligned, switching objectives should not change the object plane and the same region of the trap should be visible (albeit with a smaller field of view). However, because the

0.40 NA objective sits closer to the trap, this does perturb the ion equilibrium position and as such it may be necessary to iteratively adjust beam pointing and imaging alignment until ion florescence is once again maximized. We detected at most $\sim 5 \times 10^4$ photons per second per ion, about half an order of magnitude fewer than our back-of-the-envelope calculation in the prior section would indicate. We explore the source of this discrepancy in more depth in a Ch 5. Images of the ion trapping process are shown in Fig. 2.21. Measured transition frequencies, along with a calibration reference are displayed in Tab. 2.4.

Now that we have described the core theory and implementation behind ion trapping, we are ready to move on to doing something novel and interesting with our device.

Chapter 3

THE INTERACTION OF SINGLE ATOMS WITH FIELDS OF LIGHT

The role of this short bridge chapter is to return to the basics of atomic physics in somewhat more depth in order to establish the conceptual framework needed to understand the properties of single photons emitted from atoms. We will provide a treatment, with an experimentalist's eye, of the relevant topics in quantum optics which will serve as the basis for the tools we develop for photon control.

Specifically, we need to understand how to model the interaction between optical pulses with a resonant atomic transition of a given lifetime. Additionally, we must be able to describe the photon produced when the excited state relaxes, including its polarization, emission statistics, and temporal behavior.

We begin by outlining the pertinent questions we must be able to answer to arrive at a full treatment of this problem:

1. What determines the characteristics of an atomic excited state such as its lifetime and branching ratio to lower lying levels?

2. How can we model the interaction of an atom with an electromagnetic field which drives a transition between states?

3. When an excited state relaxes, how do we characterize the mode of the electromagnetic field during this process?

4. How can we properly account for the emission statistics of photons during an atomic excitation of finite duration?

We will proceed by answering these questions one by one, after which we will have a clear picture of the tools at our disposal for the control of single photons.

3.1 Understanding atomic transitions

We first need to understand how different atomic energy levels couple to each other. There are two main concepts we need to unpack to get there: the reduced dipole matrix element and Clebsch-Gordan coefficients.

We will start by considering a simple 2-level systen where a ground state $|g\rangle$ is coupled to an excited state $|e\rangle$ separated by $\hbar\omega_0$, driven by a classical optical field $\vec{\mathcal{E}}(t) = \vec{E(t)}\cos(\omega t)$. To quantify the strength of coupling, we define the electric-dipole matrix element,

$$\vec{d} = \langle e|\,\hat{d}\,|g\rangle \tag{3.1}$$

where $\hat{d} = e\hat{r}$ is the dipole operator, e is the charge of the electron and $\hat{r}$ is the position operator.

We use the electric–dipole matrix element because, at optical wavelengths, it is the dominant way light couples to internal atomic structure. In this regime, the electric field varies slowly across the size of an atom as the atomic length scale is of order $\sim a_0$ while the optical wavelength λ is $\sim 10^3 a_0$. Physically this means the light's spatial structure is nearly uniform across the electron cloud. Thus the dominant effect of the field will be to push positive and negative charges in opposite directions generating a dipole moment. Strongly allowed optical lines are therefore overwhelmingly dipole in character. Higher-order transitions exist, but their strengths are smaller as they derive from much weaker forces due to the magnetic dipole of the oscillating field or field gradients, so they need much higher intensities and longer times to see.

As the magnitude of this dipole interaction will be proportional to the amplitude of the electric field, the coherent rotation rate between two states must be proportional to both the dipole matrix element and $|\vec{E}|$. We can quantify the rate at which population transferred

between $|e\rangle$ and $|g\rangle$ using the Rabi frequency, defined as

$$\Omega(t) = \frac{\vec{d} \cdot \vec{E}(t)}{\hbar}. \tag{3.2}$$

Assuming a Bloch sphere picture, we can calculate the time that a constant Rabi frequency field must be turned on in order to rotate the state vector by a polar angle, θ, as $t_\theta = \theta/\Omega$.

We will consider how to actually calculate the relevant quantities for the $^2\mathrm{S}_{1/2} \rightarrow ^2\mathrm{P}_{1/2}$ transition of $^{174}\mathrm{Yb}^+$ as this is our system of choice for photon shaping experiments. The relevant quantum numbers are orbital angular momentum L, total angular momentum $J = L + S$ (with $S = 1/2$), and its projection along $\hat{z}$ axis m_J. Electric dipole transitions only couple states with opposite parities yielding the selection rule $\Delta L = \pm 1$. Additionally, as the optical wave front is essentially flat across the electron cloud, changes in angular momentum can only be accomplished by the spin of the absorbed (or emitted) photon. This enforces $\Delta J = 0, \pm 1; J = 0 \nrightarrow J = 0$. The polarization of the absorbed (emitted) photon sets the conditions: $\sigma^\pm \Rightarrow \Delta m_J = \pm 1; \pi \Rightarrow \Delta m_J = 0$.

For a particular set of states with quantum numbers J and m_J, we find the coupling strength using the Wignar-Eckart theorem:

$$\langle J', m'_J|\ e\hat{r}\ |J, m_J\rangle\ = \frac{\langle J'|\, e\hat{r}\, |J\rangle}{\sqrt{2J'+1}}\, \langle J, m_J;\, 1, q|J', m'_J\rangle. \tag{3.3}$$

We see that this formulation separates the radial and angular components of the wavefunction overlap. The inner product on the right hand of the Eqn. 3.3 is referred to as the Clebsch-Gordan coefficient. Interpreting this term, the bra term labels the total angular momentum (J) and its projection about the quantization axis (m_J) of the final atomic orbital whereas '1' is the angular momentum of the emitted photon due to its intrinsic spin which may have ± 1 or 0 projection onto the quantization axis, corresponding to the value of q'. J' and m_J' are the labels for the initial state.

The radial component containing $\hat{r}$ leads to the 'reduced' matrix element. What this means is the reduced matrix element encodes the information about the coupling between two orbitals with no regard for angular momentum and thus can be used to calculate things

like the lifetime of an excited state. On the other hand, the Clebsch-Gordan coefficients capture the *relative* strength between multiple decay pathways from the same level based on angular momentum conservation.

Derivation of the reduced dipole element requires sophisticated atomic structure calculations. Alternatively, it can be determined from the experimentally measured lifetime. Clebsch-Gordan coefficients can be readily looked up in tables of different angular momentum couplings. For the case of decay from the $^2\mathrm{P}_{1/2}, m_J = +1/2$ to the Zeeman levels of the $^2\mathrm{S}_{1/2}$ state, we consider the cases when the emitted photon has right circular (i.e.$q = +1$ corresponding to a σ^+ transition), left circular ($q = -1, \sigma^-$) or linear ($q = 0, \pi$) polarization.

This leaves us with: $\langle \frac{1}{2}, \frac{1}{2}; 1, +1 | \frac{1}{2}, -\frac{1}{2} \rangle = \sqrt{2/3}$ and $\langle \frac{1}{2}, \frac{1}{2}; 1, 0 | \frac{1}{2}, +\frac{1}{2} \rangle = \sqrt{1/3}$, corresponding to the emission of a σ^+ and π photon, respectively. The coefficient for σ^- is 0 expected as there is no manifold with $m_J = 3/2$. The relative branching ratio is calculated by taking the square of these coefficients.

We have now answered, at least schematically, the first two questions laid out at the beginning of this chapter. We have partially answered the third question as it relates to the polarization and frequency of photons produced during relaxation. To further characterize the photons during atomic transitions, we will need to consider the full temporal evolution of an atom undergoing an external drive.

3.2 Modeling the time evolution of atoms

Now that we understand how to describe the coupling between states, we need to address how fields actually drive these transitions. We already have some sense for this based on the Rabi frequency, but to model the complete time evolution of a quantum system undergoing a dynamic drive we need to return to some more fundamental quantum optics.

If we had a two-level system with an infinite excited-state lifetime, then we could simply set initial conditions (i.e.the starting populations for the system) and integrate a time-dependent Rabi frequency to find how much population has been transferred between the states at any future in time. This amounts to solving the time-dependent Schrödinger equation for the

system. We could then reconstruct the full temporal evolution of the two-level system in such a manner. However, real atomic dynamics are more complex. The finite lifetime of excited states leads to spontaneous emission of photons, a process which fundamentally degrades the coherence between states while removing population from the excited state.

The explanation for why spontaneous emission leads to dephasing but stimulated emission does not requires a full quantum field theoretical treatment and is beyond the scope of this thesis. However, a schematic picture is that spontaneous emission arises from the atom's coupling to vacuum fluctuations of the quantized field. This leads to amplitude damping of population at rate Γ and to loss of coherence at a rate $\Gamma/2$. This limits our ability to drive a transition with full control. Additional environmental noise, such a magnetic field fluctuations or experimental limits on parameters like the linewidth of the laser used to drive a transition also contribute to effective dephasing between the atomic state and the experimental control. By contrast, stimulated emission by a coherent drive transfers population while preserving a well-defined phase relation between the atom and the driving mode as the combined atom-field remains in a pure state with a well defined phase relationship, so the process itself does not decohere the atomic superposition. However, as stimulated emission only results in the production of photons back into the driving field, these photons cannot be experimentally useful for typical remote entanglement or network applications as they are 'carried away' with the excitation pulse.

To account for the effects of spontaneous emission and noise, we introduce 'dissipation' operators, also known as 'jump' or 'collapse' operators, into our system evolution. These dynamics are captured in the Lindblad master equation. In the density matrix formalism, we can write the time dependent form of ρ as

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] + \sum_k \left(L_k \rho L_k^\dagger - \frac{1}{2} \left\{ L_k^\dagger L_k, \rho \right\} \right) \tag{3.4}$$

H, L and k are the full system Hamiltonian, the jump operators and index denoting the jump paths.

We will now describe how to write each term of the master equation for the Λ system

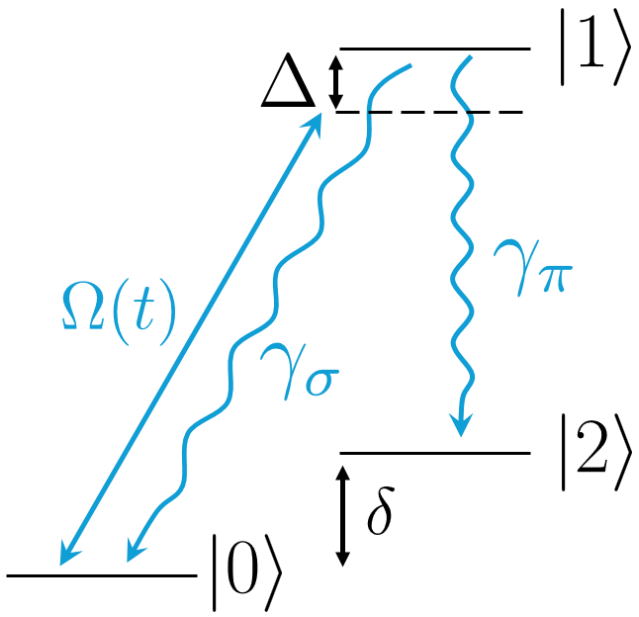

Figure 3.1: A simplified representation of the Λ system relevant for photon shaping showing levels, detunings, excitation and decay pathways.

used in our photon control experiments. In our experiments, we consider the $^2S_{1/2}$ and $^2P_{1/2}$ manifolds of $^{174}Yb^+$. In fact, we can simplify the system even further. Rather than treating all four levels in these orbitals, we will only drive the σ^+ transition from $[^2S_{1/2}, m_J = -1/2] \equiv |0\rangle \Rightarrow [^2P_{1/2}, m_J = +1/2] \equiv |1\rangle$. The P level has two decay pathways leading to the emission of a σ^+ polarized photon in situation that the electron relaxes back to $|0\rangle$ or a π photon when the electron relaxes to $[^2S_{1/2}, m_J = +1/2] \equiv |2\rangle$.

We can write the state of our atom in the Λ configuration as $|\psi\rangle_{atom} = c_0 |0\rangle + c_1 |1\rangle + c_2 |2\rangle$. A representation of the Λ level structure and relevant parameters is shown in Fig. 3.1. The density matrix formalism makes clear the relevant populations and coherence between basis states:

$$\rho \equiv \begin{pmatrix} \rho_{00} & \rho_{01} & \rho_{02} \\ \rho_{10} & \rho_{11} & \rho_{12} \\ \rho_{20} & \rho_{21} & \rho_{22} \end{pmatrix}, \qquad \rho_{ij} \equiv \langle i|\rho|j\rangle = c_i c_j^*, \; i,j \in \{0,1,2\}. \tag{3.5}$$

Populations of $|0\rangle, |1\rangle$ and $|2\rangle$ are given by the diagonal terms where as coherences between states are given by the off diagonal terms.

For a Λ system with $|0\rangle \leftrightarrow |1\rangle$ driven by a laser with envelope $\Omega(t)$, frequency ω_L and controllable phase offset $\phi(t)$, after applying the rotating wave approximation, we can write the interaction Hamiltonian as

$$H_I = \frac{\hbar}{2}\Omega(t)(e^{i(\omega_L t+\phi(t))}\,|0\rangle\langle 1| + e^{-i(\omega_L t+\phi(t))}\,|1\rangle\langle 0|). \tag{3.6}$$

The atomic Hamiltonian simply captures the splitting between the excited state and ground states,

$$H_0 = \hbar\omega_0\,|1\rangle\langle 1|\,. \tag{3.7}$$

where we have initially assumed that the two S Zeeman manifolds are degenerate in energy as is the case when no external magnetic field is present. Thus, in matrix form, we have $H_{total} = H_I + H_0$ which can be explicitly written as

$$H = \frac{\hbar}{2}\begin{pmatrix} 0 & e^{i(\omega_L t+\phi(t))}\Omega(t) & 0 \\ e^{-i(\omega_L t+\phi(t))}\Omega(t) & 2\omega_0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{3.8}$$

Adopting a frame which rotates at frequency ω_0, we can easily include terms for time dependent laser detuning, given by $\Delta(t) = \omega_L(t) - \omega_0$ and any ground state splitting, δ. This yields the rotating frame Hamiltonian

$$\tilde{H}(t) = \frac{\hbar}{2}\begin{pmatrix} 0 & \Omega(t)\,e^{+i\phi(t)} & 0 \\ \Omega(t)\,e^{-i\phi(t)} & -2\big[\Delta(t)+\dot{\phi}(t)\big] & 0 \\ 0 & 0 & 2\delta \end{pmatrix}. \tag{3.9}$$

Because we treat changes in controllable phase as instantaneous, $\dot{\phi}(t) = 0$. The addition of the $e^{\pm i\phi(t)}$ term accounts for relative phase between the atomic states and local oscillator (i.e.laser). As we will only ultimately control the phase-parity of the laser, this means that $\phi(t)$ will always either be 0 or π. In other words, the $\tilde{H}_{0,1}$ and $\tilde{H}_{1,0}$ will be allowed to switch

sign. Intuitively, we see that an addition of a π phase will reverse the direction of population transfer between $|0\rangle$ and $|1\rangle$. The ability to modulate the relative phase will be important when describing photon control in later chapters.

The excited state scatters photons at a rate given by $\Gamma = 1/\tau \approx 1.23 \times 10^8 s^{-1}$ with Γ_π and Γ_{σ^+} determined by the atomic branching ratio. We now introduce collapse operators $L_0 = \sqrt{\Gamma_{\sigma^+}}|0\rangle\langle 1|$ and $L_2 = \sqrt{\Gamma_\pi}|2\rangle\langle 1|$ where the relative amplitude between σ^+ and π is 2:1.

Collapse operators capture the behavior of incoherent channels in system evolution. Beyond spontaneous emission, we can add terms representing the finite linewidth of the laser used to drive this transition. However, as the free-running linewidth of an ECDL is on the order of $\sim$ 10-100 kHz, this is much smaller than the dephasing rate intrinsic to the atomic relaxation and can typically be neglected. If we choose to include additional dephasing terms, we could then write

$$\hat{L} = \hat{L}_{decay} + \hat{L}_{dephasing} \tag{3.10}$$

with

$$\hat{L}_{decay} = \begin{pmatrix} \Gamma_{\sigma^+}\rho_{11} & -\frac{\Gamma_{\sigma^+}}{2}\rho_{01} & 0 \\ -\frac{\Gamma_2}{2}\rho_{10} & -(\Gamma_{\sigma^+}+\Gamma_\pi)\rho_{11} & -\frac{\Gamma_\pi}{2}\rho_{12} \\ 0 & -\frac{\Gamma_\pi}{2}\rho_{21} & \Gamma_\pi\rho_{11} \end{pmatrix}. \tag{3.11}$$

Since our laser only drives the $|0\rangle \Leftrightarrow |1\rangle$ transition, we can write the dephasing due to finite laser linewidth as

$$\hat{L}_{laser} = \begin{pmatrix} 0 & -\gamma_{01}\rho_{01} & 0 \\ -\gamma_{10}\rho_{10} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \tag{3.12}$$

where $\gamma_{0,1} = \gamma_{1,0}$ and equal to the linewidth of the laser. Similar terms could be used to include $\vec{B}$-field noise leading to decoherence between $|0\rangle$ and $|2\rangle$. We understand laser linewidth induced dephasing as a loss of coherence between qudit states and the local oscillator coupling those states due to the phase noise of the oscillator.

All relevant terms can now be entered into the master equation and the time evolution of the system can be solved. In practice, this is almost always done numerically.

3.3 *Describing the photons generated during spontaneous emission*

Photon statistics during a driven transition can be derived from examining the population of the scattering state as a function of time. In the most general form, for an excited state, $|1\rangle$, with lifetime $\tau = 1/\Gamma$ we can write

$$\text{photon flux, } I(t) \;\propto\; \Gamma\,\rho_{11}(t). \tag{3.13}$$

For the Λ configuration we consider, we have multiple decay channels corresponding to the polarization of emitted photons, $q \in \{\pi, \sigma^{\pm}\}$. This yields

$$I_q(t) \;\propto\; \Gamma_q\,\rho_{11}(t), \qquad \sum_q \Gamma_q = \Gamma. \tag{3.14}$$

If we consider an excitation pulse of finite duration, normalizing $I(t)$ yields the temporal intensity envelope of the photons emitted during this process

$$|g_q(t)|^2 = \frac{\Gamma_q\,\rho_{11}(t)}{\langle N_q\rangle}, \qquad \int |g_q(t)|^2\,dt = 1. \tag{3.15}$$

Here, $\langle N_q\rangle$ is the mean photon emission number into polarization mode q given by:

$$\langle N_q\rangle \equiv \int_{t_0}^{t_f} \Gamma_q\,\rho_{11}(t). \tag{3.16}$$

For the case of single emissions, we can write the state of the resultant photons using polarization-time mode creation operators,

$$|1_{g_q}^{(q)}\rangle \;=\; \int dt\; g_q(t)\;\hat{a}_q^{\dagger}(t)\,|vac\rangle\,, \qquad \int dt\,|g_q(t)|^2 = 1, \tag{3.17}$$

where $\hat{a}_q^{\dagger}(t)$ is the creation operator for a photon in polarization mode, q, at time, t.

Because the emission of a π photon projects the atom into a dark state (i.e. σ^+ polarized light cannot excite a photon from $|2\rangle$), we have the condition $\langle N_\pi\rangle \leq 1$. As the atom can produce at most one π photon, we have a two-dimensional subspace spanned by $\{|\text{vac}\rangle,\ |1_{g_\pi}^{(\pi)}\rangle\}$ with $P_0 = 1 - \langle N_\pi\rangle$ and $P_1 = \langle N_\pi\rangle$ respectively. The same approach could be used to characterize the mean photon emission number for σ decay. However, since the emission

of a σ polarized photon projects the atom back into $|0\rangle$, re-excitations are possible and the inequality does not hold. In practice, it is usually sufficient to curtail this basis to just a few polarization-time basis states depending on the strength and duration of the excitation pulse.

For low spontaneous emission probabilities, neglecting any amplitude of multiple emissions due to possibility of re-excitation following a σ path decay, and setting $\delta = 0$ such that Zeeman levels are degenerate, we can write the entangled atom-photon state as

$$|\Psi\rangle = \sqrt{1 - \langle N\rangle}\,|0_A\rangle\,|vac\rangle + \sqrt{\langle N_\sigma\rangle}\,|0_A\rangle\,|1^\sigma_{g_\sigma}\rangle + \sqrt{\langle N_\pi\rangle}\,|2_A\rangle\,|1^\pi_{g_\pi}\rangle \tag{3.18}$$

where $\langle N\rangle = \langle N_\pi\rangle + \langle N_\sigma\rangle$. If we trace over the atom and any unobserved degrees of freedom, the state of the π field is

$$\rho_\pi \;=\; P_0\,|\mathrm{vac}\rangle\langle\mathrm{vac}| \;+\; P_1\,|1^{(\pi)}_{g_\pi}\rangle\langle 1^{(\pi)}_{g_\pi}|. \tag{3.19}$$

This is a mixed state: the off-diagonal coherence between $|\mathrm{vac}\rangle$ and $|1^{(\pi)}_{g_\pi}\rangle$ is lost once we discard information about the atomic state and the detailed emission history. If instead we *condition* (herald) on detecting a π photon in mode g_π, the post-selected photonic state is the pure single-photon wavepacket $|1^{(\pi)}_{g_\pi}\rangle$.

In principle one can always write the full atom–field state as a pure state in a joint Hilbert space. However, once we allow for the possibility of multiple emissions, even in the restricted case where at most a single π photon can be produced while any number of σ photons may be scattered, the pure state becomes a superposition over many different photon polarization-temporal modes and emission histories. The atom may emit zero, one, two, or more σ photons at different times before a terminating π emission (or no π emission at all), and each such sequence of jumps leaves the atom in a different correlated state with $(n_\sigma, n_\pi) \in \{0, 1, 2, \dots\} \times \{0, 1\}$. The exact pure state is therefore a highly entangled sum over all allowed (n_σ, n_π), with amplitudes that depend on the full time record of the jumps, and cannot be expressed solely in terms of the mean photon numbers $\langle N_\sigma\rangle$ and $\langle N_\pi\rangle$. Practically, for cases where the probability of multiple photon emissions is not negligible, the best we can do is to represent the emitted photons as a mixed state.

