

Quantum Control of Be^+ Ions

Edward K. Kleiner

Advisor: Professor David A. Hanneke
May 5, 2016

Submitted to the
Department of Physics and Astronomy
of Amherst College
in partial fulfillment of the
requirements for the degree of
Bachelors of Arts with honors

© 2016 Edward K. Kleiner

Abstract

Precisely controlling the quantum states of atoms is a delicate and complex endeavor. In this thesis, I describe the methods used to prepare Be^+ ions for quantum logic experiments, with an emphasis on the trapping and cooling of the ions. We can already produce trapped ions cooled to near the Doppler limit, and we perform calculations in preparation for the use of resolved side-band cooling to bring the ions below the Doppler temperature.

This thesis also outlines the design and construction of a tapered amplifier apparatus intended to provide increased laser power in order to improve our ability to manipulate the states of trapped ions. This device receives 10 mW of 940 nm light from a seed laser and outputs over 110 mW from a fiber.

Acknowledgments

This thesis is indebted to many people, but it owes far and away the most to Professor Hanneke. After three classes, three summers in the lab, and a thesis I can say without exaggeration that he is chiefly responsible for everything I know about physics. Without his support, both in the lab and outside of it, this thesis would not exist.

That is not to give short shrift to the rest of the physics and astronomy department, both students and professors, who have helped me in innumerable ways during my time at Amherst. Specific thanks are due to Jim Kubasek, who was always there to machine a part for me, fix my SolidWorks design, and then remachine that part for me; and to Norm Page, who always knew what my circuit needed even if I didn't.

Thanks as well to all those who worked in Prof. Hanneke's lab with me – you all made the lab feel like a second home. Acknowledgments are also due to all those who came before me, without whose work my thesis would not be possible – to you, I bequeath a citation and a spot in my bibliography.

I am incredibly thankful for the friendship and support of my dear suite-mates – it was a pleasure to live with you. You were always there to listen to me complain about my thesis, as long as I was willing to return the favor.

And to my dear Sara, who kept me lighthearted throughout the process, I owe you the profoundest debt of all.

And finally, I am tremendously grateful for the love and generosity of my dear family, who helped me every way they could, with care packages, advice, and even a few crucial edits. Mom, Dad, Matthew, and Aaron: thank you.

This material is based upon work supported by the Amherst College Dean of the Faculty and by the National Science Foundation under CAREER Grant No. PHY-1255170.

Contents

1	Introduction	2
1.1	Quantum Logic Spectroscopy	2
1.1.1	Proton-Electron Mass Ratio	3
1.2	Ion Trapping	5
1.3	Doppler Cooling	7
1.4	Resolved Sideband Cooling	8
1.4.1	Repumping	12
2	Lab Overview	19
2.1	Ultra-High Vacuum	20
2.2	The Trap	22
2.2.1	Helical Resonator	29
2.3	Frequency Tripling	30
2.3.1	Pound-Drever-Hall Lock	32
2.4	Laser Stabilization and Tuning	36
3	Laser Amplification	40
3.1	Design	40
3.2	Apparatus	43
3.2.1	Tapered Amplifier	43
3.2.2	Lenses	46
3.2.3	Other Components	48
3.3	Beam Shape	50
3.4	Coupling	52
4	Conclusion	57
4.1	Tapered Amplifier Design Notes	57
4.2	Future Work	60
4.2.1	Box	61
4.2.2	Doppler Experiment	62

4.2.3 Resolved Sideband Cooling	62
A Frequencies and Detunings	64

List of Figures

1.1	Picture of the Linear Paul Trap	6
1.2	Sideband Frequency Plot	11
1.3	Repumping diagram	14
1.4	Boltzmann distribution at the Doppler limit	15
1.5	Sideband cooling and repumping	17
1.6	Pulse sequence example	18
2.1	First Ion Crystal	25
2.2	Be ions in trap with varying U_0	26
2.3	Be and BeH ions	27
2.4	Diagram of the oscillator for the PDH	33
2.5	Block diagram of the PDH	34
2.6	Pound-Drever-Hall Signal	37
3.1	Tapered amplifier design	44
3.2	Cartoon of the TA chip	45
3.3	Lens diagram	48
3.4	Output beam shape in the horizontal	51
3.5	Output beam shape in the vertical	52
3.6	Zoom on the brass adjuster designs	53
3.7	Plot of beam power	55
4.1	Picture of the current apparatus	61