3.3.1 Quantum jumps approach to photon statistics

The Lindblad master equation treats spontaneous emission as a *continuous* loss of population and coherence. It is ideal for predicting $\rho_{11}(t)$ and, through Eq. (3.13), the *average* photon flux. In the case of a terminating π decay channel this already provides the full temporal profile of the emitted photon and the total emission probability. However, by itself it does not give a very transparent picture of the *distribution* over different photon-number events (e.g. the relative frequencies of zero-, one-, and multi-photon outcomes), which play an important role in photon control for networking applications. Although working in the small $\langle N \rangle$ (i.e. $\ll 1$) regime allows us to ignore most multi-photon processes, many practical atom–photon entanglement protocols benefit from stronger excitations. In that regime it is useful to develop numerical tools that directly access the emission statistics.

To do this, we use a family of closely related techniques variously known as the quantum trajectory approach, quantum jumps methods, or Monte Carlo wave-function simulations. All of these are different 'unravelings' of the same master equation. They generate stochastic realizations of the system's evolution whose ensemble average reproduces the deterministic Lindblad dynamics. In the quantum jumps picture one thinks in terms of an ideal photon detector monitoring the environment. Between detection events the system evolves under an effective non-Hermitian Hamiltonian, and a detection ('click') at time t induces a discrete jump via the appropriate collapse operator L_q for channel $q \in \{\pi, \sigma^{\pm}\}$. The probability for a jump in channel q in a short interval $[t, t + dt]$ is

$$p_q(t)\, dt = \langle L_q^{\dagger} L_q \rangle_t \, dt = \Gamma_q \, \rho_{11}(t)\, dt, \tag{3.20}$$

so that the *distribution of jump times* for the first emission in channel q is proportional to $\Gamma_q \rho_{11}(t)$. In the terminating π channel, this first-jump distribution is precisely the photon's temporal intensity envelope $|g_\pi(t)|^2$.

From a numerical perspective, the Monte Carlo wave-function method simply provides a way to simulate many such quantum trajectories: we propagate the state under the non-Hermitian effective Hamiltonian, draw random jump times according to the rates $\Gamma_q \rho_{11}(t)$,

apply the corresponding collapse operators, and repeat. By collecting statistics over a large ensemble of trajectories, we can directly estimate quantities such as the probability of emitting zero, one, or multiple photons in each channel, the joint distribution of emission times, and the temporal mode of photons conditioned on the prior history of the system.

Put more colloquially, these methods implement can be summarized as coherent drive punctuated by randomly timed, photon emissions whose detection-time histogram *is* the wave packet, at least for a terminating channel or the first emission in a given channel. This is exactly the viewpoint we use later to (i) validate measured temporal waveforms and (ii) quantify multi-photon statistics that are not as easily accessed from the master equation alone. In this way, the master equation and quantum jumps approaches play complementary roles in our photon-control toolbox, and we will draw on both in the following chapters.

Chapter 4

OPTIMAL CONTROL OF SINGLE PHOTONS

This chapter represents the application of the theory laid out in Ch. 3 to the problem of optimal photon control for networking applications. Whereas Ch. 3 tells us how to describe photons generated during a particular emission process, we now need to develop tools that allow us to identify the correct exciting field necessary to produce photons with desired properties.

Subsequent sections detail the construction of such tools. Specifically, we develop novel methods to identify the excitation pulse needed to generate a photon of any temporal profile from any emitter. We also employ established tools in the family of quantum trajectory theory to new ends and describe how to calculate photon statistics for optimal shot selection in quantum networking experiments.

We compare our new techniques to established methods and highlight both the advantages of our approach in flexibility, generalizability and low infrastructure cost while also taking stock of the relative strengths of other alternative strategies for photon control. We discuss the application of photon shaping for hybrid and same qubit networks and quantum state transfer.

4.1 Why finding the optimal excitation pulse is hard

While we can find the temporal evolution of a quantum system provided we know the form of the interaction Hamiltonian by solving the master equation, it turns out the reverse problem is significantly more challenging. The basic reason why is that dissipation/decoherence leads to non-unitary evolution of the system. This is especially apparent if we recall the quantum jumps interpretation of system evolution described in the previous chapter.

The continuous damping of atomic correlations induced by dephasing via spontaneous emission can equally be treated as an aggregation over jumps at many times. The entire point of Monte Carlo methods is that this ensemble picture and the master equation approach lead to identical predictions for system observables.

Returning to the basic axioms of quantum mechanics, we know that, in the absence of measurement, the evolution of a quantum system is unitary. Of course, there are a few different but equivalent ways to express the unitary of a quantum system. We can equally say the evolution of the system is deterministic, meaning that for any given input state and Hamiltonian, the state of the system at a future point can always be calculated. Deterministic evolution also implies that the reverse is true. For any state at time t, we could find the state at any earlier time.

Measurement breaks the condition of unitary evolution. For example, if upon observation we find a two level system to be in $|0\rangle$, without any other information, we have no way of saying whether the state of the system before observation was $\Psi = |0\rangle, \Psi = \sqrt{1/10}\,|0\rangle + \sqrt{9/10}\,|1\rangle$ or any other arbitrary superposition state. In other words, solutions to the photon control problems we want to solve will be non-unique.

We will not be able to find singular, analytic expressions for interaction Hamiltonians to produce photons of a desired form. Instead, we will need to use optimization protocols to test many different Hamiltonians and feed back on solutions until an appropriately good form is found to satisfy our experimental requirements. Before describing those tools, we should establish, in greater depth, why anyone would care about the problem of photon temporal waveform control in the first place.

4.2 Motivation for photon control: robust networks, hybrid systems and state transfer

The ability to control the state of single photons represents a key tool for building quantum networks and enabling mixed-architecture qubit platforms. For example, the production of indistinguishable photons from distinct qubit platforms would allow the application of

established entanglement generation protocols to be used to realize hybrid quantum systems [45, 46]. Such systems could leverage the strengths of distinct qubit types such as the long coherence times of ions or the fast gate speeds of solid state spin donors to out compete any single-qubit type platform [47].

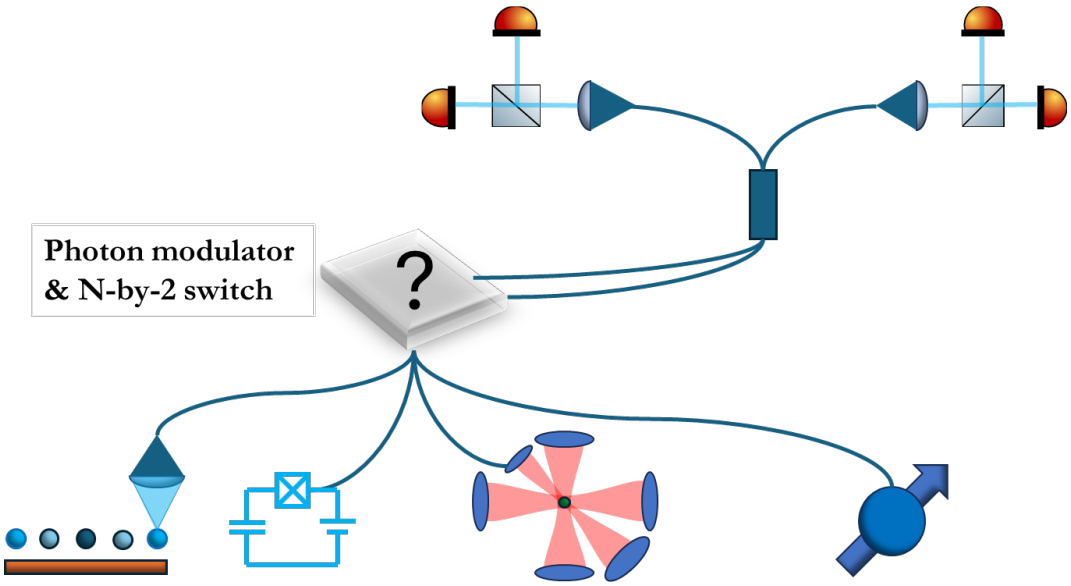

Figure 4.1: Schematic of a future hybrid quantum network. To implement remote entanglement between distinct qubit types, the ability to generate photons identical in all bases will be required. Thus, we need will need some arbitrary photon modulator (APM) capable of either in-situ or post emission photon control. Our work on photon temporal waveform shaping and number state modulation, in tandem with existing frequency and polarization modulation tools, completes a the picture for such an APM.

Building photon-mediated hybrid systems requires control over the state of a flying qubit in all bases including frequency, polarization, number, spatial mode and time bin [48]. Advances in quantum frequency conversion provide a promising avenue to bridge large spectral mismatch between emitters [49, 50]. Photon polarization can be easily controlled via mapping dipole emission fields to lab-frame field orientations [16]. Meanwhile, spatial

mode control can be achieved via the use of cavities [51, 52], fiber based waveguides [19] or simple masking techniques [53] while spatial light modulators present reconfigurability and flexibility for a range of protocols [54].

In tandem with these techniques, control over the photon temporal waveform would allow any two quantum emitters to generate identical photons. By matching the photon detection probability of both nodes by either modulation of emission probability or attenuation of photon number in one arm of the photon mixing apparatus, various single [48, 55, 56] and two [57] photon detection methods, well established in same type qubit networks, could be used to entangle mixed-architecture quantum systems. Such remote entanglement techniques have been independently implemented in neutral atoms [58], solid state systems [59] and trapped ions [60]. A conceptual representation of such a hybrid quantum network is shown in Fig. 4.1. However, to date, no such protocol has been demonstrated between *distinct* qubit platforms. Previously, interference of photons from a trapped ion and (telecom converted) neutral atom ensemble has been achieved; however, photon mixing relied on temporal gating to ensure high contrast [61]. As the efficiency of gating based approaches decreases as the mismatch of excited state lifetimes grows, it is desirable to find an alternative approach for ensuring high similarity between emitted photons. Figure 4.2 and Tab. 4.1 illustrate how temporal gating could be used to generate photon interference for the case of two emitters (in this case a Yb ion and ZnO bound donor qubit) who share a common transition frequency but with a difference in excited state lifetimes of a factor of ~ 8.

Full control of the temporal envelope of single photons would also allow optimization of processes such as quantum state transfer in repeater protocols. By matching the time-reversed natural decay shape of an atomic line, the probability that a photon emitted by one node in a network will be absorbed by a subsequent node can be maximized [62, 63]. Similarly, in atom-cavity systems, tailoring photon shapes to impedance match an optical resonator can be used to maximize flying-stationary qubit coupling [64].

Same qubit type networks can also benefit from task specific tailored single photons. In trapped ion photonic interconnects, for example, photons from distributed traps are

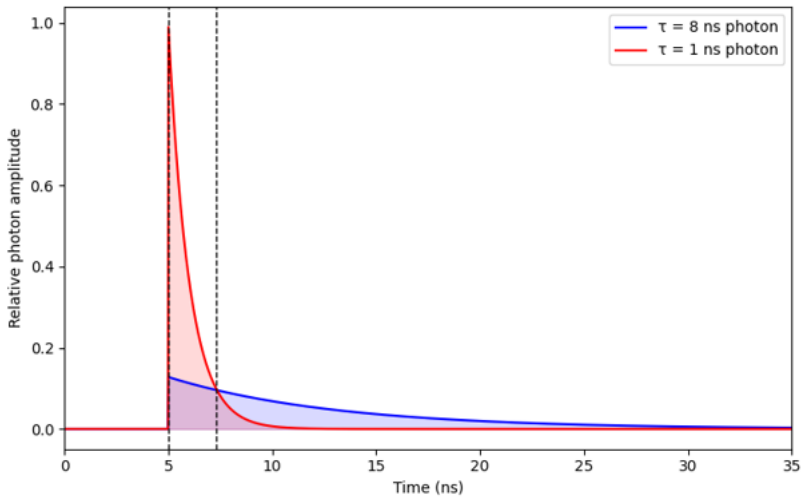

Figure 4.2: Temporal gating approach to hybrid system photon interference. In order to account for differences in excited state lifetime, one approach to generating two photon interference is simply to only detect photons which arrive during a certain time window. This figure illustrates the concept for the case of remote entanglement between a Yb ion and ZnO donor solid state qubit where the lifetimes vary by a factor of ∼8. With such an approach the indistinguishability of photons during a fixed duration is traded off with the photon production probability during the same period. For example, by choosing a short detection window much less than either state lifetime, the shape of the photons from both emitter during that period will be nearly identical. However, the corresponding photon interference rate will be severely limited a factor of $(\int_0^{t_{window}} e^{-t/(8ns)} dt)^2$. This limit is derived from the photon production probability of the slow emitter in this period and the need to match the probability of detection from both qubits.

Two-photon interference rate	Two-photon indistinguishability
0.01	0.997
0.05	0.982
0.10	0.957
0.25	0.867
0.50	0.723
1.00	0.629

Table 4.1: Two-photon indistinguishability as a function of the interference rate for qubits represented in Fig. 4.2.

most often produced via strong excitation using a pulsed laser [26, 65]. In such systems, the resultant photon waveform closely matches the Fourier transform–limited Lorentzian line-shapes of the excited state, yielding a temporal distribution with an exponentially decaying tail characterized by the state lifetime. While photons from two such nodes will interfere perfectly assuming no timing or interferometric error sources, the fidelity of remote entanglement operations for such experiments may be degraded by non-common mode path fluctuations of each arm of the photon detection apparatus [48]. The deleterious effects of path dependent phase fluctuation can be minimized by appropriate selection of photon distributions with gaussian waveforms being the most robust to path length instability [66]. Two representative target photon distributions optimized for state transfer and interferometric stability are shown in Fig. 4.3.

Timing jitter of excitation pulses may lead to similar effects. For instance, in the excitation protocol we describe in Ch. 5, the dominant source of timing error is set by the sample rate accuracy of the arbitrary waveform generator (AWG) used to synthesize optical pulses. Assuming the trigger source is synchronized to the AWG clock, typical root-mean-square

pulse jitter is on the order of 150 ps. If we calculate the indistinguishability of two photons generated by two excitation pulses with a 150 ps delay between nodes from emitters with 8 ns lifetimes, typical of optical transitions used for remote entanglement in atoms/ions, we can get a sense of the relative error induced by such timing jitter. In the case that we use the typical short pulse protocol leading to the generation of (nearly) exponentially decaying photons, this yields an indistinguishability of 0.963. However, when excitation is driven such that the resultant photons have distributions described by a gaussian with a full-width-half-max of 15 ns, the indistinguishability increases to 0.9997. This corresponds to a reduction of error from timing jitter by more than a factor of 10^2. As such, we see that the ability to produce photons of arbitrary temporal shape can prove crucial in long distance, high fidelity distribution of flying qubits.

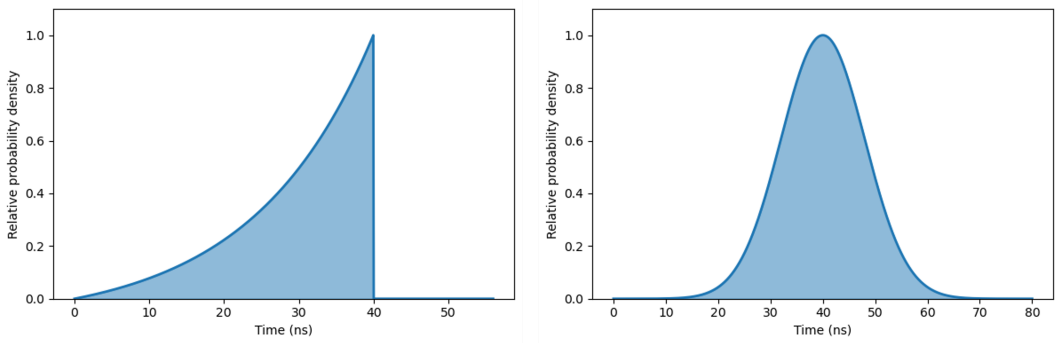

Figure 4.3: Target photon distributions optimized for quantum network applications. **Left:** exponentially rising distributions are desirable for quantum state transfer experiments as such photons have the maximum likelihood to be absorbed by quantum nodes (either free space atoms, or cavity coupled systems) with matched lifetimes. **Right:** photons with gaussian distributions are most robust to interferometric timing error and phase noise.

Beyond temporal gating, discussed above, alternative approaches to the control of photon temporal waveform have made use of post-emission tools such as circuit based pulse shaping via dispersive optics [67], electro-optical modulation of photon wave packets [68–70] or

measurement based post selection techniques [71]. However, such methods are not widely applicable across all frequency bands owing to hardware limitations. Additionally, post-emission shaping may introduce additional noise or loss into the system. Table 4.2 describes some of the current approaches for photon temporal control including relevant strengths and weaknesses.

Technique	**Strengths**	**Limitations**	**Requirements**
Cavity STIRAP	Deterministic emission into a well-defined cavity mode; shaped photons via adiabatic control	Requires high cooperativity and adiabaticity; added system complexity; wavelength tied to cavity	High-finesse cavity, stable lasers, active locking
Dispersive optics / time lenses	Flexible spectral and temporal reshaping with good efficiency	Phase-matching and bandwidth limits; dispersion loss; typically wavelength-specific hardware	Dispersive media, high-speed phase modulators, synchronized waveform sources
Temporal gating	Applies to any emitters; simple to implement	Severe efficiency loss for large lifetime mismatch	Fast detectors or Pockels-cell gating; precise synchronization
Post-emission modulation (EOM/AOM)	Arbitrary waveform shaping; reconfigurable	Insertion loss; finite modulation depth; feature size limited by bandwidth	High-bandwidth modulators and drivers; stable RF/optical sync

Table 4.2: Comparison of photon-shaping techniques: strengths, limitations, and requirements.

Approaches which aim to control photon waveform during the emission process by embedding the emitter in a cavity have been widely demonstrated or proposed in many qubit platforms including ions [72], quantum dots [73] and neutral atoms [74]. Cavity-based Stimulated Raman Adiabatic Passage (STIRAP) experiments modulate driving fields to control photon shape during emission with the adiabatic condition imposing constraints on achievable waveforms. All such approaches require specialized optics and are necessarily coupled to the resonator modes [75–77]. These constraints place limits on the rate of of photon generation and face practical challenges for scaling due to the need for additional infrastructure in already crowded qubit platforms. Moreover, while emission into a cavity mode is advantageous for some networking tasks, it can be less suitable for protocols that encode atom–photon entanglement in multiple orthogonal modes (e.g., polarization or frequency two-photon detection schemes), unless the resonator is explicitly engineered to support multiple qubit emission paths [78].

Instead, the technique we demonstrate makes use of a free space quantum emitter wherein shaping of the photon temporal waveform is realized by deterministic modulation of the population of the state undergoing radiative decay. This control is realized by combining optimization of the pulse envelope with parity phase control for a field which couples a ground and excited state.

4.3 Any emitter, any photon temporal waveform

For our photon control experiments, we use $^{174}Yb^{+}$ as our platform of choice. The ultimate motivation for the use of Yb is to implement a hybrid quantum system as described in 'Photon-mediated entanglement scheme between a ZnO semiconductor defect and a trapped Yb ion' [46]. The protocol described therein makes use of the hyperfine splitting of the $^{2}S_{1/2}$ in $^{171}Yb^{+}$ to encode the qubit logical basis for the ionic node. However, we opt to use the 174 isotope for photon shaping demonstrations due to its simpler level structure as described in Chapter 3. As discussed in that chapter, we will use σ^{+} polarized light to drive the $^{2}S_{1/2}, m_J = -1/2 \equiv |0\rangle \Rightarrow {}^{2}P_{1/2}, m_J = +1/2 \equiv |1\rangle$ transition and collect π photons emitted

as the state relaxes to ${}^2\mathrm{S}_{1/2}, m_J = +1/2 \equiv |2\rangle$. We have already derived the pertinent form of the interaction Hamiltonian for this Λ system and described how to model the temporal evolution of the system. The atomic level diagram for our experiment, including state labels and representations of the intensity profiles of emitted photons, is presented in Fig. 4.4.

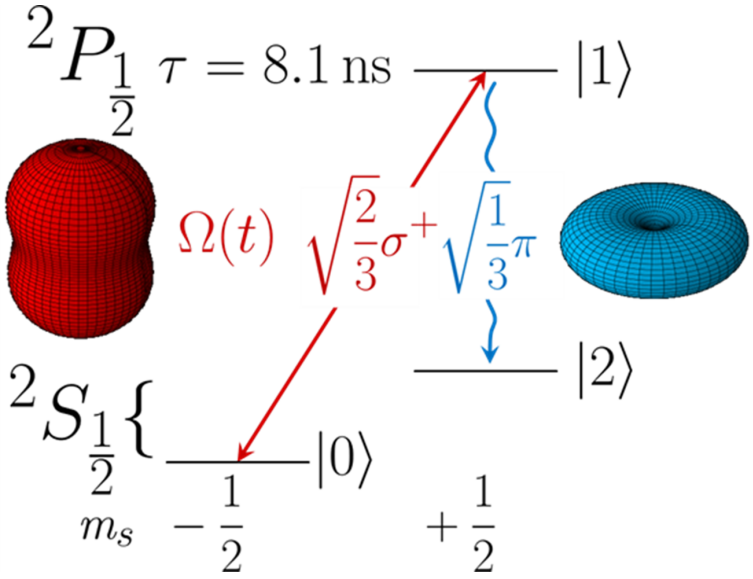

Figure 4.4: Partial level structure of ${}^{174}\mathrm{Yb}^+$ relevant for photon control experiments. Optical pumping to $|0\rangle$ via the application of σ^- polarized light (not shown) initializes the system. The application of a time-dependent complex Rabi frequency, $\Omega(t)$, drives the transition from $|0\rangle$ to $|1\rangle$ to control the population, $\rho_{11}(t)$. Spontaneous emission from $|1\rangle$ generates photons with a temporal waveform at rate $\Gamma\rho_{11}(t)$ in a superposition of σ^+ and π polarizations with a relative intensity of 2:1.

The key element of this work is its ability to be applied to a wide range of qubit platforms. Our broadly applicable approach has only two requirements for a candidate qudit: (1) control of the phase-parity and (2) modulation of the amplitude of a field coupling a ground state to an excited manifold which produces a photon during relaxation. By adapting the Hamiltonian describing the time evolution of this system, this approach can be applied to

higher dimensional Hilbert spaces corresponding to more general multi-level systems. Below we provide a generalizable 'recipe' to find the optimal interaction Hamiltonian to produce photons of any desired temporal waveform.

4.3.1 Photon shaping recipe

Our basic approach to produce a photon of a desired shape is as follows:

1. Initialize the system in the $|0\rangle$ state, corresponding to $|^2S_{1/2}, m_J = -1/2\rangle$. Experimentally, this is achieved by the application of appropriately polarized light based on the orientation of the quantization axis.

2. Define an interaction Hamiltonian which represents a driving field coupling $|0\rangle$ to $|1\rangle$ (i.e. $|^2P_{1/2}, m_J = +1/2\rangle$). Experimentally, this represents a shaped, resonant, and σ^+ polarized laser pulse. In order to simulate realistic parameters, we restrict the maximum Rabi frequency to 250 MHz based on standard available laser power and optics for typical dipole transitions in ion trapping.

3. Solve the system time-evolution to find $P_1(t)$, where $P_1(t) = |\langle 1|\psi(t)\rangle|^2$. As $|1\rangle$ undergoes spontaneous emission, the temporal distribution of the resulting photon follows $P_1(t)$. This process is described by solving the Lindblad master equation. Numerically, this is achieved by modeling the system using Qutip, a python based quantum optics library. Experimentally, the laser pulse is applied and the resultant time of detection of the photon is recorded. The state is then reinitialized in $|0\rangle$ and the process is repeated until a sufficiently dense photon histogram can be constructed.

4. Once the form of the emitted photon is determined, iteratively update H_{int}, or equivalently, the laser pulse shape, utilizing a variational algorithm until the predicted or measured photon shape converges with the desired distribution.

Depending on the application of photon shaping, particular details of wavefunction optimization will vary; however, the basic approaches and considerations described throughout this and subsequent sections continue to apply. For example, in our remote entanglement application, we describe photons optimized specifically for single-photon remote entanglement schemes such as the paradigmatic 'Cabrillo' protocol [79]. In such experiments, two remote matter qubits undergo weak, synchronized, pulsed excitation in order to produce identical photons with low single photon emission probabilities. These photons are then combined on a beam splitter and counted with single photon detectors such that a single photon measurement projects the joint matter qubit system into a Bell state. As such, we aim to optimize the temporal wave packet of photons corresponding only to π emission and trace out the contribution of σ^+ events. However, for applications such as the standard 2-photon polarization based remote entanglement schemes commonly employed in trapped ions [80], the entire photon temporal wavefunction corresponding to both π and σ^+ must be considered. For such an experiment, the same basic approach to waveform optimization can be employed with the only differences being the need to track the temporal waveforms of both polarization paths and different merit functions for photon emission statistics.

4.3.2 Limits on photon shaping

Before we apply our 'recipe', it is worth deriving a few more relations describing our photon control process. Although photons of any temporal waveform can be produced by any emitter, it is not true that for any waveform we can arbitrarily specify the number state basis of the emission field.

We can, however, arrive at certain bounds which guide us in understanding what the maximum possible single photon emission probability is for a target temporal distribution, $g(t)$. We can derive this bound by relating the target distribution to the limits of the excited state probability. We can write

$$P_1(t) = s \times g(t) \tag{4.1}$$

where s is a scaling parameter that enforces $P_1(t)_{max} \leq 1$. To find an upper limit for $\langle N_\pi \rangle$, we set $s = 1/\mathbf{max}(g(t))$. This yields

$$\langle N_\pi \rangle_{\max} \leq \frac{\Gamma_\pi}{\mathbf{max}(g(t))} \int_{t_0}^{t_f} g(t) \left[1 - e^{-\Gamma(t_f - t)}\right] dt. \tag{4.2}$$

The bracketed element in the above equation accounts for the fact that only a fraction of the population at some time, t, will contribute to photon emission by the completion of the pulse cycle at time t_f. The inequality arises due to the incoherent nature of spontaneous emission leading to a limit on the ability to deterministically populate (or depopulate) the excited state at any time, $t > t_0$. The amplitude of this effect relies heavily on the particular form of $g(t)$ and the characteristic dephasing time of a particular emitter with no general analytical form for any arbitrary distribution. However, emitters with relatively higher intrinsic dephasing time, T_2^*, and photon distributions with heavier weights at early times will more closely approach this limit.

We will discuss the effects of dephasing on photon shaping in greater depth in subsequent sections. Throughout the rest of the text, we define fidelity of photon waveform between the desired distribution, $g(t)$ and the predicted, or measured, shape, $\varepsilon(t)$ where both distributions are normalized to be the mode overlap as given by

$$\mathcal{F} = \left(\frac{\int g(t)\, \varepsilon(t)\, dt}{\sqrt{\int g(t)^2\, dt} \sqrt{\int \varepsilon(t)^2\, dt}} \right) = |\langle g(t) | \varepsilon(t) \rangle|. \tag{4.3}$$

As such, our fidelity metric does not explicitly account for $\langle N \rangle$. Instead, we can consider the maximum achievable value of $\langle N \rangle$ at a particular $\mathcal{F}$ threshold for a given target waveform, $g(t)$.

4.3.3 Photon classes

Before moving on to discuss results of numerical pulse shape optimization, we provide a quick framework for categorizing photon distributions. The utility of these classifications will become apparent in later sections. The basic need for categorization arises from the diverse

system dynamics which places limits on certain photon shaping parameters, such as fidelity of temporal waveform preparation and number state, depending on the target distribution. As such, the associated photon classes provide a qualitative framework for understanding what degree of control we have to generate photons which fall into each class. Representative photon distributions of each class are shown in Fig. 4.5.

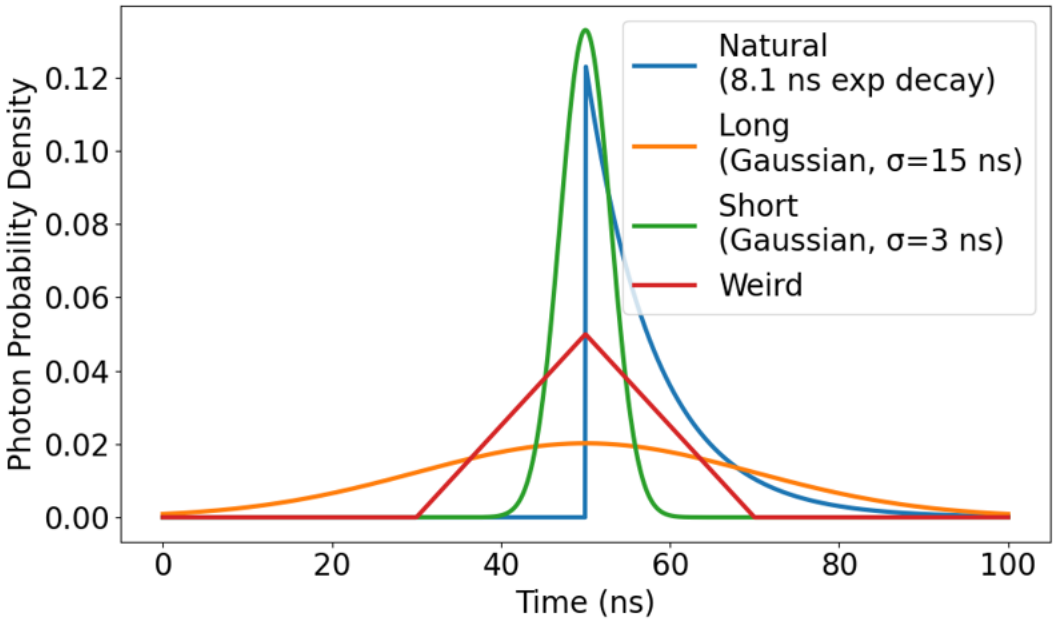

Figure 4.5: Representative distributions for different photon classes. 'Natural' photons have a profile determined by the Fourier transform of the intrinsic atomic linewidth, leading to an exponentially decreasing amplitude. 'Long' photons show greater relative weight in the tails than the natural distribution. 'Short' photon have a narrower distribution than natural. 'Weird' photons contain discontinuities.

1. *Natural*: The lifetime-limited spontaneous-emission waveform, i.e. an exponential decay.

2. *Long*: Distributions with heavier tails than the natural case (greater relative weight at late times; effectively longer duration than the lifetime set by the natural decay).

3. *Short*: Distributions narrower than the natural case (compressed in time relative to the lifetime-limited width).

4. *Weird*: Distributions with discontinuities, such as triangular distributions or exponentially rising waveforms. By definition we exclude the natural photon waveform from this class.

4.4 Simulating photon emission

We are now ready to return to the photon optimization recipe. To actually implement this recipe, we need to concretize exactly how steps (3) and (4) can be implemented. The details to do so are as follows:

1. A normalized target photon shape and emission probability is specified. We discretize the distribution into a fixed number of time bins. We typically use a bin width of 1 ns as this is about one order of magnitude shorter than the excited state lifetime which sets the characteristic atomic response time.

2. An initial interaction-Hamiltonian is defined to capture the time-dependent laser pulse used to drive the atomic excitation. As a first solution to 'seed' our optimization algorithms, we find that a simple copy of the target photon distribution works well as a *relative* scaling of the excitation pulse. This initial solution encodes the relative amplitude of $\Omega(t)$.

3. In order to achieve the desired emission probability, the initial solution is scaled so that the pulse integral leads to the correct value of $\langle N \rangle$. This is done using a simple bracketed root-finding algorithm (Brent's method). For each scaling, the interaction Hamiltonian is updated and the master equation is solved to find $\langle N \rangle = \int \Gamma \times \mathrm{P}_1(t) dt$.

4. The appropriately scaled initial laser pulse is then used to compute $\mathcal{F}$.

5. Depending on the application, we may care more or less about controlling for both $\mathcal{F}$ and $\langle N \rangle$. In cases where we care only the temporal profile of the emitted photon, $\mathcal{F}$ serves as the function to optimize for. If emission probability is also a concern, then

$$O = \mathcal{F} \times \frac{min(\langle N_{target} \rangle, \langle N_{predicted} \rangle)}{max(\langle N_{target} \rangle, \langle N_{predicted} \rangle)} \tag{4.4}$$

   is used as the fitness function.

6. One of two optimization techniques, described in the next subsection, is used to iteratively update the form of the interaction Hamiltonian until a target state preparation fidelity is reached.

4.4.1 Numerical methods

We experimented with a variety of optimization techniques to identify the best methods for arriving at an optimal form of the interaction Hamiltonian. Methods tested included Bayesian optimization, Nelder-Mead, particle swarm optimization, simulated annealing and evolutionary strategy algorithm based approaches. Of these, we found simulated annealing and genetic algorithms to be the most efficient in optimizing the system, measured by time taken for solutions to be found which achieve a desired fidelity. However, we should note that as the author's expertise in selection and design of optimization algorithms is limited, this should not be understood as ruling out the possibility that of any of the other methods name above or others not included in this list may actually be the best choice for this application.

As the principal intended audience for this work is other experimental physicists whose familiarity with such algorithmic methods is not assumed, we will sketch out a brief overview of how simulated annealing and evolutionary algorithms can be used to solve this problem.

For evolutionary optimization, we specifically use the Covariance Matrix Adaptation Evolution Strategy (CMA-ES). In our optimization we treat the laser envelope as a vector of parameters and let a 'population' of candidate waveforms explore that high-dimensional space. This is done by injecting small, random changes into the amplitude of Ω at various

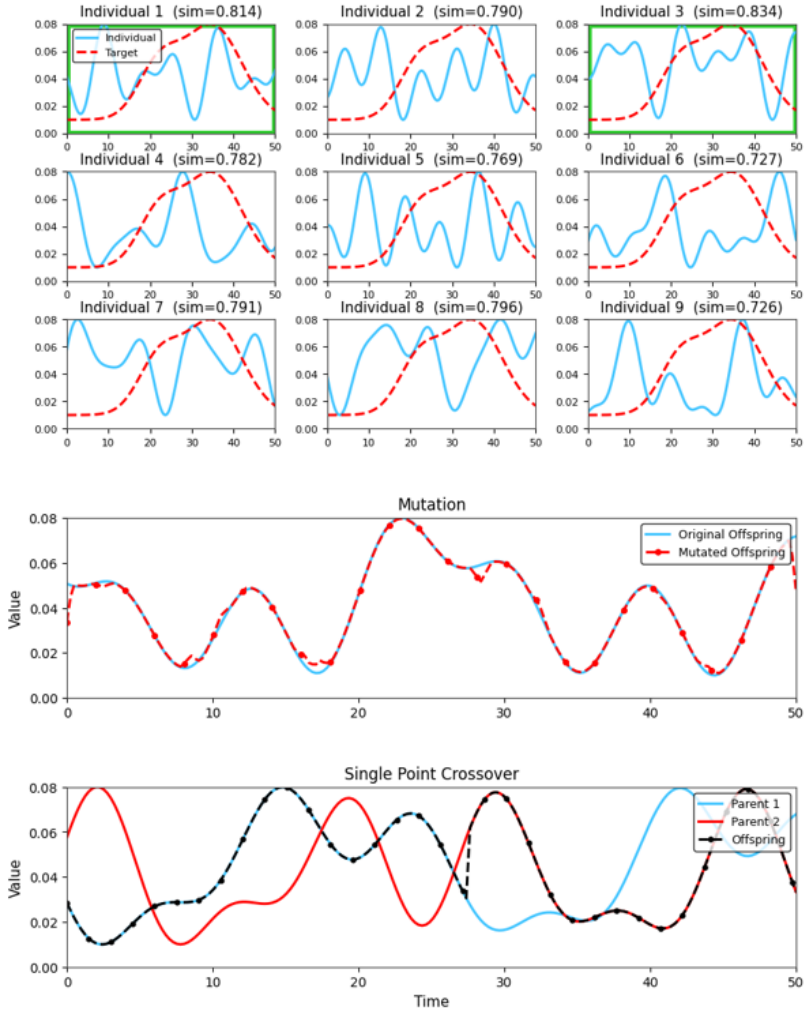

Figure 4.6: Conceptual approach to waveform optimization using evolutionary algorithms. **Top:** an initial population of candidate envelopes ('Individuals 1–9', blue) is compared to a target waveform (dashed red). The individuals with the best match to the target waveform (boxed in green) are selected as the 'parents' for future generations. **Middle:** mutation occurs when one candidate is perturbed locally to create a slightly different offspring. **Bottom:** single-point crossover occurs when two parent waveforms exchange segments to form an offspring that inherits features from both. The algorithm iterates selection, mutation, and crossover, promoting high-fitness candidates so the population converges toward the target.

points in time. The genetic analogy here is that these changes act like mutations of the parent genome, consisting of the discretized laser pulse shape/Rabi frequencies. CMA-ES samples new candidates by drawing Gaussian perturbations around the current best guess. After we evaluate each candidate with our fitness metric, the algorithm keeps the more successful ones and updates both the center and the full covariance of the sampling Gaussian. This is akin to selection of fittest individuals in a given generation. The fittest individual solutions are then mixed and mutated to produce the next generation of solutions. This continues until a solution is found which produces a photon that satisfies the fidelity constraint we apply. A conceptual representation of evolutionary algorithims is shown in Fig. 4.6.

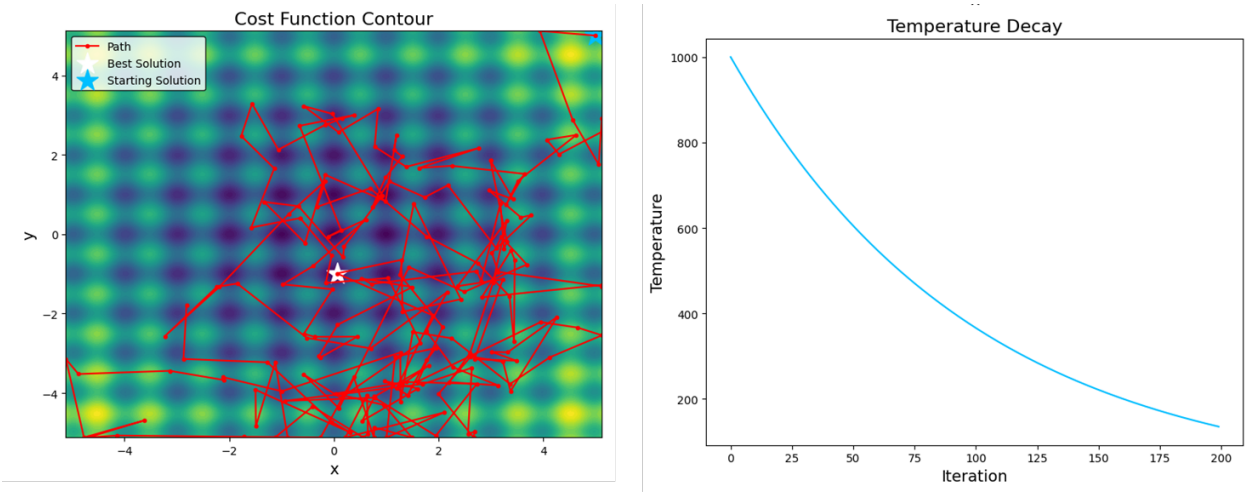

Figure 4.7: Conceptual approach to waveform optimization using simulated annealing algorithms. **Left:** a search trajectory (red) wanders over a cost landscape. Cooler temperatures correspond to higher fidelities solutions. Early in the run, high temperature allows occasional uphill moves, preventing the trajectory from becoming stuck in local minima. As cooling progresses, steps localize and the path settles near a minimum. **Right:** the corresponding exponential temperature decay curve.

Simulated annealing (SA) draws inspiration from the cooling of hot metals wherein a

crystalline atomic lattice rearranges to find lower energy configurations. We start from an initial envelope and make small random changes to the waveform. This is implemented as a small Gaussian bump, increasing or decreasing a range of Ω values, centered at a random time index. Each new form of $\Omega(t)$ is evaluated by the same fitness function. If the candidate produces a better match to the desired target photon shape, we accept it. If the solution is worse, we may still accept it with probability $e^{\Delta f/T}$ where Δf is the fitness change and T is a 'temperature.' Early on, T is high, so the algorithm freely explores and can jump out of local minima. As T is reduced by a cooling schedule, acceptance of worse moves becomes rarer, letting the search settle into a good solution. We track and return the best waveform found. A conceptual representation of function optimization using SA methods is shown in Fig. 4.7.

For both genetic and simulated annealing based algorithms, we may apply some smoothing of the final waveform to generate cleaner solutions.

4.4.2 Optimization results with $\Omega(t)$ amplitude modulation

Initial attempts at waveform optimization relied solely on control of the instantaneous Rabi frequency amplitude to optimize predicted photon distributions. Numerical results demonstrate that that pulse amplitude modulation is sufficient to produce long photons of any $\langle N \rangle$ with arbitrarily high fidelities. Such control is already sufficient for the realization of hybrid system entanglement protocols such as that describe by Lilieholm et al [46]. For any two emitters of different lifetimes, a target distribution which is long relative to both qubits natural photons is readily achievable.

However, empirically we established that as target photon distributions are made shorter and shorter vis-à-vis the state lifetime, the 'amplitude modulation only' approach breaks down. In other words, we cannot generate short or weird photons with high fidelities by control of $|\Omega(t)|$ alone. Figure 4.8 illustrates the results of amplitude only optimization for typical long and short photons. The reasons for these limitations are threefold:

1. The contribution of spontaneous emission to $|0\rangle$ at times early in the excitation process

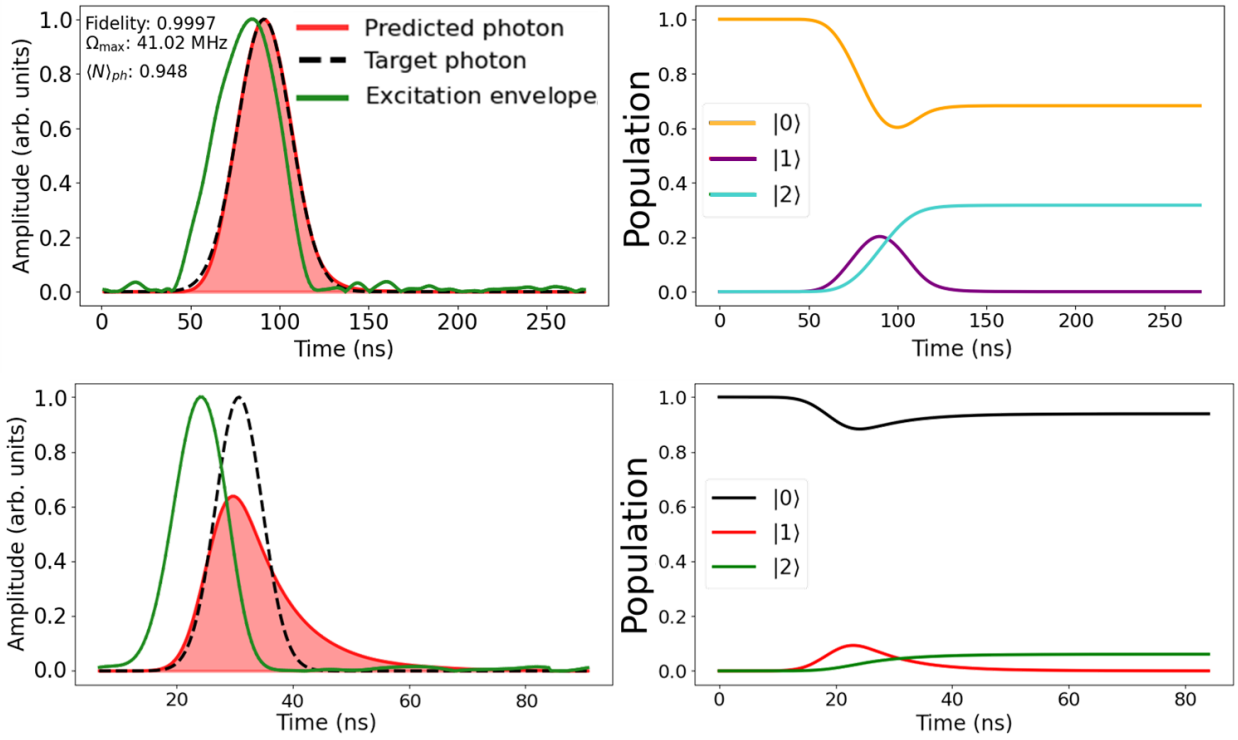

Figure 4.8: Numerical results of photon shaping with Rabi frequency amplitude modulation for a long (**top**) and short (**bottom**) photon. Long photons can be produced with arbitrarily high fidelity and photon emission probabilities. However, short photons show the characteristic stretched tail beyond the target distribution owing to the inability to depopulate the excited states faster than the exponential decay determined by the excited state lifetime.

leads to some amplitude of incoherent re-excitation. This error grows with $\langle N \rangle$ and is minimized for emitters with a small $|0\rangle : |2\rangle$ branching ratio. This contribution is also minimized for small ratios of photon width to excited state lifetime, τ, as such distributions suppress the probability of multi-photon events.

2. Intrinsic dephasing time between $|0\rangle \rightarrow |1\rangle$ is limited, in the absence of any environmental noise, by the relaxation rate of the excited state. This can alternatively be understood as a loss of coherence between the qudit transition and the driving field. This also leads to a contribution of incoherent excitation with an amplitude which grows with $\langle N \rangle$.