Chapter 1

Introduction

1.1 Quantum Logic Spectroscopy

The field of atomic trapping has played a prominent role in experimental physics over the last several decades. The ability to confine atoms and ions has paved the way for the investigation of a diverse range of subjects; trapped atoms or atomic ions have been used to create types of matter ranging from Bose-Einstein condensates[1] to plasmas[2] and to break new ground in fields like parity violation[3] and quantum computing.[4]

But the field of molecular trapping is still in its infancy. While trapping molecular ions is fairly straightforward – or at least no harder than trapping atomic ions – it is significantly more difficult to manipulate their states, since the addition of rotational and vibrational states adds multiple degrees of complexity. Without the ability to cool and probe the trapped molecular ions, there is very little physics that can be done.

In our lab, we plan to study molecular transitions using a relatively new technique known as quantum logic spectroscopy (QLS).[5] QLS relies on trapping two ion species simultaneously and coupling them using the Coulomb force. One ion, the “logic ion,” has transitions that are useful for cooling and state manipulation, while the other, called the “spectroscopy ion,” has the transitions we are interested in. In our case, we plan to use an atomic ion, ${}^9\text{Be}^+$, as the logic ion, on which we will perform cooling procedures in order to prepare our spectroscopy ion, ${}^{16}\text{O}_2^+$, in a known state so that we can investigate its transitions.

Using QLS to study trapped molecular ions will allow researchers to study types of transitions that were previously inaccessible. This has applications across a range of subjects nearly as diverse as those studied using atomic trapping, from astrophysical spectroscopy to quantum chemistry.[6] In our experiments, we will use QLS to investigate a question related to the fundamental constants of the universe: how do you weigh a proton?

1.1.1 Proton-Electron Mass Ratio

The difficulty of weighing a proton lies in the fact that our measure of mass is based on the proton – the official kilogram (a cylinder of platinum-iridium alloy kept in a vault in France) is about 40% protons by weight, and any conceivable scale would of course be made out of protons. Thus, the official weight for the proton, $m_p \approx 1.67 \times 10^{-27}$ kg, is in some ways a circular definition, since the values $1 m_p$ and 1 kg are interrelated. One way around this problem is to measure the proton mass as a function of some other fundamental mass,

like the mass of the electron m_e . By measuring the dimensionless ratio $\mu = m_e/m_p$, we can avoid using the kilogram at all and thus dodge the problem of interdependent units.

Under the standard model, of course, this is a relatively unexciting debate. Many people have very precisely measured the values of both m_p and m_e , and NIST's website lists a value of μ to ten decimal points. However, certain unified field theories that predict the existence of more than three spatial dimensions leave open the possibility that m_p , which is almost entirely a function of the strong force binding together the three quarks that make up a proton, is in fact varying as a function of time. On the other hand, m_e is a function of the electron's interaction with the Higgs field, which is governed by the weak force. Since, for all the reasons mentioned above, there are issues with measuring m_p directly, we can measure time variation in μ (we will in fact measure $d\mu/\mu$) as a proxy for determining whether there is a change in the relative strengths of the strong and weak field.

To perform these experiments, we first have to be able to control the temperature and state of our logic ion, Be^+ . This thesis outlines the steps taken so far in trapping and cooling Be^+ ions, including a description of my own efforts to construct a tapered amplifier system, and presents a plan for the next steps for manipulating the quantum states of Be^+ .

1.2 Ion Trapping

Earnshaw’s theorem states that it is impossible to create an electrostatic trap – that is, there is no arrangement of stationary charges that produce a stable potential minimum.[7] Instead, ion trapping employs changing electric fields, which can be timed and configured such that the average force on the particles is toward the center of the trap.