3. Decay dynamics cannot be inverted without a coherent π pulse of unphysically short duration preventing turning points or discontinuities in $P_1(t)$.

For a given atomic system, points (1) and (2) are unavoidable and resultant from the fundamental characteristic of the transition employed in photon shaping. The magnitude of their contributions to infidelity in photon generation depend highly on the transition chosen, the desired photon shape and target $\langle N \rangle$.

However, as both effects result from the intrinsic dephasing resulting from spontaneous emission, their contributions to total infidelity can be limited by operating in a low $\langle N \rangle$ regime if a decrease in photon generation rate can be tolerated. The next subsection illustrates how a protocol could be optimized to maximize remote entanglement rate under such constraints for a hybrid system.

4.4.3 Amplitude only modulation for hybrid entanglement

Let us us consider the case of a hybrid 2-photon based remote entanglement schemes where each node is composed of an emitter of distinct lifetime and the operator only has control over photon waveform of the emitter with the shorter lifetime. The operator wants to optimize the process to maximize total entanglement generation rate given by the product

of the 2-photon indistinguishability and total 2-photon detection rate. Indistinguishability decreases with $\langle N \rangle$, due to loss of full coherent control as dephasing effects accumulate with greater emission probability. Detection rate benefits from increasing $\langle N \rangle$ by a factor of $\mathbf{min}(\langle N \rangle_{qubit1}, \langle N \rangle_{qubit2})^2$. The selection of an intermediate $\langle N \rangle$ value will thus be ideal.

A realistic merit function, representing total stationary qubit-photon entanglement generation purity for one node, may take the form

$$R = \langle N \rangle \cdot \mathcal{F} \cdot P_{\mathrm{H,V}|\pi,\sigma} \cdot P_{1|1 \text{ detected}}, \tag{4.5}$$

where $P_{\mathrm{H,V}|\pi,\sigma}$ quantifies the mode overlap of π and σ photons to a defined lab polarization. $P_{1|1 \text{ detected}}$ accounts for the fraction of single photon detection events that correspond to emission processes where exactly one photon was generated (i.e. accounting for multi-photon probability). $\mathcal{F}$ denotes the photon temporal waveform preparation fidelity.

Thus, the full entanglement generation rate will be given by $R_1 \times R_2 \times (1/T_{shot}) \times \eta$ where R_1 and R_2 are the photon entanglement purity per node, T_{shot} is the full time for each experimental trial (including necessary initialization, reset, etc.) and η encodes the complete system transmission and detection efficiency. As $\mathcal{F}$ typically decreases sub-linearly with $\langle N \rangle$, the corresponding 2-qubit entanglement generation rate will increase with $\langle N \rangle$. However, the fidelity of the entangled states will decrease. As such, there will be an additional tradeoff between entanglement generation rate and total 2-qubit entangled state preparation fidelity. The optimal selection of a corresponding $\langle N \rangle$ (and by extension limit of $\mathcal{F}$) will therefore necessarily be dependent on particular requirements specific to each experiment.

4.4.4 Phase parity control and coherent de-excitation

While the generation of long photons proves suitable for many applications, ideally we would like to have the ability to produce photons of any temporal distribution, regardless of the lifetime of the emitter used. We have already seen how the scattering rate puts a hard limit on the emission probability for any photon emission process. However, if we only consider the relative temporal profile, we can in fact 'beat' the excited state lifetime to produce short

and weird photons.

Our technique to generate short photons relies on an additional degree of freedom in the interaction Hamiltonian describing excitation, beyond simply modulating the instantaneous Rabi frequency. We introduce the additional flexibility of modulating the relative phase of the driving field over the course of a single excitation pulse. Experimentally, this amounts to advancing or delaying the wavefront of the excitation laser pulse by an (ideally) instantaneous fixed phase during a single shot.

We motivate this approach by analogy to the interaction between a light field and a Fabry-Perot resonator wherein a resonant optical wave packet excites a single cavity mode with light leaking out of the cavity at a rate determined by the resonator quality factor. By advancing the phase of the laser by π part way through the pulse duration, the mode field amplitude can be reduced to 0 at a rate faster than would otherwise be possible. Returning to the single atom picture (which is after all just a quantum mechanical resonator), in the rotating frame we see that a step $\phi \rightarrow \phi + \pi$ is equivalent to $\Omega \rightarrow -\Omega$. In other words, a negative Rabi frequency pumps population out of the excited state through stimulated emission.

This phase modulation technique can also be understood by adopting a the Bloch sphere picture for the transition between $|0\rangle$ and $|1\rangle$ where in the effects of spontaneous emission are to induce dephasing and relaxation of the state vector. For the 3-level system which describes the $^{174}Yb^{+}$ ion, the picture is more complicated as decay to $|2\rangle$ leads to non-conservation of the probability current in the sub-space of $\mathrm{Tr}_{|2\rangle}\rho$. However, because we only care about the temporal evolution of $|1\rangle$, these effects can be ignored. The application of a time varying Rabi frequency drives the state vector from $|0\rangle$ to $|1\rangle$. Advancing the phase of $\Omega(t)$ by π can be thought of either as a virtual phase gate (ie. a rotation about $\hat{z}$) or a reversal of the direction of motion of the state vector on the Bloch sphere. Thus, by controlling the relative rate and direction of rotation of the state vector, a solution can be found for any desired normalized profile of $P_1(t)$. A conceptual representation of this process is shown in Fig. 4.9.

The inclusion of arbitrary phase in our system Hamiltonian leads to a quadratic increase

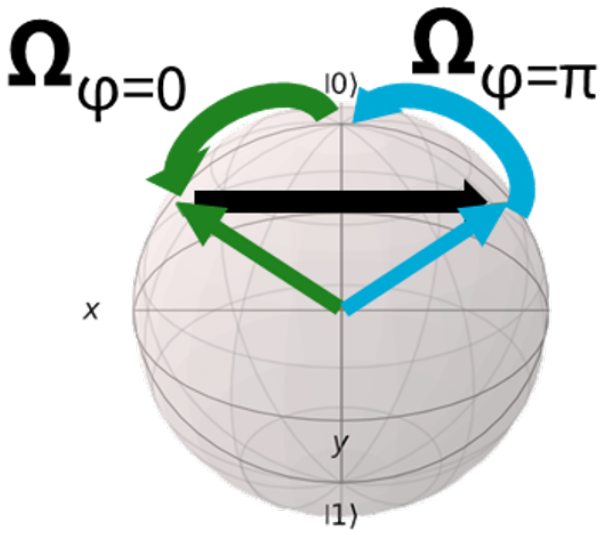

Figure 4.9: Concept behind arbitrary photon shaping with Rabi frequency phase and amplitude modulation shown on the Bloch sphere. First, a pulse with time varying amplitude (blue) drives the state vector along an arc from $|0\rangle$ to $|1\rangle$. Then, part way through the pulse, the phase of the driving field is immediately advanced by π radians, represented by the black arrow. In the qubit frame, this amounts to a virtual phase gate which instantaneously rotates the state vector by π radians about $\hat{z}$. Finally, the out-of-phase pulse component (green) drives the state vector back to $|0\rangle$.

in computational cost and introduces additional complexity for practical realization. Instead, recognizing parity control alone of phase is sufficient to depopulate the excited state for any time scale, we allow the driving field to take on negative Rabi frequencies in our simulation, corresponding to out-of-phase components of the laser. This allows efficient computation of optimized excitation pulse shapes without increasing the dimensionality of our solutions. Figure 4.10 shows snap shots for the optimization process to generate a short photon.

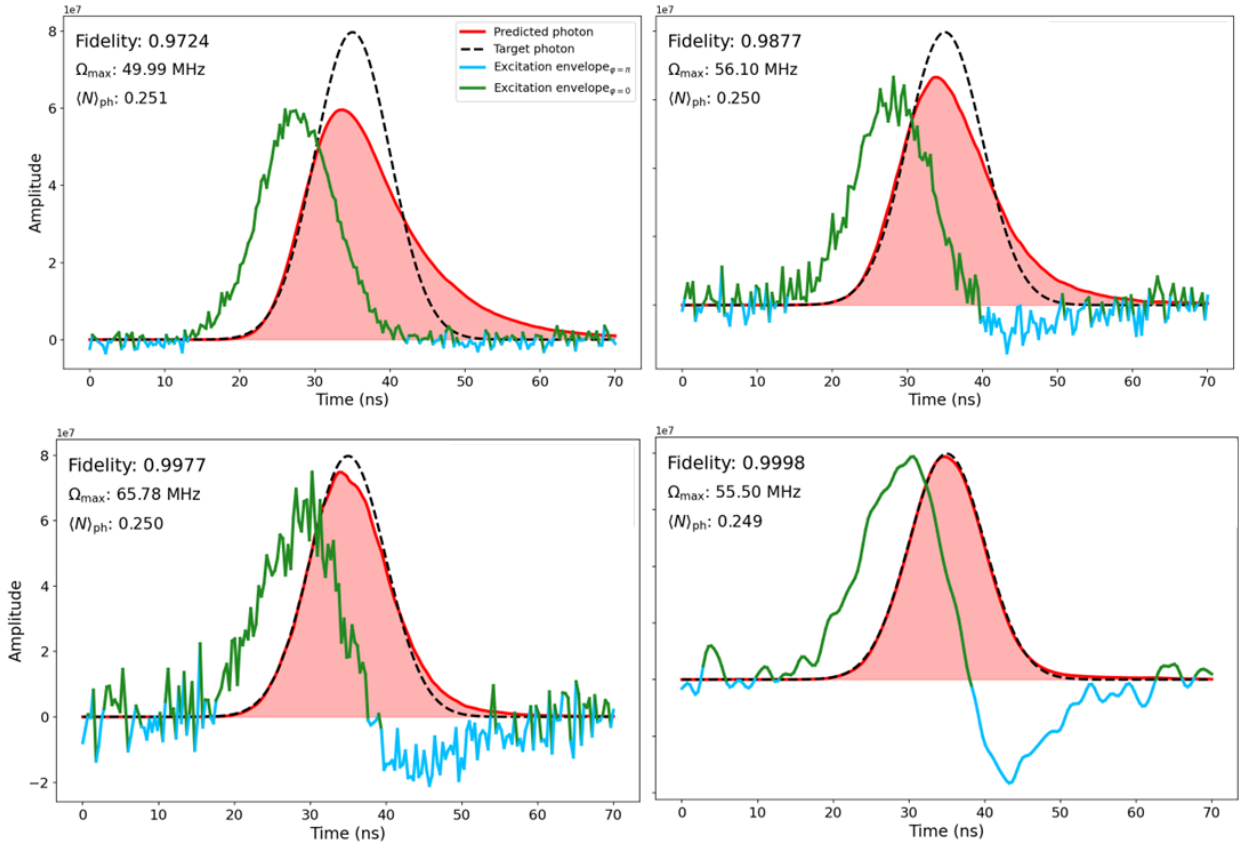

Figure 4.10: Numerical optimization process showing pulse shape and predicted photon at various stages. The progression of pulses from initial to final is top left → top right → bottom left → bottom right.

In tandem with amplitude modulation, bit-wise phase control allows any emitter of any branching ratio and lifetime to produce photons of any target profile with an arbitrarily high fidelity at the trade off of lower $\langle N \rangle$. In other words, the only *intrinsic* error to photon waveform generation stems from the lifetime of the excited state. Thus, we imagine the quantum network user evaluating the combined photon generation rate and waveform fidelity merit function and then optimizing for a specific task. The numerical methods utilized to calculate optimal excitation pulses can easily be made to either maximize $\langle N \rangle$ for a fixed fidelity threshold or alternatively enforce a desired photon emission rate and calculate the best possible photonic waveform under such a constraint. We present some representative numerical results for photon waveform shaping in Figs. 4.11 and 4.12.

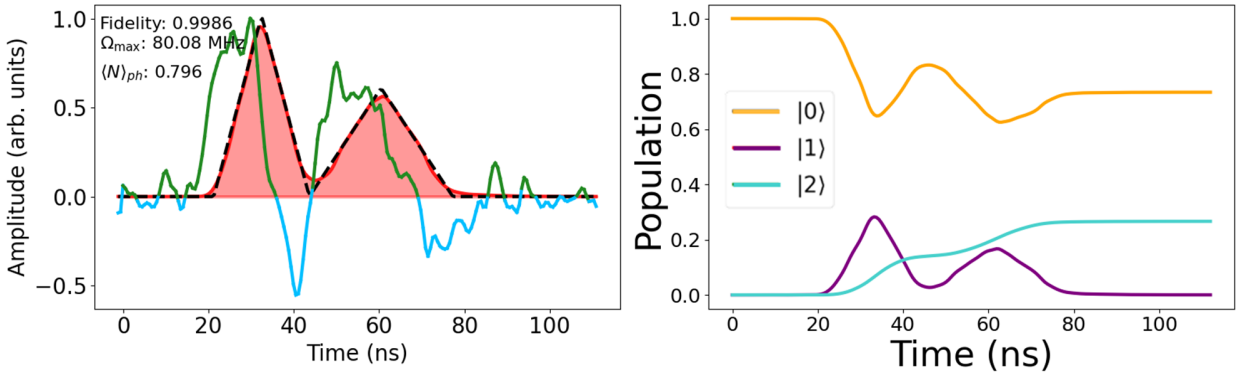

Figure 4.11: Predicted optimal laser pulse envelope including in and out-of-phase components and best photon shape to produce a double-triangular distribution .

By varying system parameters such as lifetimes and branching ratios, we can explore how predicted laser pulses must change to generate a photon of a desired distribution. As expected, we find that as radiative state lifetimes increase relative to characteristic photon time scales, the ratio of the in and out-of-phase pulse integrals becomes closer and closer to 1. This can be seen in Fig. 4.13. A summary of the effects of varying different parameters

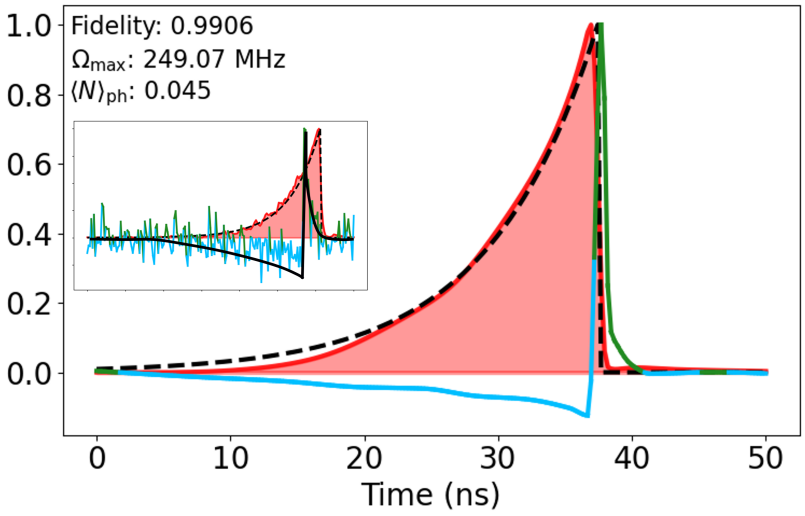

Figure 4.12: Predicted optimal laser pulse envelope including in and out-of-phase components and best photon shape to a produce an exponentially rising distribution. Inset shows final pulse envelope before manually fitting waveform shape (thick black line). Manually assigning a final envelope is performed to generate a smooth, piecewise excitation waveform. For non-discontinuous distributions, smoothing is performed via averaging or gaussian filtering. However, for solutions to generate rising exponential, this smears out discontinuities and degrades fidelity.

as well as notes on dominant error sources and limitations on photon shaping for different photon classes can be found in Tab. 4.3.

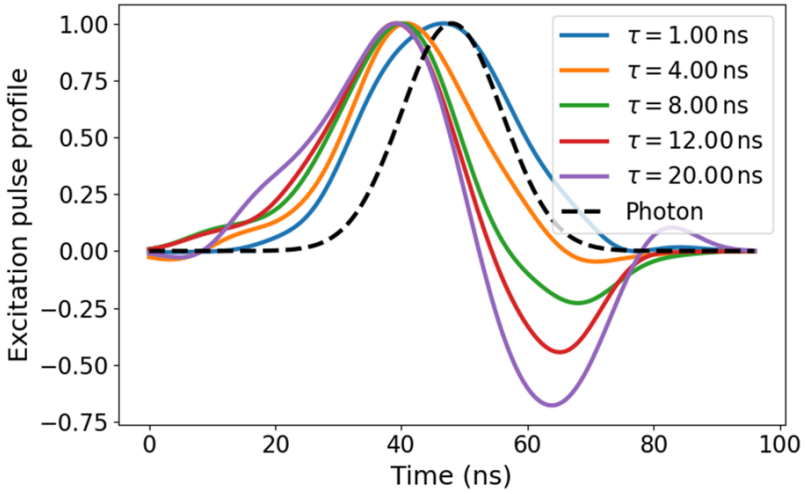

Figure 4.13: Optimized excitation pulse shapes to produce a desired gaussian photon ($\sigma = 8ns$) for emitters of a variety of lifetimes. As τ grows, the corresponding pulse shape shows earlier phase flips and greater contribution of out-of-phase contributions to $\Omega(t)$.

Photon Class	**Waveform characteristics**	**Requirements for excitation control**	**Dominant error sources**	**Notes on fidelity, $\mathcal{F}$**
Natural	$\propto e^{-\Gamma t}$	$\tau_{\text{pulse}} \ll \tau_{\text{atom}}$	–	(1) $\mathcal{F}$ limited by τ_{pulse}, pulse area.
Long	High probability density at large t, relative to natural	amplitude of $\Omega(t)$	None (for smooth shapes, $\langle N \rangle \leq 1$)	(1) $\mathcal{F}$ remains high for $\langle N \rangle \leq 1$; (2) largely insensitive to branching ratio.
Short	Low probability density at large t, relative to natural	bit-wise phase and amplitude of $\Omega(t)$	(1) Incoherent re-excitation; (2) Dephasing	(1) $\langle N \rangle \uparrow \Rightarrow \mathcal{F} \downarrow$; (2) $\tau_{\text{atom}} \uparrow \Rightarrow \mathcal{F} \uparrow$; (3) $\Gamma_{1,2}/\Gamma_{1,0} \uparrow \Rightarrow \mathcal{F} \uparrow$.
Weird	Discontinuous or non-monotonic	bit-wise phase and amplitude of $\Omega(t)$	(1) Incoherent re-excitation; (2) Dephasing	(1) Error in pulse area $\rightarrow$ over pumping and coherent re-excitation; (2) $\tau_{\text{atom}}/\tau_{\text{pulse res.}} \uparrow \Rightarrow \mathcal{F} \uparrow$; (3) $\langle N \rangle \uparrow \Rightarrow \mathcal{F} \downarrow$; (4) $\tau_{\text{atom}} \uparrow \Rightarrow \mathcal{F} \uparrow$; (5) $\Gamma_{1,2}/\Gamma_{1,0} \uparrow \Rightarrow \mathcal{F} \uparrow$.

Table 4.3: Summary of photon classes, their waveform characteristics, control requirements, dominant error sources, and fidelity dependencies.

Similarly, if we instead consider an emitter of a single lifetime but vary the target $\langle N \rangle$ for a target photon, we see the effects of incoherent re-excitation and as predicted in Sec. 4.4.2. Specifically, the presence of an exponentially decaying tail determined by state lifetime becomes more pronounced as $\langle N \rangle$ increases. One such representative set of photon shapes is shown in Fig. 4.14.

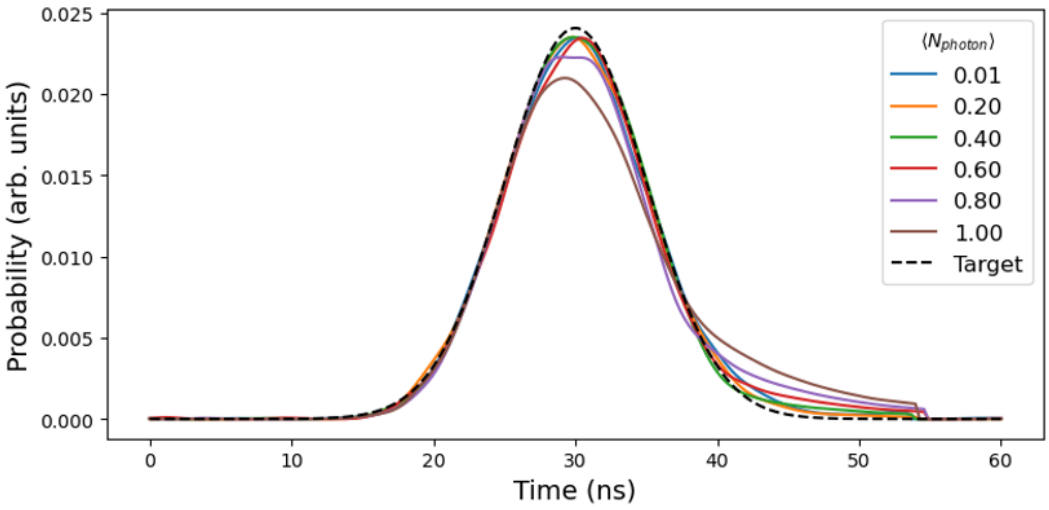

Figure 4.14: Predicted photon shape for a single target distribution as mean photon emission number, $\langle N \rangle$, is varied for an emitter with $\tau = 8.1$ns. We see that even for a short photon, dephasing due to incoherent re-excitation leads to an inability to perfectly match the desired photon temporal waveform. These effects worsen with increasing $\langle N \rangle$.

However, while the master-equation-only approach can calculate the mean photon emission number for a given excitation pulse, and we can see the effects of multiple excitations manifest as $\langle N \rangle$ increases, we cannot determine the full emission statistics corresponding relative probability of producing a n-photon event. This information is critical for optimization of actual qubit-photon entanglement protocols as multi-excitation events lead to state preparation errors. What is worse, because of non-unit photon detection efficiency, these errors may not be heralded. Consider a one-photon detection protocol where two distributed nodes are driven

by identical excitation pulses. If there is a non-zero possibility that a given excitation may produce two or more photons from one of the nodes, provided only one photon is actually detected, any multi-emission shot will be incorrectly heralded as generating entanglement. As entanglement verification cannot be done without destroying the state of the system, this will lead to the result of the entire circuit being wrong. As such, understanding how to adjust excitation parameters and predict photon statistics for an excitation process is of principal importance for network applications. To gain greater insight here, we must employ quantum trajectory theory.

4.5 Quantum Trajectory Theory approach to photon statistics

Our basic approach to determining emission statistics is to employ quantum Monte Carlo methods, as described in Sec. 3.3.1, to simulate the evolution of an excitation process many times and to track when jumps occur. Specifically, we model each trajectory as a coherent evolution of the wavefunction for our system under a time-dependent interaction Hamiltonian. For each trajectory, we allow discrete and instantaneous jumps from $|1\rangle$ to either $|0\rangle$ or $|2\rangle$.

Between jumps the (unnormalized) state $|\tilde{\psi}(t)\rangle$ propagates with the non-Hermitian effective Hamiltonian

$$H_{\text{eff}}(t) \;=\; H(t) \;-\; \frac{i\hbar}{2}\sum_q L_q^\dagger L_q \quad \Rightarrow \quad \frac{d}{dt}\,|\tilde{\psi}(t)\rangle = -\frac{i}{\hbar}\,H_{\text{eff}}(t)\,|\tilde{\psi}(t)\rangle,$$

where $L_\pi = \sqrt{\Gamma_\pi}\,|2\rangle\langle 1|$ and $L_{\sigma^+} = \sqrt{\Gamma_{\sigma^+}}\,|0\rangle\langle 1|$. The norm $\|\tilde{\psi}(t)\|^2$ gives the probability that no photon has been emitted prior to time, t.

When they occur, jumps acts as a projective measurement, leaving the state in either $|0\rangle$ or $|2\rangle$, and coherent evolution begins again from this point. In order to generate sufficiently dense histograms of times of emission, thousands or tens of thousands of such trajectories are usually modeled. We see that by examining the aggregate distribution of all jumps, we recover a photon temporal profile that matches what we predict using the master-equation-only approach.

As the excitation protocol used in our Λ system prevents re-excitation after the emission of a π photon, so long as no measurement is made of σ photons, the temporal waveform predicted by either the trajectory averages or master equation represents a pure state. However, if σ decays are also counted, the detection of a σ photon at some time, t_0, thus changes the state of the system, leading to distinct trajectories and π photons with of various temporal waveforms. Practically, as any irreversible interaction with the environment constitutes a form of measurement (whether the PMT registers a photon click or not) this means that any multi-emission process will lead to the creation of a mixed state for the resultant photon field. Therefore, it is important to be able to characterize the probability that if we detect a photon for a given excitation pulse, this photon truly came from a one-photon event, rather than being the only photon actually detected by the experimental apparatus. This is where the Monte Carlo methods become useful.

For each trajectory, we track the time at which every jump occurs. With this data, we are able to characterize the full emission statistics for an excitation process. Returning to the desired photon shown in Fig. 4.14, by examining the distribution tails we can see that even for distributions where the probability of re-excitation is suppressed by low mean photon emission numbers and a short characteristic time scale vis-à-vis the state lifetime, multi photon processes lead to significant dephasing. This limits the ability to achieve arbitrarily high temporal waveform preparation fidelities. Previously, we could only find $\langle N \rangle$, but not the relative frequency of n-photon emissions[1]. Now, by tracking jumps, we can assign relative probabilities to n-photon events. The results of tracking jumps for the corresponding excitations shown in Fig. 4.14 are recorded in Tab. 4.4. We can see that even in the case of $\langle N \rangle = 0.2$, we have some non-zero instances of 2 photon events. A quantum network user could employ the technique described above to optimize excitation parameters for a given protocol.

[1]Actually, treating $\langle N \rangle$ as the mean field number operator in the number state basis is only legitimate in the instance where the entire atom-field system exists as a pure state. Otherwise, we can not write a number state representation for the field and would instead resort to the mixed state density operator. However, as quantum trajectory theory tells us, either approach will yield the same results in this case.

$\langle N_{\mathbf{photon}} \rangle$	P(**N total** = 0)	P(**N total**) = 1	P(**N total** = 2)	P(**N total**) $\geq$ 3
0.01	99.08%	0.92%	0.00%	0.00%
0.20	80.48%	19.30%	0.22%	0.00%
0.40	62.14%	36.44%	1.38%	0.04%
0.60	44.88%	50.92%	4.02%	0.18%
0.80	27.36%	65.52%	6.82%	0.30%
1.00	10.80%	80.32%	8.30%	0.58%

Table 4.4: Monte Carlo photon event statistics showing probabilities (%) for different mean photon numbers $\langle N_{\text{photon}} \rangle$ corresponding to Fig. 4.14. P(N total) denotes the probability of trajectories where exactly

4.6 New techniques for post selection

The application of quantum trajectory methods to modeling photon emission processes turns out to have more to offer than a full accounting of the relative incidences of n-photon events alone. We can use the same data generated by recording the times of emission for many trajectories to make conditional inferences about the likelihood that a photon detected at any given time came from an n-photon event. This can allow a quantum network user to more intelligently make decisions about which shots to throw out.

If we return to the situation of a 1-photon entanglement protocol with non-unity detection described in the prior section wherein the occurrence of a multi-photon emission event at one node leads to a false-positive entanglement heralding, we can reduce the frequency at which such false-positives are accepted by applying certain thresholds on photon time of detection.

Figure4.15 illustrates this concept. Let us say that we have a single-photon remote entanglement scheme where we want to impose a constraint that if a photon is detected, the

probability that this photon came from a single photon emission event is greater than 0.90. Most real protocols would probably target lower error rates to increase Bell state preparation fidelity; however, we choose this threshold for clarity in plotting photon distributions in Fig. 4.15. Prior, the best we could do was to calculate the corresponding ratio of multi-photon events for a given $\langle N \rangle$. However, we have more information available to us than just this ratio. We can also use this trajectory data to determine correlations between time of arrivals in multi photon events.

The intuition for our approach is straightforward. Consider the case where the qubit at each node has the structure of our Λ system where in the emission of a photon at a fixed polarization (let say π for consistency) projects the atomic system into a dark state. For such a node, if a π photon is detected at very early time in the excitation period, we expect the probability that this photon came from a 2-photon process to be much lower than if we detect the π photon very late in the pulse duration.

Now, using trajectory data, we can quantify that intuition by looking at the ratio of n-photon events across the entire range of time bins. We can find the time after which a photon which is detected has is more than X% likely to be drawn from a multi-photon emission event and throw out all shots where the time of detection is after this cutoff.

One subtlety to note here is that the act of throwing out certain counts changes the temporal distribution of photons detected during 'good' counts. Returning to the trajectory picture, this makes sense as we now only consider a subset of trajectories with an upper time limit on photon emission. As such, the associated temporal waveform for this set will be given by the ensemble average of trajectories which satisfy the time-of-arrival condition. As such, care must be taken to ensure the conditioned temporal waveform from both nodes remains identical. For same-qubit type systems this is trivially achieved by applying symmetric thresholds. However, for hybrid systems, this requires distinct thresholds to ensure the resultant distributions continue to match. Distributions for first and second emitted photons, as well as the distribution of first photons conditioned on single and multiple emission events are shown in Fig. 4.16, illustrating how the temporal profile of first photons is affected by

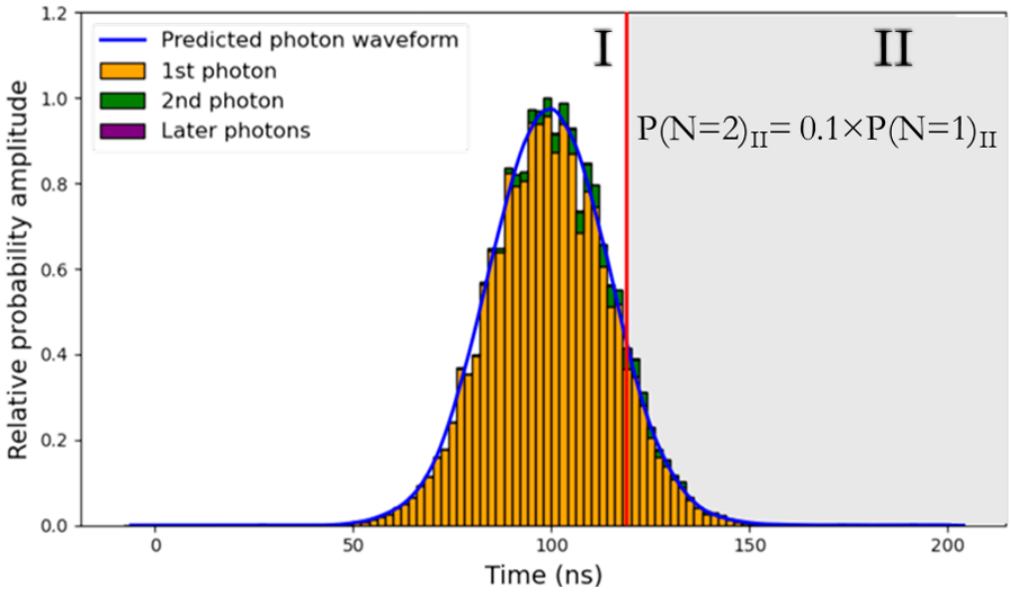

Figure 4.15: Illustration of time-based post-selection. Histogram shows simulated click-time densities for single-photon shots ($p_1(t)$, orange) and multi-photon shots ($p_2(t)$, green; $p_{\geq 3}(t)$, purple), normalized to unit area over clicks. The blue line shows the expected temporal distribution for the same process based on solution of the system master equation. The vertical line marks a threshold for shots where the probability that a photon detection came from a multi-photon emission event is greater than 0.10, $P_{N\geq 2} = 0.10$. This is used to established a cut time $t^\star$. Accepting only shots in region I (ie. $t \leq t^\star$) raises the fraction of true single-photon heralds while reducing the acceptance rate.

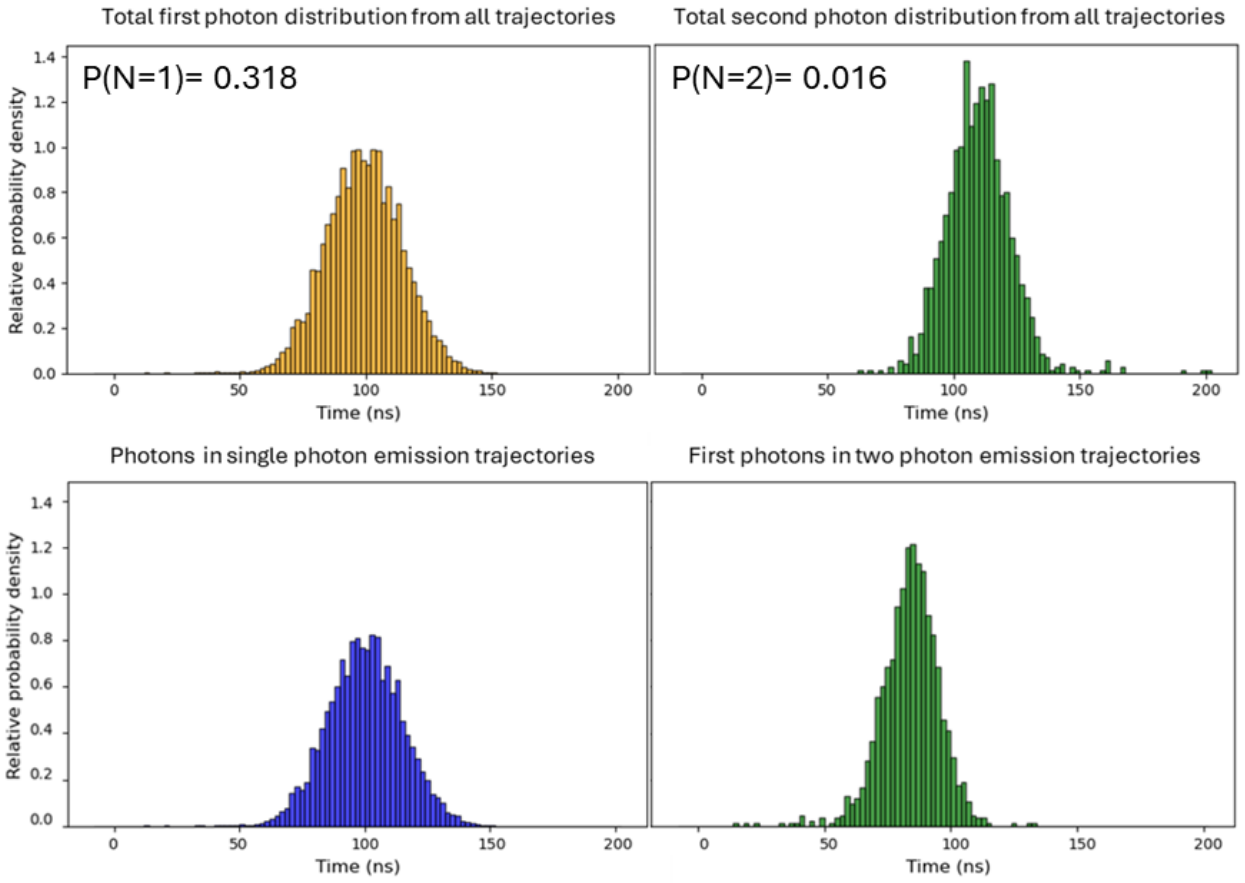

Figure 4.16: Comparison of photon distributions for one and two photon emission events for Fig. 4.15. **Clockwise from top left:** normalized distribution of all first photon times-of-arrivals; normalized distribution of all second photon times-of-arrival; normalized distribution of all first photon times-of-arrivals conditioned on the emission of only one photon; normalized distribution of all first photon times-of-arrivals conditioned on the emission of two photons. We observed expected behavior for single and multi-photon events including: photon anti bunching (seen by comparing top right and bottom right panes) and the a tighter, earlier photon distribution arising from two photon emission events vis-à-vis those coming from single photon events (seen by comparing bottom left and bottom right panes).

controlling for total number of emissions.

Between the additional emission ratio statistics derived in the prior section and the ability to identify thresholds suitable for suppressing errors from multi-photon events, quantum Monte Carlo methods have allowed us to develop two novel tools for optimal photon control. However, we have now gone as far as we can with simulation alone. The next chapter deals with how to apply the temporal shaping to the production of actual photons from a trapped ion node.

Chapter 5

GENERATING SHAPED PHOTONS FROM AN ION NODE

The heart of any experimentalists thesis must necessarily be the application of theory to a real physical system. To that end, this chapter describes efforts to realize arbitrary photon shaping in our trapped ion system. As we have already detailed the construction of the ion trap and support infrastructure in prior chapters, the contents contained here are somewhat narrower in scope. Specifically, we aim to demonstrate the ability to modulate an excitation pulse driving the $|^2\mathrm{S}_{1/2}\ m_J = -1/2\rangle \Rightarrow |^2\mathrm{P}_{1/2}\ m_J = +1/2\rangle$ transition in $^{174}\mathrm{Yb}^+$ to control the temporal distribution of photons produced during π decay to the $|^2\mathrm{S}_{1/2}\ m_J = +1/2\rangle$ state. We have already described the pertinent theory to model the time evolution of this system and optimization methods to predict the ideal form of an applied laser pulse to generate a photon of a desired distribution in Ch. 3 and 4. respectively. Instead, this chapter deals with how we can experimentally synthesize matching optical pulses. We also describe synchronized state preparation, pulsed excitation and detection to reconstruct the temporal waveform of photons scattered during this process.

We describe the implementation of and experimental limits for the generation of photons with arbitrary temporal waveforms. We close this chapter with some prospects for future work and applications of the techniques established to ongoing fields of research in quantum networking.

5.1 Overview of the photon shaping and measurement apparatus

In order to implement our photon control experiment, we need to expand the trapping infrastructure described in Ch. 2. In broad categories this includes additional laser optics, quantization axis steering, control electronics, photon collection optics and experimental

timing equipment. A schematic representation of the full experimental set up is shown in Fig. 5.1. A brief sketch of the new apparatus follows below with greater detail in subsequent sections.

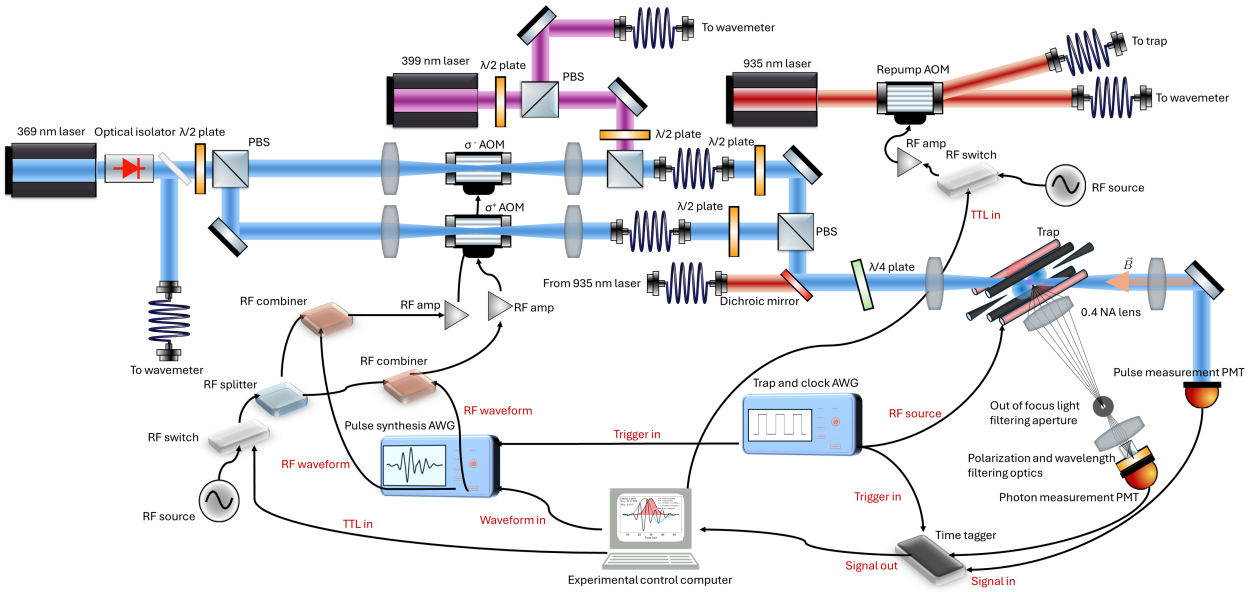

Figure 5.1: Experimental design for photon control from a trapped ion node. Implementation of photon temporal waveform shaping requires the use of two circularly polarized 369 nm beams to drive optical pumping and pulsed excitation. Each beam is focused to a narrow waist in an AOM which acts as a fast pulse shaper. This is accomplished by sending an RF waveform from an AWG to each AOM. In order to synchronize the experimental duty cycle with the ion trap, a second AWG provides both the trap RF and the experimental clock such that laser pulses and time-tagged photon detection are locked to the ion. Additional optics paths for ionization and repump lasers, as well as control electronics are also shown.

For photon shaping, we now need the addition of a second 369 nm beam to drive alternate optical pumping and excitation pulses. These beams must be of opposite circular polarizations such that, with the correct orientation of the quantization axis, we drive only σ^+ and σ^-

transitions. To set the quantization axis, we use three magnetic field coils to allow an arbitrary $\vec{B}$ pointing at the ion location. To shape the two σ^+ and σ^- beams to match the desired pulse, we tightly focus each beam through an acousto-optic modulator (AOM) that is driven by a variable amplitude and phase RF waveform supplied by an arbitrary waveform generator (AWG). The AWG is controlled by computer via a graphical user interface (GUI) we developed to run the experiment. The virtual interface for this GUI is shown in Fig. 5.2

As ion motion during a trap period will lead to a phase dependent instantaneous Doppler shift in the ion frame, we synchronize the experimental duty cycle to the trap. This is done by triggering the pulse shaping AWG using the trap RF source.

In order to select a single polarization to reconstruct the photon temporal waveform, π decay is mapped to a linear lab-frame representation and a Glann Thompson polarizer is added to the collection optics directly before the camera/PMT. The identification of the correct polarizer axis is accomplished by measuring the angular dependence of detected photon flux as twice as many σ photons are produced compared to π. Photon histograms are assembled by sending the output of the PMT to a fast time resolving pulse counter (the time tagger), also triggered by the trap RF source.

With the full infrastructure in place, the experimental duty cycle is as follows:

1. Single, or at times several, ions are trapped and cooled until crystals form. Doppler cooling is performed by applying a constant amplitude RF waveform to each AOM.

2. Next, target pulse waveforms are uploaded to the AWG, the Doppler cooling source is turned off and repeated optical pumping followed by excitation pulses are applied and the resultant photon time-of-arrivals are recorded.

3. The initial photon histogram is reviewed, and appropriate delays are added to both channels to ensure good temporal separation between pump and excitation pulses. Additional alignment of excitation pulses relative to the RF period is performed to minimize the effects of instantaneous Doppler shift due to ion motion.

Figure 5.2: Graphical User Interfaced designed for AWG control in photon shaping experiment. We built this GUI to allow easy experimental control of applied waveforms. A csv is selected which represents the envelope for a desired pulse. These envelopes are generated through the numerical simulations described in Ch 4. The user can than adjust a variety of parameters to set the correct optical power (via AWG voltage) and delay times for proper synchronization between pumping end excitation and the trap frequency. An optional power correction function can be used to account for system gain nonlinearities by apply an inverse mapping for a provided gain curve. Additional phase flips can be added at user specified times.

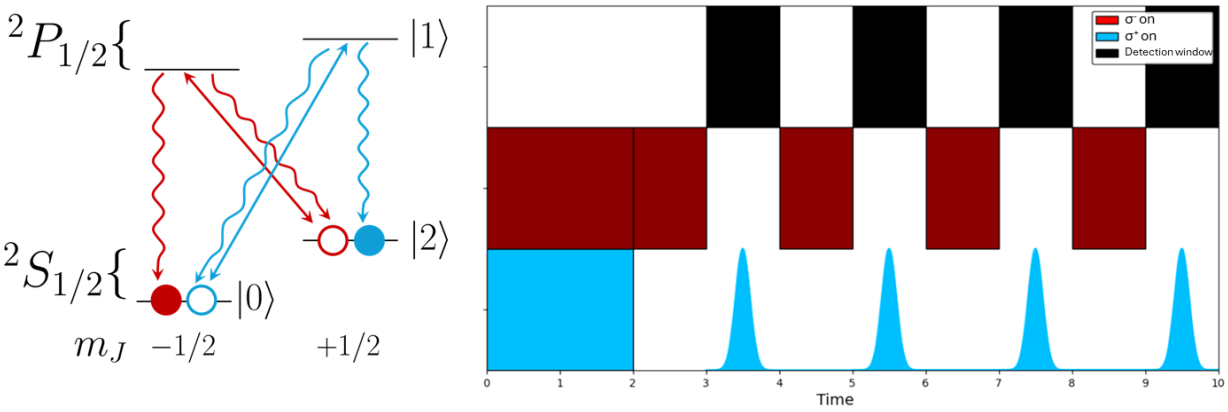

Figure 5.3: Experimental duty cycle for control of photon temporal waveforms. **Left:** optical pumping and pulsed excitation transitions shown in red and blue respectively. Correspondingly colored open circles represent population at the start of each operation while filled circles show population at after completion. Depending on the amplitude and duration of each operation, pumping and excitation may not be complete. **Right:** representation of alternating pulses plus detection window in time.

4. After a sufficiently dense histogram can be reconstructed, 935 nm repump light is turned off so that ion scatter is reduced to virtually zero. A background is collected.

5. The background subtracted photon waveform can then be generated and compared to the target photon shape.

6. A second PMT is placed at the exit port of the vacuum chamber. This is used to measure the applied optical pulses.

After this sequence, the process may begin again for a new target photon shape. A representation of the transitions pulse sequence for our duty cycle is displayed in Fig. 5.3.

5.2 Pulse synthesis

In order to generate temporally shaped pulses, we must be able to modulate the amplitude of the 369 nm laser at speeds faster than ∼10 ns. The rationale for this time scale is that we must be able to change the instantaneous Rabi frequency seen by the ion at least as fast as the characteristic time scale of the atomic response, determined by the 8 ns lifetime. Generating such pulses from a constant wave laser is not trivial. Fast control via the use of a high-bandwidth, high central frequency electro optical modulator (EOM) could be possible. However, such devices are relatively expensive and introduce additional complexity as they act by adding two frequency sidebands above and below a carrier with which they co-propagate.

As such, we opt to use AOMs powered by a time variable RF amplitude. While the use of AOMs for pulsed atomic excitations is not uncommon [81], typical applications only consider the pulse integral which determines the total excitation probability. However, as we need to control the gradient of these pulses to deterministically control $\rho_{11}(t)$, our approach must be somewhat more sophisticated.

5.2.1 *Effects of electronic and optical systems on pulse shape: AOM shuttering, AWG waveforms and interference*

A typical AOM used in atomic physics experiments consists of tellerium dioxide or quartz crystal connected to a piezo electric transducer. The principle of operation is to apply an alternating voltage to the transducer, generating small, cyclical expansions and contractions. These vibrations excite phonon modes within the crystal. As the speed of sound in such materials is on the order of 5000 meters per second, achieving shutter speeds comparable to the characteristic atomic response time scale requires beams focused down to $\approx 10 - 30$ microns at the AOM diffraction plane. The optical beam interacts with the sound wave within the crystal and some component of the beam is diffracted. After the AOM, the output beam consists of the non-diffracted (0 order) mode and one or more higher or lower order beams. These beams correspond to the addition or subtraction of energy at the acoustic frequency or one of its higher order harmonics and have an angular separation from the 0-order determined by the order and Bragg angle, set by crystal lattice properties. As such, each order picks up a frequency shift of $n \times f_{RF}$ and can be spatially resolved from the carrier and other orders. The amplitude of diffraction depends on the strength of the acoustic excitation. Thus, by varying the amplitude of the RF-carrier, pulses with a specified intensity profile can be created. Each diffracted beam also picks up an additional phase set by the RF-carrier. For our photon shaping experiment, we thus modulate the amplitude of the instantaneous Rabi frequency by adjusting the scale of the RF waveform supplied to the AOM. We also attempt to test phase modulation by advancing the RF signal by π radians part way through the applied RF pulse.

The strength of phonon-laser interaction depends on how well the spatial modes of both waves are matched. Most AOMs specify the use of laser beam diameters on the order of $\sim$100-200 microns which correspond to shuttering times of $\approx$100 ns. To achieve shorter shuttering times, we instead use a tightly focused beam which results in a moderate decrease in diffraction efficiency. At this speed, we also approach the fundamental limit set by the RF

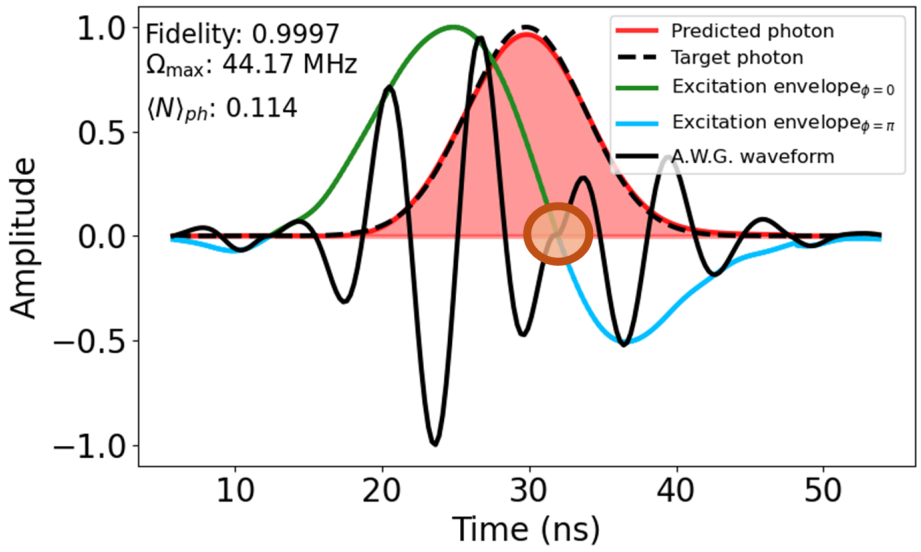

Figure 5.4: Numerical results for short photon optimization showing laser pulse envelope and associated AWG waveform with π change at 0-crossings (circled region).

carrier frequency which is set to 130 MHz. For a single full wave excitation, the minimum pulse time would be set by $1/f_{RF} \approx 8$ns. Higher central frequency AOMs exist, or multipass designs could be used to increase this speed limit. However, with current control electronics we would not be able to support a faster AOM as we are bandwidth limited by the AWG we use to drive the AOM.

In order to reduce any leakage from other AOM orders, we use an aperture at the output of the device to mask light not from the intended order. Without these apertures, we can see the effect of optical interference at the RF frequency by inspecting optical pulses as shown in Fig. 5.5. Figure 5.6 shows a conceptual picture of how misalignment to fibers can induce pulse shape deformation as well as a measured RF vs optical power curve which can lead to optical pulse nonlinearities as compared to the applied RF waveform.

To ensure a consistent single spatial mode at the trap, output beams from the AOM are coupled to a single mode fiber before being sent to the ion. We find that alignment to the

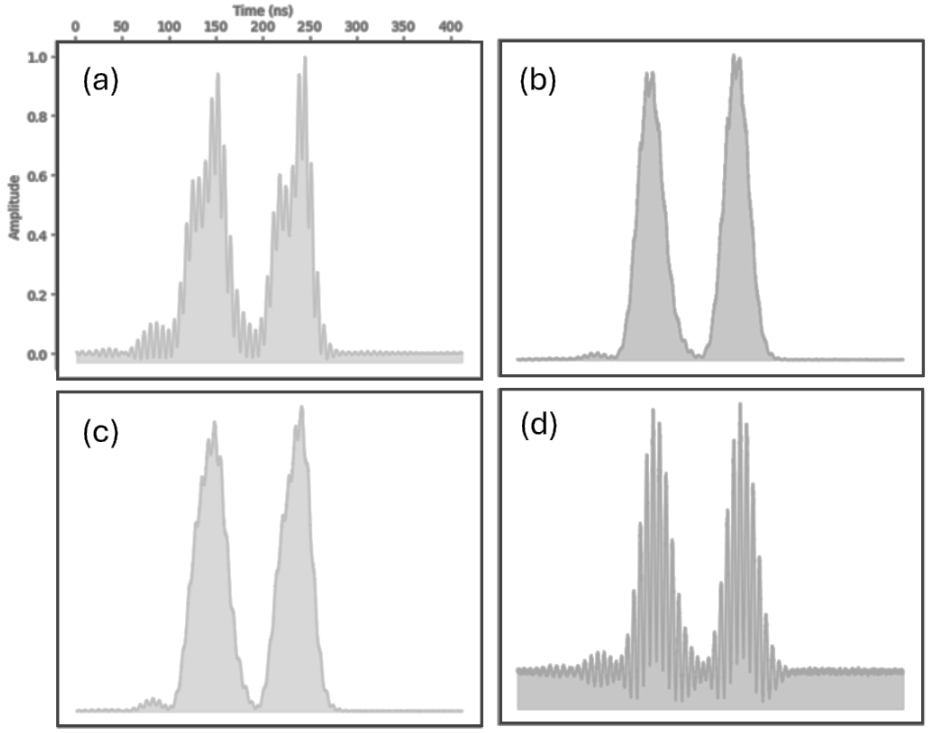

Figure 5.5: Effects of optical interference and fiber misalignment on measured pulse shapes. A double gaussian AWG envelope is used to generate the measured pulses. **(a)** no aperture is present, leading to beating of the optical pulse at the AOM frequency due to coupling of multiple diffracted orders into the single mode fiber. **(b)** the same pulse with the addition of an aperture showing a redution in interference. **(c)** and **(d)** the same pulse at two slightly different angles of incidence to the single mode fiber.

fiber also leads to similar pulse distortion effects. Because any change to the RF waveform takes time to propagate across the beam, the RF pulses become 'smeared' out in the resultant optical pulse with a spatial/temporal shear. Differences in coupling efficiency across the beam width at the fiber input also contribute to pulse distortion. Mapping of the spatial profile of the beam before the fiber to the wavefront after is a challenging problem which cannot be properly accounted for in the context of an optics assembly without active feedback or full mode propagation modeling. The aggregate effect of these contributions is to degrade effective optical pulse synthesis accuracy with any feature in the RF waveform being averaged over some short (on the order of several nanoseconds) time scale at the ion.

A variety of electronic effects can also reduce how well a generated optical pulse matches a programmed RF waveform envelope. The AOM diffraction efficiency and the RF amplifier used after the AWG both have nonlinear responses to applied RF power. A curve of optical power versus RF amplitude from the AWG can be measured and used to correct this nonlinearity in a steady state context.

However, larger issues arise from limits of the AWG used. Our device has a sample rate of 1.2GaS/s and a bandwidth of 80 MHz. As we run our AWG at a carrier frequency of 130 MHz, this pushes the device outside of its reliable operating limits. Taken all together, this leads to waveforms showing a wide range of distortions. The relative amplitude of these distortions depends heavily on the particular waveform characteristics and as such cannot be generally accounted or corrected for by simply applying a single mapping function that captures all system nonlinearities.

Instead, we take advantage of the existing optimization infrastructure built to support identification of optimal pulses numerically and adjust system inputs and outputs to allow us to use the same algorithms to empirically optimize pulse shapes. The results and limitations of feedback based synthesis are described in the next section.

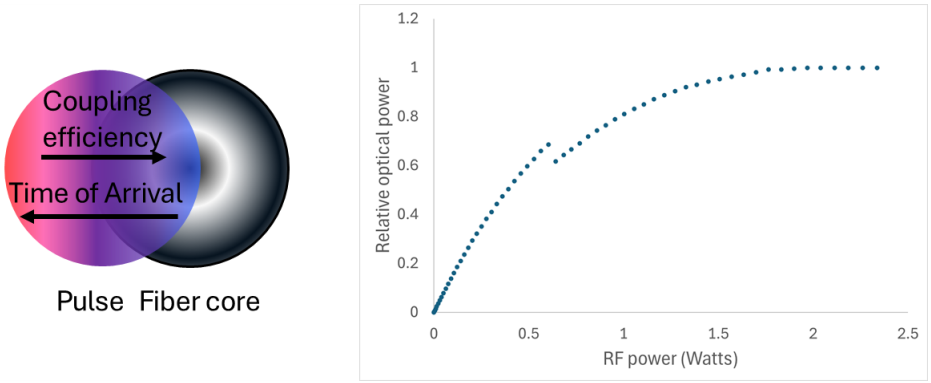

Figure 5.6: Error sources leading to a mismatch between target and produced optical pulses include **Left:** schematic representation of mapping acoustic wave propagation across laser wavefront leading to 'smeared' pulse shapes. Different coupling efficiency across wavefront at fiber further contributes. **Right:** nonlinear response in RF amplifier and AOM diffraction efficiency leads to flattening gain curve. We aim to operate in range well approximated by a linear fit. Alternatively, the power curve can be mapped back to AWG waveform.

5.3 Optimizing laser pulses

Our approach to optimizing laser pulses is conceptually straightforward. We have already implemented suitable waveform optimization algorithms developed for numerical prediction of target laser pulse profiles to generate a given photon. By adjusting the inputs and outputs of these algorithms to communicate with experimental hardware, we can use the same approach to measure laser pulses and update AWG waveforms until suitably high fidelity excitation pulses are synthesized. The principal experimental challenge becomes one of developing the necessary infrastructure to accomplish this.

5.3.1 Pulse measurement and feedback process

We modify the work flow of the pulse optimization loop to do the following:

1. The target distribution is updated to allow us to specify a desired laser envelope.

2. The desired laser envelope is converted to an AWG readable format and and sent to the AWG.

3. The AWG is triggered in tandem with a PMT which measures the optical pulse generated. Many such shots are used to construct the laser pulse histogram.

4. The measured pulse histogram is interpolated and compared to the target distribution. The optimization algorithm of choice is used to update the AWG waveform and the measurement process is repeated.

5. Feedback is continued until a user specified fidelity, defined by the cosine similarity of the target and desired pulse amplitudes, is achieved.

As compared to the numerical optimization of the interaction Hamiltonian, experimentally synthesizing high fidelity lasers pulses presents an additional problem. Our feedback signal is now only the intensity of the laser light at each time bin. We cannot directly verify the relative optical phase of the light, needed for coherent de-excitation. Although, an interferometer could theoretically be built to verify how changes in RF phase map to the optical phase, implementing the necessary infrastructure into our experimental feedback was not a viable option. Instead, for pulses which rely on phase modulation, we opted to optimize the absolute value of the pulse envelope and apply a post hoc π phase change at times when the desired Rabi frequency switches signs.

We find that this optimization method allows the preparation of arbitrarily shaped laser pulses with fidelities greater than 0.999 for smooth, mono-modal distributions. Distributions with multiple local maxima or discontinuities proved harder to synthesize. However, fidelities greater than 0.99 were still readily achievable. Some representative results for pulses with no phase modulation are shown in Fig. 5.7 while an illustrative example including post hoc phase change for negative Rabi frequencies is shown in Fig. 5.8.

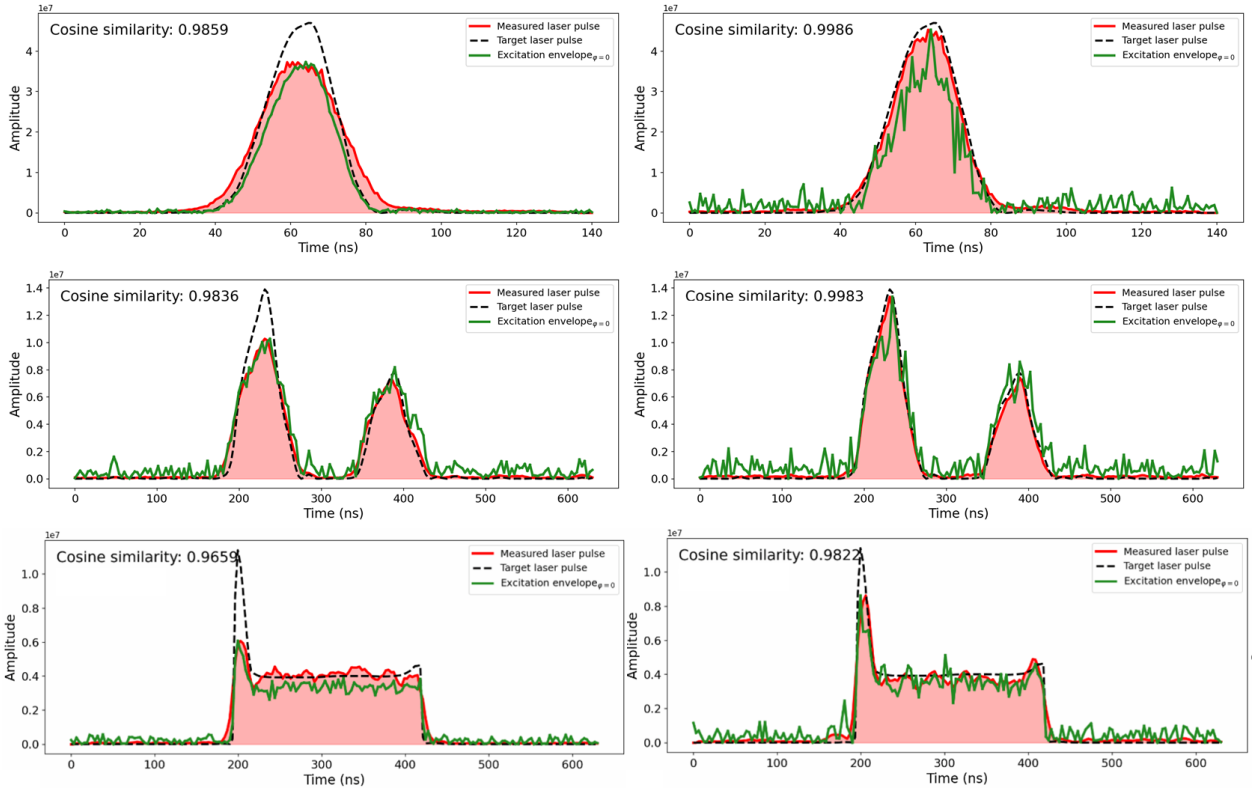

Figure 5.7: Results of laser pulse amplitude optimization for three representative waveforms. For each pair, the left hand side represents the initial pulse generated by the using the target pulse shape as the AWG envelope. The right hand side shows the best measured pulse after optimization. The excitation envelope shown in green corresponds the AWG waveform before being multiplied by the corresponding RF frequency and being sent to the AOM. Smooth distributions show relatively low distortions and as such the fidelity of the initial pulse is much higher. Distributions with discontinuities result in much worse initial waveforms. Through optimization, we are able to reduce pulse error by $\sim$7.5 and $\sim$30 for the case of smooth, mono-modal and discontinuous distributions respectively.

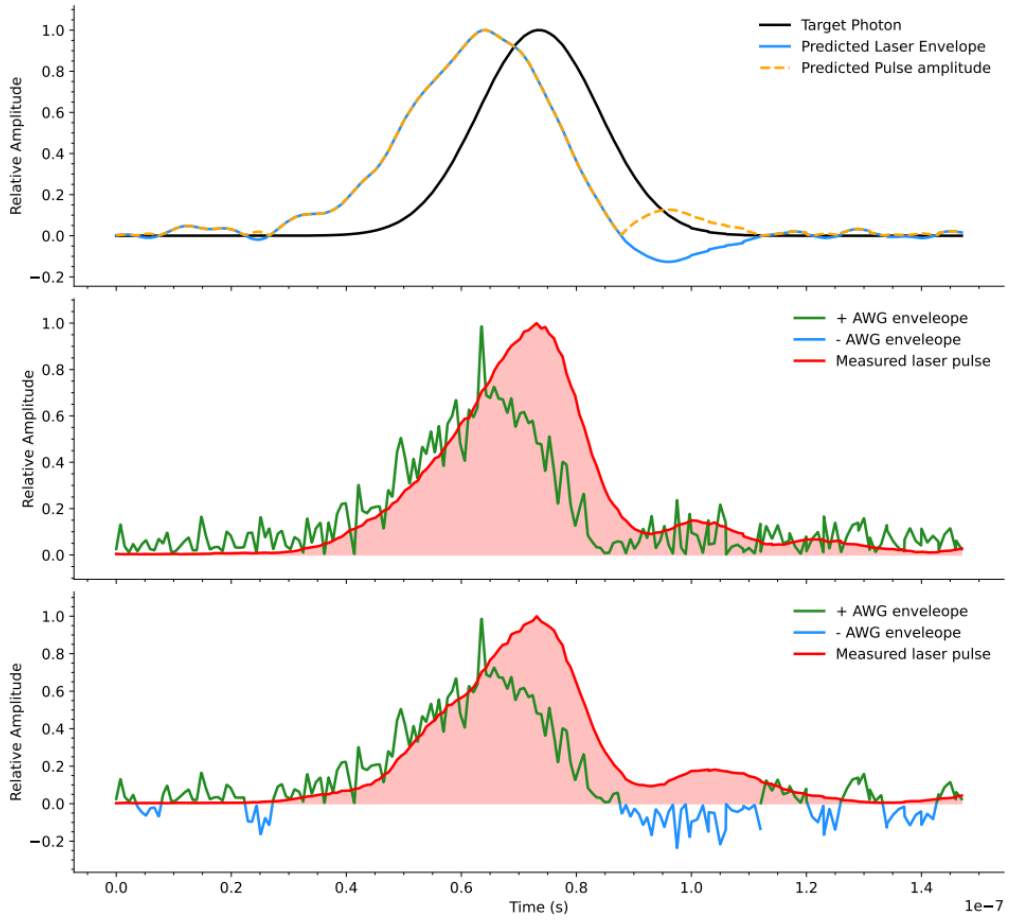

Figure 5.8: Results of laser pulse optimization for a representative waveform with optical phase modulation. **Top:** the predicted laser envelope and expected pulse amplitude needed to produce a target photon waveform. **Middle:** optimization of the AWG waveform to produce an optical pulse with the target amplitude. **Bottom:** The addition of out of phase components in the RG waveform produced by the AWG also leads to slight changes in measured pulse amplitude.

5.3.2 Effects of phase modulation and f_{RF} on optical pulses

While flipping the phase of the RF waveform used to generate optical pulses should, in principle, not effect the measured intensity of the resultant optical pulses, experimentally, we find that phase-flips do lead to some distortion in measured waveform. Specifically, we see that when a target envelope is synthesized for a pulse which has some out of phase component, upon the final addition of phase flips, the optical pulse is observed to decrease in amplitude more quickly than for the no-flip applied case. We attribute this to the out-of-phase RF causing the acoustic wave in the AOM to be extinguished through destructive interference more quickly than would otherwise occur based on device rise/fall times. At this stage, we begin to see indications that experimental realization of phase modulation may not be readily achievable with our experimental hardware. As we only aim to apply phase changes at zero crossings where the sign of the Rabi frequency flips, the measured amplitude of the laser pulse would ideally go to zero part way through the pulse duration. However, we are not able to achieve this even after optimization. We attribute this to the finite rise and fall times of the AOM and AWG primarily.

In our model, the finite RF frequency is ignored for intensity modulation as this only sets the carrier. Amplitude modulation is determined by varying the envelope of the waveform. The RF frequency should not, ideally, manifest in the features we see for intensity of the optical pulse. The only caveat is that to generate smooth pulse shapes, we need waveforms to be longer than at least a single full RF pulse period. In other words, even with sufficiently fast electronics, the shortest optical pulse we could produce must be at least $1/f_{RF}$. For a carrier of 130 MHz, this corresponds to about 8 ns in duration.

5.4 Measuring single photons

Now that we have established

1. How optimal laser pulses shapes can be predicted to produce a desired photon,

2. How such pulses can be experimentally synthesized using feedback and optimization techniques, and

3. How to construct the proper infrastructure to synchronize experimental shots

we are ready to implement our approach by measuring single photon waveforms.

Because ions experience time dependent forces synchronized with the trap drive as a result of micromotion, in order to generate the most accurate estimate possible of single photon temporal distributions we use trigger excitation pulses in tandem with the trap RF. If this is not done, we will instead be measuring the aggregate photon distribution generated by a mixed state of ions with many different time dependent detunings. In reality, even with synchronization to the trap, the photon histograms we construct will still have contributions from an admixture of different ion states and thus not perfectly represent a true pure state measurement. There are two main reasons for this.

First, as all lasers have a finite linewidth, this will introduce phase noise over the course of an excitation which will not be identical shot to shot. The free running linewidth of these devices is quite narrow ($\sim$100 kHz), especially compared to the corresponding linewidth and as such the effects of this environmental dephasing will be very minimal. Of more concern, laser frequency may drift on the order of 10 MHz during the several minutes required to reconstruct a photon waveform. The simplest approach of feeding back on laser measured laser frequency using the wavemeter reading as an error signal is limited by the resolution of the device which is 10 MHz. Additionally, other factors such as collisions with background gases may lead to different atomic states during the process of generating of photon waveform measurement.

Finally, we operate above the Lamb-Dicke threshold as discussed in Ch. 2. Therefore, photon absorption and emissions is likely change the ions motional state from shot to shot. By increasing or decreasing the ions velocity due to secular motion, these events will lead to slightly different instantaneous Doppler shifts of the laser frequency in the ion frame. We can estimate these effects by considering trap parameters. The radial and axial secular

frequencies are on the order of 1 MHz and 200 KHz with corresponding mean motional mode occupation numbers of 10 and 50 respectively for our trap as discussed in Ch. 2. We calculate the corresponding energy as

$$E \approx \hbar\omega(\bar{n} + 1/2) \tag{5.1}$$

and equivalent maximum velocity as

$$v_{\max} = \sqrt{2E/m} \tag{5.2}$$

where m is the mass of our ion. This yields a velocity of only ≈ 0.1m/s and corresponding maximum instantaneous Doppler shifts of $\approx$0.3 MHz. Plugging this detuning into Eqn. 2.15 and assuming we operate at the saturation intensity, this will lead to a modulation of the atomic scattering rate of less than 0.001%.

5.4.1 Characterizing micromotion

Fast effects due to micromotion pose a more significant problem, thus necessitating synchronization of shots to the trap frequency. While locking the timing of excitation and measurement to micromotion does ensure shot-to-shot consistency, this still means that within a given shot ions will effectively be subjected to a chirped pulse due to the instantaneous Doppler shift associated with motion due to the RF drive. As detuning from the transition frequency leads to a modulation in scattering rate, the effects of will be to degrade the similarity between the target photon and measured photon as laser pulse solutions are generated with an assumption of zero time dependent detuning.

In order to characterize the amplitude of micromotion, we first measure the ion scattering rate as a function of phase during a single RF period while Doppler cooling. This can be accomplished by sending the PMT output to the time tagger, triggered using the same AWG that supplies trap RF. The results of a representative scan are shown in Fig. 5.9. Ideally, with zero ion motion, the resultant RF synchronized scan would show a constant amplitude scattering rate. Instead, we seem a high contrast with ion florescence varying by a factor of $\sim$25 throughout a single trap period. We can estimate the maximum instantaneous detuning

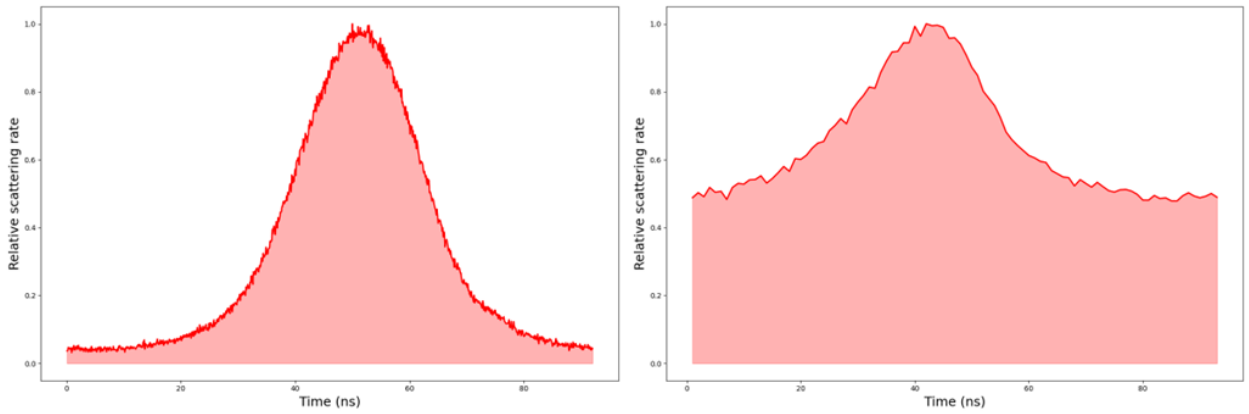

Figure 5.9: Photon scattering rate as a function of phase in RF period before (**left**) and after (**right**) partial micromotion compensation. Ion florescence was calibrated to the phase of the RF period to characterize the effects of micromotion. We measure a nearly 25× difference in scatter rate over the course of one period before compensation. After adjusting electrode voltages, this contrast was reduced to 2×. Our trap lacks adequate electrodes to fully compensate micromotion.

caused by excessive micromotion by referencing Eqn. 2.15. We operate close to saturation intensity so we can write

$$\frac{R(\Delta)}{R(0)} = \frac{1+s}{1+s+(\frac{2\Delta}{\Gamma})^2}. \tag{5.3}$$

At saturation intensity, this yields $\Delta = \Gamma/2\sqrt{24(1+s)} \approx 68$MHz.

This detuning arises from the instantaneous Doppler shift thus we can calculate the corresponding velocity $\nu_{mm}(t) = v(t)/\lambda$. Plugging in our detuning and laser wavelength, this yields a max velocity of ∼28 m/s. Assuming sinusoidal motion, we can estimate the total spatial spread of our ion due to micromotion as $1/2 \times (28m/s) \times (1/f_{RF}) \sim 1.3\mu m$. This is approximately in line with minimum resolvable ion spreads using our EMCCD camera.

We attribute this primarily to misalignment of end cap electrodes. As end caps are not perfectly aligned the center of the four rods which generate the RF potential, these electrodes act to push the ion off of the RF null. As our trap does not have dedicated compensation electrodes, we are not able to fully minimize this micromotion. In principle, the micromotion induced detuning could be included in numerical simulations, and laser pulse shapes could be adjusted accordingly. However, because these effects rely on specific trapping details such as electrode voltages which are regularly adjusted, accounting for these effects apriori during pulse design is not realistic. Instead, we can minimize effects by aligning timing so that the center of the excitation pulse is applied when instantaneous velocity due to micromotion is zero. Experimentally, this is accomplished by adjusting the delay time at which our excitation pulse is generated after being triggered by the trap synchronized AWG. During single photon collection, the delay is repeatedly increased until the the maximum scattering rate is measured, indicating the excitation pulse is well aligned with the trap. While we do not consider the effects of micromotion induced time dependent detuning during initial pulse design, we do seek to match measured photon characteristics with these effects in Sec 5.5.

5.4.2 Optical pumping and pulse timing

The final steps before we are ready to measure the temporal distribution of photons is to implement the necessary elements for state preparation and to synchronize state preparation and pulse excitation. State preparation is performed via optical pumping using circularly polarized light with a quantization axis parallel to the laser propagation vector. Referring to the experimental layout in Fig. 5.1, we first combine both 369 nm laser beam path on a high extinction-ratio polarizing beam splitter and then send both beams through a quarter wave plate to achieve appropriate polarized beams of opposite handedness. As the quarter wave plate is designed for 355 nm and thus has a retardance between the fast and slow axes of slightly less than 90° for 369 nm light at normal angle of incidence, we slightly adjust the orientation of the wave plate. We select the correct rotation and angle of the wave plate by using a second PBS to ensure that both 369 nm beams are optimally circularly polarized by balancing the output powers on each port of the PBS for each beam separately. We then optimize optical pumping efficiency by blocking either the left or right circularly polarized path and adjusting the current of three magnetic field coils placed to provide the ability to set an arbitrary direction for the quantization axis. We adjust these coils until ion florescence is minimized when only a single 369 nm polarization is applied.

We can then calibrate optical pumping efficiency by applying alternating pulses of σ^+ and σ^- polarized light and observing the rate of photons scattered by the ion over these pulse durations by synchronizing the AWG and photon detection. We fit an exponential decay to the measured scattering rate curve allowing us to estimate the time to pump a target fraction of the population into $|0\rangle$. The results of this fit are shown in Fig. 5.10. Considering the effects of micromotion discussed prior, for perfect polarizations of pumping light, we would expect an exponentially decaying scattering rate modulated by the RF-phase synchronized signal. However, we see that while the max scattering rate for each period in the optical pumping curve initially decreases exponentially, pumping efficiency is reduced at later times in the scan. While we have not definitively identified the explanation for this

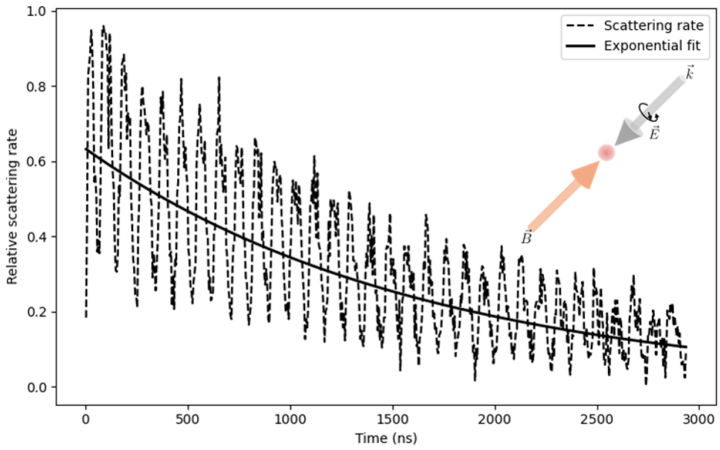

Figure 5.10: Optical pumping for initial state preparation. We exponentially fit a curve with characteristic decay time of $\tau = 1580$ ns to our measured scattering rate.

observation, we suspect stress induced birefringence on the entrance view port diminishing the polarization purity of the laser may be responsible. While pumping efficiency may suffer from such experimental effects, we benefit from an experimental scheme where we target lower emission probabilities for shaped photons such that pump strength is higher than the subsequent excitation.

With pumping implemented and laser pulses synthesized, the last experimental detail to address is to ensure correct temporal alignment of optical pumping and shaped excitation pulses. We do this by adjusting the relative delay between pulses through the GUI interface and measuring the time of arrival of the resultant laser pulses on a second PMT which is located on a path after the laser exit viewports. A scan of aligned pulses is shown in Fig. 5.11.

5.5 Photon shaping results

We are finally ready to apply a series of excitation pulses and to measure the resultant photon distributions. Ultimately, due to the large effects of ion micromotion in our trap, we first apply alternating pump-pulse cycles and observe the time-of-detection histogram for target

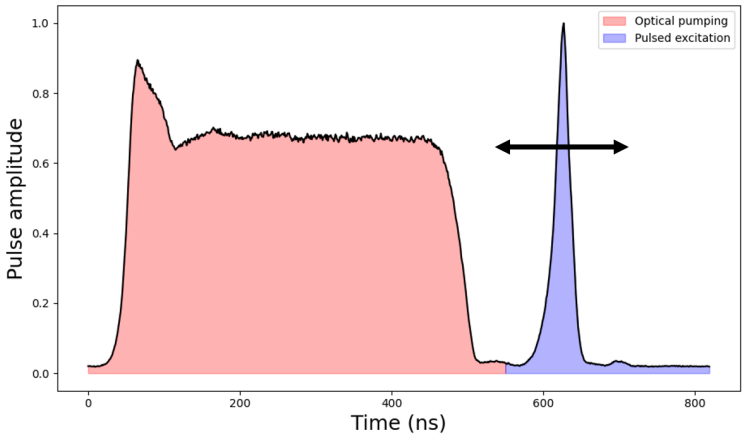

Figure 5.11: Measured laser pulse shapes used for pumping-excitation temporal alignment. We adjust the delays between the two pulses to ensure that excitation and pumping do not overlap in time.

photon shapes and adjust the delay of pulses in order to align excitation to the phase of the trap period which minimizes measured photon distortions vis-à-vis the target distribution.

For some target distributions, we also detune our laser frequency from the central transition frequency. We do this for two reasons. First, collecting sufficiently high numbers of photons to create dense histograms takes several minutes. If we do not sufficiently detune from resonance, this leads to heating of the ion over the course of a data collection cycle. This leads both to a decrease in count rate and can cause the ion crystal to melt. Additionally, by parking the laser frequency farther off resonance in the red-detuned side of the Lorentizan atom line, this reduces the relative effect of cyclical, micromotion induced detuning on the atomic scattering rate, yielding a higher fidelity for the photon temporal waveform measured.

In a trap with the ability to fully compensate micromotion, high temporal waveform fidelity photons could be produced on resonance and as such the need to detune is a system specific limitation. However, it is briefly worth noting the effects of detuning which can be

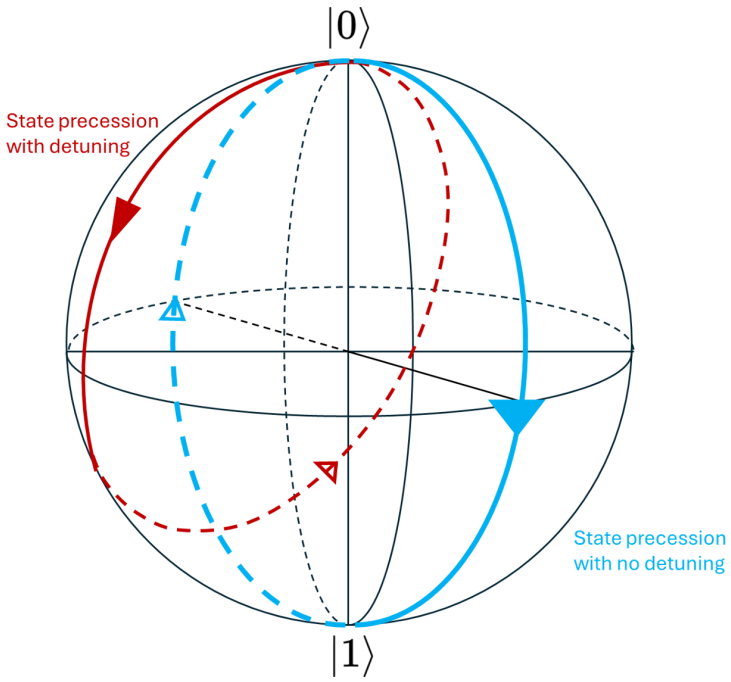

Figure 5.12: Bloch–sphere representation of state evolution under resonant and detuned driving. The blue trajectory shows state precession for a resonant drive ($\Delta = 0$), which follows a great circle connecting $|0\rangle$ and $|1\rangle$. The red trajectory shows precession for a finite detuning ($\Delta \neq 0$). The effective rotation axis is tilted toward $|0\rangle$, so the state follows a smaller circle and accumulates an additional phase relative to the resonant case.

understood via the generalized Rabi frequency,

$$\Omega(t) = \sqrt{\Omega_R(t)^2 + \Delta(t)^2} \tag{5.4}$$

where $\Omega_R(t)$ is the 'regular' Rabi frequency and $\Delta(t)$ is the time dependent detuning. In the rotating frame, we see that a time dependent detuning leads to a continuous precession about $\hat{z}$ of the state vector on the Bloch sphere. This leads the state to rotate about the titled axis $(\Omega, 0, \Delta)$. Thus rather than driving along a great circle from $|0\rangle$ to $|1\rangle$, with a constant Rabi frequency and detuning, the state vector instead traces out closed loops with a shearing angle. A conceptual representation of this is shown in Fig. 5.12.

If we consider the effects of such a detuning on single-photon remote entanglement protocols using the Zeeman levels of ion to encode the qubit basis (with logical basis states still referred to as $|0\rangle$ and $|2\rangle$ to maintain consistency with our prior labeling and excited state $|1\rangle$), detuning leads to phase accumulation between $|0\rangle$ and $|1\rangle$ for over the pulse duration. However, as the interference of photons at a beam splitter eliminates path information so long as the same detuning is applied to both emitters photons remain indistinguishable and resultant node-wise phases accumulated during detuned excitation are identical and can be canceled as a shared, common-mode global phase. On the other hand, when two nodes are driven by different time dependent detunings, this leads to a relative phase of the entangled state which will manifest in subsequent operations on the two nodes. Even in this case, if the relative detuning between the two nodes is known, this phase can be accounted for and canceled by the correct application of phase gates on the qubits prior to future operations. In the worst case, when the detunings are noisy or untracked, the effect will equivalent to phase noise, washing out the Bell-state phase and reducing entangled state fidelity. In the two-photon protocol, coincident detection of two photons in orthogonal modes projects the ions into an entangled state, and this projection is insensitive to the overall optical phase from the detuned excitation, which can therefore be ignored.

For our system, we find even with the addition of micromotion induced time dependent detunings not initially accounted for in numerical simulations, we can still achieve high

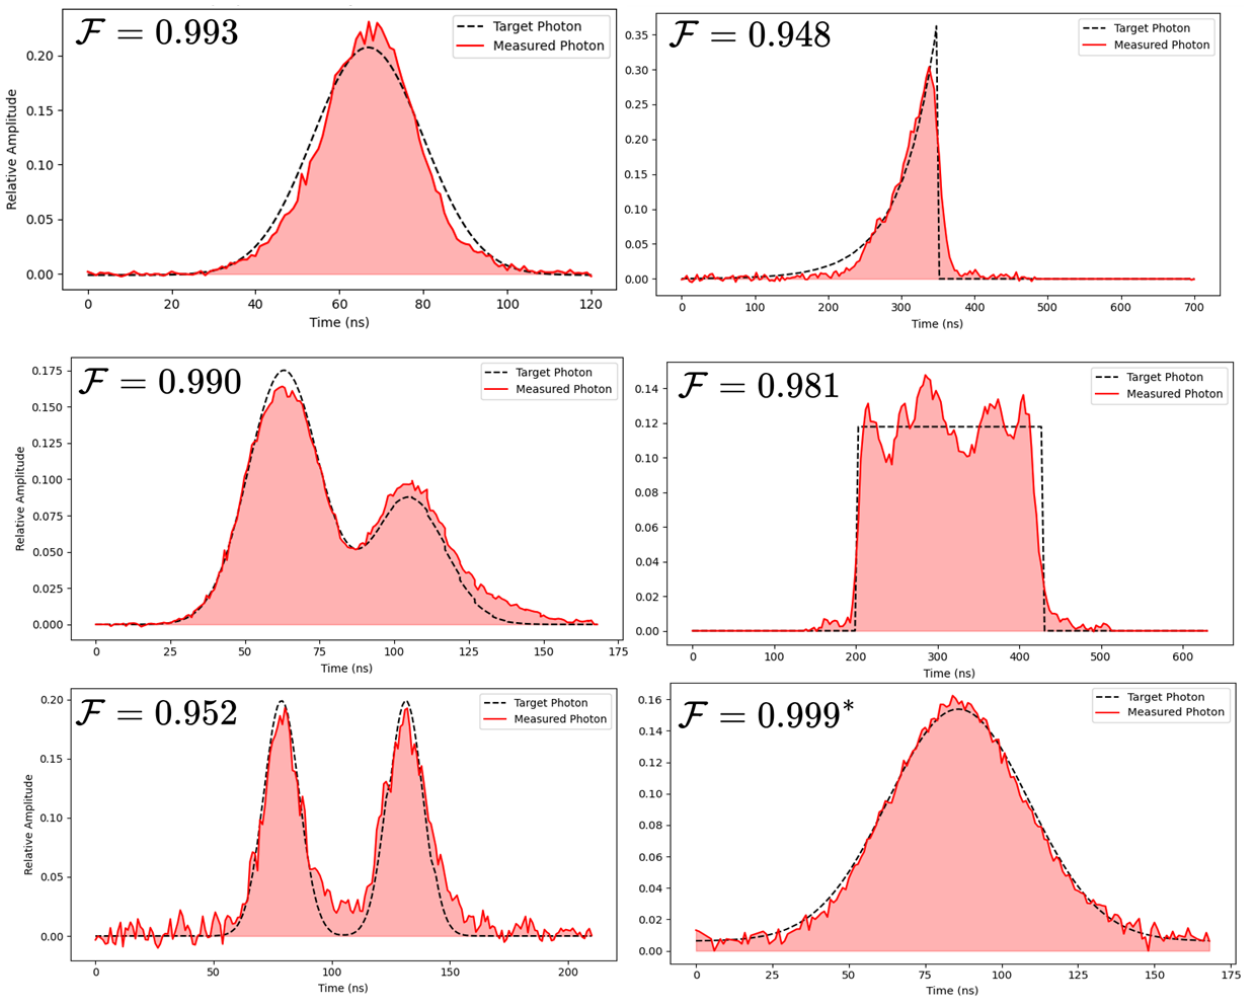

Figure 5.13: Selected single photon distributions showing shapes optimized for hybrid systems, quantum state transfer and interferometric stability. The first five distributions show the fidelity, defined by the mode overlap, between the initial target and final measured photon distributions. The final single photon measurement, indicated by an '*' was matched to an optimally fitted Gaussian after measurement.

fidelities between target and measured photon temporal waveforms as defined by their mode overlap. Experimentally, we calculate the fidelity of a measured photon distribution by taking the square root of the normalized photon count histogram and corresponding target distribution and calculating the inner product. As single photons generated during spontaneous do not have a defined phase, this is the correct metric to assess measured and target state fidelities. Some selected measured single photon distributions are shown in Fig. 5.13. We demonstrate the ability to produce Gaussian photons with fidelities as high as 0.999. Such photon distributions should prove suitable for proposed hybrid entanglement schemes [46].

Furthermore, while we are experimentally limited in our ability to generate short, discontinuous target distributions such as rising exponentials and flat top photons due to hardware limitations, we have achieved fidelities as high as 0.948 and 0.981 respectively. We assess the principal limitation on our ability to produce these distributions as resulting from the low bandwidth (80 MHz) and rise/fall times (8.4 ns) of the AWG employed in our experiment. For true sharp features, target synthesizer speeds should be much faster than the relevant atomic response time set by the excited state lifetime. We attribute the most likely source of an inability to see the effects of phase modulation from applied pulses to the same hardware limitations, although we cannot make this claim definitively as we did not employ a homodyne detection apparatus to measure the relative phase of our excitation pulses across the pulse period. Additionally, the effects on laser wavefront phase coherence resulting from the optical system down stream from our modulators (including coupling into single mode fibers) has not yet been modeled. With sufficiently high time resolution modulators and synthesizers, such as fast EOMs and an AWG only moderately higher bandwidths, demonstration of this additional resource for photon temporal waveform control should be possible.

However, we still demonstrate the ability to generate photons with temporal distributions shorter than the lifetime limited excited state decay via the effects of dephasing. Returning to the generalized Rabi frequency described above, we see that we are able to invert the direction of population transfer faster than would otherwise possible by taking advantage of the speed

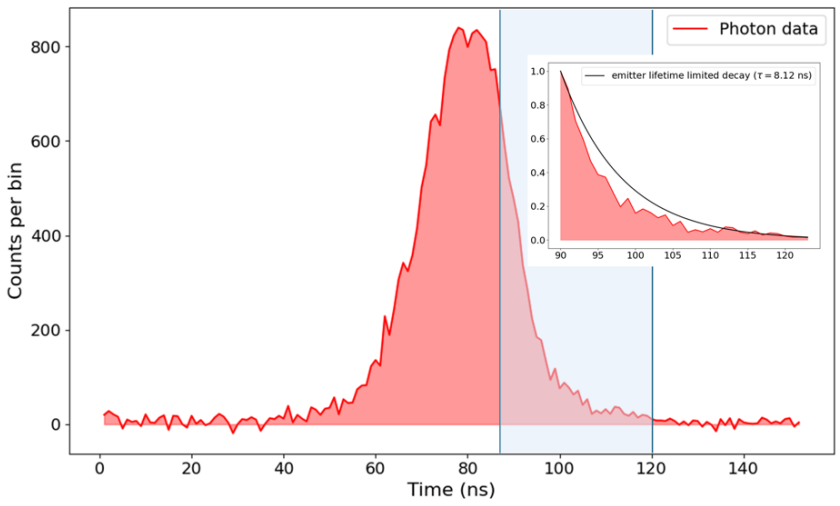

Figure 5.14: A single photon measurement with characteristic distributions shorter than the excited state lifetime limit produced via time dependent detuning. The inset shows the area highlighetd in blue along side the shorter decay that would be possible for a photon limited by the non-detuned atomic transition linewidth.

up granted by the additional detuning term without the use of phase. Experimentally, we find that we can produce distributions of shorter time scales than natural photons with one representative result shown in Fig. 5.14. This opens the door towards other methods for photon temporal control.

So long as the phase accumulation due to detuning is trackable, this method could be used to generate photons shorter than excited state lifetime with only modest Rabi frequencies and could be used to generate photons with low emission probabilities suitable for hybrid single-photon remote entanglement protocols. For same qubit type networks, intentionally leveraging detuning could be an even more flexible tool owing to the ability to treat detunings as a cancelable global phase.

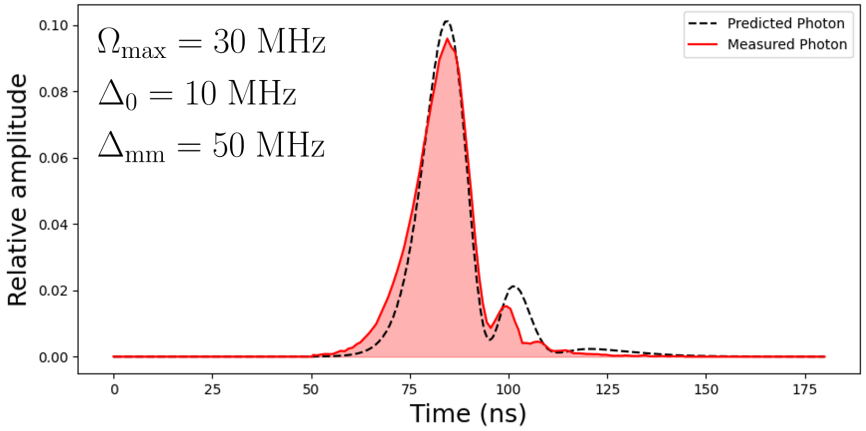

Figure 5.15: A photon temporal waveform measurement showing Rabi flopping under the effects of time dependent detuning. In both experimental and predicted photon distributions, we observed a small partial recovery of the photon temporal amplitude indicating the quantum state has begun precessing back towards the excited state, albeit along some intermediate and time dependent axis of rotation.

Finally, as a point of scientific interest, we present one more photon distribution which demonstrates the presences of Rabi flopping at a generalized Rabi frequency over the duration of a single atomic emission. Figure 5.15 shows both an experimentally measured photon

distribution and an attempt to numerically match the experimental results by simulating our three level system. The simulation includes a RF synchronized detuning modeling our trap behavior, application of a time dependent Rabi frequency corresponding to the excitation pulse, and adjustment of the timing of the pulse relative to the RF period as well a constant lab frame laser detuning. Although the displayed match is imperfect, we recreate the basic features seen in our measured photon, most interestingly being the partial revival of excited state population and corresponding increase in photon scattering rate. So far as we know, direct measurement of Rabi flopping in the signature of single photon temporal waveforms has not previously been observed. Therefore, while this flopping is not complete in the sense of a full population between the ground and excited states, we still find this to be satisfying coda to our work towards experimental control of photon temporal waveforms.

5.6 Looking back and looking forward

The work of my PhD has been primarily dedicated towards contributing in some small way to advancing the performance of quantum networks. As quantum information sciences continues to grow and develop, it is clear that the field must continue to push beyond single isolated quantum simulators or computers with tens, or maybe some day soon hundreds of logical qubits, located in the basements of university departments or behind company walls. While much work still remains to be done on local devices, I believe that the greatest hurdle towards universal, powerful quantum devices will come in the form of achieving robust, high rate exchange of quantum information between distant network nodes. To date, the best demonstrated remote entanglement rates are still in the low hundreds. This number will need to grow by at least several orders of magnitude if we are to envision many of the most promising prospects that quantum computing theorists tell us are possible.

To that end, this dissertations details the development of new methods for the control of the temporal waveform of photons from single quantum emitters. We have proposed new techniques whereby emitter of any lifetime could produce photons of any temporal distribution without the need for experimental overhead in the form of bulky or expensive optics. We

have also detailed several approaches to using quantum jumps methods to identify optimal parameters for quantum network protocols. Finally, as an experimentalist at heart, the work I am most excited by has been the practical realization of photon temporal waveform control in an ion trap. We have provided a proof of principle that such control can be achieved with promising fidelities with only minimal hardware requirements.

I hope to see the ideas laid out for the use of phase modulation to allow coherent de-excitation from atomic excited states used to expand the results presented herein. With an only moderate upgrade to existing systems, this level of control should be in sight. Perhaps future directions for photon control could continue to explore methods to use time dependent detunings and other experimental knobs to find a new paths forward.

BIBLIOGRAPHY

[1]J. Preskill, *Quantum computing 40 years later*, 2021.

[2]P. Crescenzi, D. Goldman, C. Papadimitriou, A. Piccolboni, and M. Yannakakis, "On the complexity of protein folding", Journal of Computational Biology **5**, 423–465 (1998).

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

[4]D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient z gates for quantum computing", Physical Review A **96**, 022330 (2017).

[5]J. D. Wong-Campos, S. A. Moses, K. G. Johnson, and C. Monroe, "Demonstration of two-atom entanglement with ultrafast optical pulses", Physical Review Letters **119**, 230501 (2017).

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

[7]D. P. DiVincenzo, *Topics in quantum computers*, 1996.

[8]D. P. DiVincenzo, "The physical implementation of quantum computation", Fortschritte der Physik **48**, 771–783 (2000).

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

[10]J.-S. Chen, E. Nielsen, M. Ebert, V. Inlek, K. Wright, V. Chaplin, A. Maksymov, E. Páez, A. Poudel, P. Maunz, and J. Gamble, "Benchmarking a trapped-ion quantum computer with 30 qubits", Quantum **8**, 1516 (2024).

[11]C. Langer, R. Ozeri, J. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. Blakestad, J. Britton, D. Hume, W. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. Schmidt, and D. Wineland, "Long-lived qubit memory using atomic ions", Physical Review Letters **95**, 060502 (2005).

[12]P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, "Single ion qubit with estimated coherence time exceeding one hour", Nature Communications **12**, 233 (2021).

[13]C. W. Hogle, D. Dominguez, M. Dong, A. Leenheer, H. J. McGuinness, B. P. Ruzic, M. Eichenfield, and D. Stick, "High-fidelity trapped-ion qubit operations with scalable photonic modulators", npj Quantum Information **9**, 74 (2023).

[14]C. Löschnauer, J. Mosca Toba, A. Hughes, S. King, M. Weber, R. Srinivas, R. Matt, R. Nourshargh, D. Allcock, C. Ballance, C. Matthiesen, M. Malinowski, and T. Harty, "Scalable, high-fidelity all-electronic control of trapped-ion qubits", PRX Quantum **6**, 040313 (2025).

[15]L. A. Zhukas, P. Svihra, A. Nomerotski, and B. B. Blinov, "High-fidelity simultaneous detection of a trapped-ion qubit register", Physical Review A **103**, 062614 (2021).

[16]B. B. Blinov, D. L. Moehring, L. M. Duan, and C. Monroe, "Observation of entanglement between a single trapped atom and a single photon", Nature **428**, 153–157 (2004).

[17]P. Kobel, M. Breyer, and M. Köhl, "Deterministic spin-photon entanglement from a trapped ion in a fiber fabry–perot cavity", npj Quantum Information **7**, 6 (2021).

[18]A. Stute, B. Casabone, P. Schindler, T. Monz, P. O. Schmidt, B. Brandstätter, T. E. Northup, and R. Blatt, "Tunable ion–photon entanglement in an optical cavity", Nature **485**, 482–485 (2012).

[19]L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, "High-rate, high-fidelity

entanglement of qubits across an elementary quantum network", Physical Review Letters **124**, 110501 (2020).

[20]C. Auchter, C.-K. Chou, T. W. Noel, and B. B. Blinov, "Ion–photon entanglement and bell inequality violation with ˆ138baˆ+", Journal of the Optical Society of America B **31**, 1568 (2014).

[21]J. O'Reilly, G. Toh, I. Goetting, S. Saha, M. Shalaev, A. L. Carter, A. Risinger, A. Kalakuntla, T. Li, A. Verma, and C. Monroe, "Fast photon-mediated entanglement of continuously cooled trapped ions for quantum networking", Physical Review Letters **133**, 090802 (2024).

[22]D. T. C. Allcock, W. C. Campbell, J. Chiaverini, I. L. Chuang, E. R. Hudson, I. D. Moore, A. Ransford, C. Roman, J. M. Sage, and D. J. Wineland, "*omg* blueprint for trapped ion quantum computing with metastable states", Applied Physics Letters **119**, 214002 (2021).

[23]C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, "High-fidelity quantum logic gates using trapped-ion hyperfine qubits", Physical Review Letters **117**, 060504 (2016).

[24]A. C. Hughes, R. Srinivas, C. M. Löschnauer, H. M. Knaack, R. Matt, C. J. Ballance, M. Malinowski, T. P. Harty, and R. T. Sutherland, *Trapped-ion two-qubit gates with >99.99% fidelity without ground-state cooling*, Version Number: 1, 2025.

[25]M. Holzscheiter, "Ion-trap quantum computation", Los Alamos Science **27**, 264–283 (2002).

[26]S. Saha, M. Shalaev, J. O'Reilly, I. Goetting, G. Toh, A. Kalakuntla, Y. Yu, and C. Monroe, "High-fidelity remote entanglement of trapped atoms mediated by time-bin photons", Nature Communications **16**, 2533 (2025).

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

[28]L. S. Brown and G. Gabrielse, "Geonium theory: physics of a single electron or ion in a penning trap", Reviews of Modern Physics **58**, 233–311 (1986).

[29]J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, "Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins", Nature **484**, 489–492 (2012).

[30]D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, "Experimental issues in coherent quantum-state manipulation of trapped atomic ions", Journal of Research of the National Institute of Standards and Technology **103**, 259–328 (1998).

[31]W. Paul, "Electromagnetic traps for charged and neutral particles", Reviews of Modern Physics **62**, 531–540 (1990).

[32]P. K. Ghosh, *Ion traps*, Repr, International series of monographs on physics 90 (Clarendon Press, Oxford, 2007), 326 pp.

[33]G. H. Low, P. F. Herskind, and I. L. Chuang, "Finite-geometry models of electric field noise from patch potentials in ion traps", Physical Review A **84**, 053425 (2011).

[34]T. Hänsch and A. Schawlow, "Cooling of gases by laser radiation", Optics Communications **13**, 68–69 (1975).

[35]C. J. Foot, *Atomic physics*, Repr. (twice with corr.), Oxford master series in physics Atomic, optical, and laser physics 7 (Oxford Univ. Press, Oxford, 2011), 331 pp.

[36]F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Laser cooling to the zero-point energy of motion", Physical Review Letters **62**, 403–406 (1989).

[37]C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, "Resolved-sideband raman cooling of a bound atom to the 3d zero-point energy", Physical Review Letters **75**, 4011–4014 (1995).

[38]S. Ejtemaee and P. C. Haljan, "3d sisyphus cooling of trapped ions", Physical Review Letters **119**, 043001 (2017).

[39]L. Feng, W. L. Tan, A. De, A. Menon, A. Chu, G. Pagano, and C. Monroe, "Efficient ground-state cooling of large trapped-ion chains with an electromagnetically-induced-transparency tripod scheme", Physical Review Letters **125**, 053001 (2020).

[40]G. Morigi, J. Eschner, and C. H. Keitel, "Ground state laser cooling using electromagnetically induced transparency", Physical Review Letters **85**, 4458–4461 (2000).

[41]M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, "Ion-trap measurements of electric-field noise near surfaces", Reviews of Modern Physics **87**, 1419–1482 (2015).

[42]L. M. A. Nguyen, B. Bowers, and S. Mouradian, *The effect of trap design on the scalability of trapped-ion quantum technologies*, 2025.

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

[44]Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland, "Heating of trapped ions from the quantum ground state", Physical Review A **61**, 063418 (2000).

[45]S. Olmschenk, D. Hayes, D. N. Matsukevich, P. Maunz, D. L. Moehring, and C. Monroe, "QUANTUM LOGIC BETWEEN DISTANT TRAPPED IONS", International Journal of Quantum Information **08**, 337–394 (2010).

[46]J. F. Lilieholm, V. Niaouris, A. Kato, K.-M. C. Fu, and B. B. Blinov, "Photon-mediated entanglement scheme between a ZnO semiconductor defect and a trapped yb ion", Applied Physics Letters **117**, _eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/5.0019892/14541946/154002_1_online.pdf, 154002 (2020).

[47]Y. Chu, J. D. Pritchard, H. Wang, and M. Weides, "Hybrid quantum devices: guest editorial", Applied Physics Letters **118**, _eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/5.0057740/14548794/240401_1_online.pdf, 240401 (2021).

[48]H. K. Beukers, M. Pasini, H. Choi, D. Englund, R. Hanson, and J. Borregaard, "Remote-entanglement protocols for stationary qubits with photonic interfaces", PRX Quantum **5**, 010202 (2024).

[49]J. D. Siverns, J. Hannegan, and Q. Quraishi, "Neutral-atom wavelength-compatible 780 nm single photons from a trapped ion via quantum frequency conversion", Physical Review Applied **11**, 014044 (2019).

[50]S. Yu, K. Lee, S. Park, K. Kim, J. Goo, J. Park, and T. Kim, "Efficient quantum frequency conversion of ultra-violet single photons from a trapped ytterbium ion", Applied Physics Letters **126**, 084001 (2025).

[51]J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, and H. J. Kimble, "Deterministic generation of single photons from one atom trapped in a cavity", Science **303**, 1992–1994 (2004).

[52]E. N. Knall, C. M. Knaut, R. Bekenstein, D. R. Assumpcao, P. L. Stroganov, W. Gong, Y. Q. Huan, P.-J. Stas, B. Machielse, M. Chalupnik, D. Levonian, A. Suleymanzade, R. Riedinger, H. Park, M. Lončar, M. K. Bhaskar, and M. D. Lukin, "Efficient source of shaped single photons based on an integrated diamond nanophotonic system", Physical Review Letters **129**, 053603 (2022).

[53]K. Sosnova, "Mixed-species ion chains for quantum networks", PhD thesis (University of Maryland, College Park).

[54]T. Faorlin, B. Yadin, Y. Weiser, G. Araneda, S. Nimmrichter, L. Panzl, T. Lafenthaler, R. Blatt, T. Monz, and G. Cerchiari, *Controlling the spontaneous emission of trapped ions*, 2025.

[55]L. Slodička, G. Hétet, N. Röck, P. Schindler, M. Hennrich, and R. Blatt, "Atom-atom entanglement by single-photon detection", Physical Review Letters **110**, Publisher: American Physical Society, 083603 (2013).

[56]S. L. N. Hermans, M. Pompili, L. D. Santos Martins, A. R-P Montblanch, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, "Entangling remote qubits using the single-photon protocol: an in-depth theoretical and experimental study", New Journal of Physics **25**, 013011 (2023).

[57]L.-M. Duan, B. Blinov, D. Moehring, and C. Monroe, "Scalable trapped ion quantum computation with a probabilistic ion-photon mapping", Quantum Information and Computation **4**, 165–173 (2004).

[58]T. Van Leent, M. Bock, F. Fertig, R. Garthoff, S. Eppelt, Y. Zhou, P. Malik, M. Seubert, T. Bauer, W. Rosenfeld, W. Zhang, C. Becher, and H. Weinfurter, "Entangling single atoms over 33 km telecom fibre", Nature **607**, 69–73 (2022).

[59]H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, "Heralded entanglement between solid-state qubits separated by three metres", Nature **497**, 86–90 (2013).

[60]V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, "Entanglement of trapped-ion qubits separated by 230 meters", Physical Review Letters **130**, 050803 (2023).

[61]A. N. Craddock, J. Hannegan, D. P. Ornelas-Huerta, J. D. Siverns, A. J. Hachtel, E. A. Goldschmidt, J. V. Porto, Q. Quraishi, and S. L. Rolston, "Quantum interference between photons from an atomic ensemble and a remote atomic ion", Physical Review Letters **123**, 213601 (2019).

[62]M. Stobińska, G. Alber, and G. Leuchs, "Perfect excitation of a matter qubit by a single photon in free space", EPL (Europhysics Letters) **86**, 14007 (2009).

[63]A. Golla, B. Chalopin, M. Bader, I. Harder, K. Mantel, R. Maiwald, N. Lindlein, M. Sondermann, and G. Leuchs, "Generation of a wave packet tailored to efficient free space excitation of a single atom", The European Physical Journal D **66**, 190 (2012).

[64]J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, "Quantum state transfer and entanglement distribution among distant nodes in a quantum network", Physical Review Letters **78**, 3221–3224 (1997).

[65]P. Maunz, D. L. Moehring, S. Olmschenk, K. C. Younge, D. N. Matsukevich, and C. Monroe, "Quantum interference of photon pairs from two remote trapped atomic ions", Nature Physics **3**, Publisher: Nature Publishing Group, 538–541 (2007).

[66]P. P. Rohde, T. C. Ralph, and M. A. Nielsen, "Optimal photons for quantum-information processing", Physical Review A **72**, Publisher: American Physical Society, 052332 (2005).

[67]F. Sośnicki, M. Mikołajczyk, A. Golestani, and M. Karpiński, "Interface between picosecond and nanosecond quantum light pulses", Nature Photonics **17**, 761–766 (2023).

[68]P. Kolchin, C. Belthangady, S. Du, G. Y. Yin, and S. E. Harris, "Electro-optic modulation of single photons", Physical Review Letters **101**, 103601 (2008).

[69]L. J. Wright, M. Karpiński, C. Söller, and B. J. Smith, "Spectral shearing of quantum light pulses by electro-optic phase modulation", Physical Review Letters **118**, 023601 (2017).

[70]F. Sośnicki and M. Karpiński, "Large-scale spectral bandwidth compression by complex electro-optic temporal phase modulation", Optics Express **26**, 31307 (2018).

[71]R. L. Lecamwasam, M. R. Hush, M. R. James, and A. R. R. Carvalho, "Measurement-based generation of shaped single photons and coherent state superpositions in optical cavities", Physical Review A **95**, 013828 (2017).

[72]M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther, "Continuous generation of single photons with controlled waveform in an ion-trap cavity system", Nature **431**, 1075–1078 (2004).

[73]X. You and Y.-M. He, "Developing the quantum-dot single-photon sources with excellent performance by coupling asymmetric microcavity", Physical Review Research **7**, L012016 (2025).

[74]J. P. Covey, H. Weinfurter, and H. Bernien, "Quantum networks with neutral atom processing nodes", npj Quantum Information **9**, 90 (2023).

[75]G. S. Vasilev, D. Ljunggren, and A. Kuhn, "Single photons made-to-measure", New Journal of Physics **12**, 063024 (2010).

[76]A. Kuhn, M. Hennrich, and G. Rempe, "Deterministic single-photon source for distributed quantum networking", Physical Review Letters **89**, 067901 (2002).

[77]P. B. R. Nisbet-Jones, J. Dilley, A. Holleczek, O. Barter, and A. Kuhn, "Photonic qubits, qutrits and ququads accurately prepared and delivered on demand", New Journal of Physics **15**, 053007 (2013).

[78]D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance", Nature **449**, 68–71 (2007).

[79]C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, "Creation of entangled states of distant atoms by interference", Physical Review A **59**, 1025–1033 (1999).

[80]D. N. Matsukevich, P. Maunz, D. L. Moehring, S. Olmschenk, and C. Monroe, "Bell inequality violation with two remote atomic qubits", Physical Review Letters **100**, 150404 (2008).

[81]S. C. Connell, J. Scarabel, E. M. Bridge, K. Shimizu, V. Blūms, M. Ghadimi, M. Lobino, and E. W. Streed, "Ion–photonic frequency qubit correlations for quantum networks", Journal of Physics B: Atomic, Molecular and Optical Physics **54**, 175503 (2021).