We employ a design known as the “linear Paul trap,” which combines static and oscillating electric fields in order to trap ions[8, 9]. The trap is referred to as “linear” because it is elongated in one direction; the static fields are effectively caps that prevent the ions from escaping along that axis. The oscillating fields are used to trap the ions along the other two dimensions, but in accordance with Earnshaw’s Theorem they can only produce a confining field in one of the two directions. At any moment the oscillating fields produce a confining force along one axis and an anti-confining force along the other – this is in some ways analogous to the gravitational potential of a saddle: a ball placed in the center of the saddle (its “saddle point”) can’t roll up the pommel, but it can fall off the sides. However, if the saddle (or the electric field) is rotating around an axis through the saddle point within a certain range of angular velocities, then before the ball (or ion) can move very far down the potential gradient the system has undergone a 90° rotation, switching the directions of the confining and anti-confining forces and pushing the ball back toward the center. The ball/ion ends up oscillating around the saddle point in two directions, which, combined with the ion’s confinement in the elongate direction with static fields, constitutes a trap.

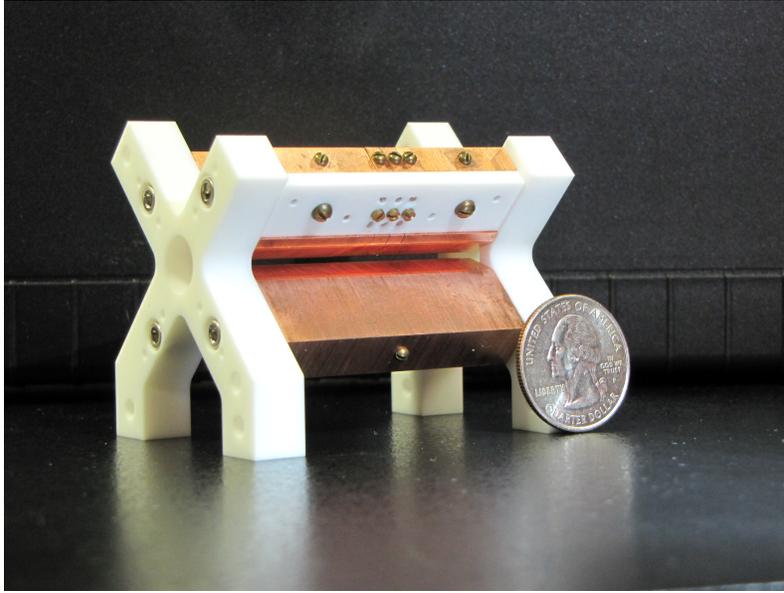


Figure 1.1: A picture of the ion trap, with a quarter for scale. The upper, segmented electrode is the DC electrode, while the RF electrode is on the bottom.

The trap we use consist of two pairs of electrodes, known as the RF and DC electrodes. The RF (short for radiofrequency) electrodes are the “spinning saddle” – or, in our case, more of a “flopping saddle,” since their voltage is oscillating around zero at about 20 MHz while the DC (direct current) electrodes remain constant. Each DC electrode is composed of five smaller electrodes, each of which can be set to a different constant voltage. For trapping, we set the first and fifth component of each DC electrode to a higher voltage (generally around 10 V) to serve as the endcaps, while the middle components are set to lower voltages (no more than a couple volts, and usually zero).

1.3 Doppler Cooling

For the purposes of quantum logic spectroscopy, it is necessary to do more than simply trap the Be^+ ions – we also have to cool them in order to prepare the coupled molecular ions in a known state. The easiest method is a form of laser cooling known as “Doppler cooling” since it relies on the Doppler effect to decrease the atoms’ temperature.

Laser cooling of atoms generally works by using photons to excite transitions from one energy level to another. Under certain conditions, the process of excitation and emission can result in a net cooling for the atoms. To Doppler cool Be^+ , we use the transition $2S_{1/2} \leftrightarrow 2P_{3/2}$, which has a frequency of 313.132 nm in the ultraviolet (see Appendix A for a table of relevant frequencies).[10] We shine a laser into the trap that is detuned to the red by about 9.7 MHz and that is therefore less likely to excite transitions in an atom at rest. But for an atom that is moving opposite the direction of beam propagation, the beam appears bluer due to the Doppler effect, meaning that the red-detuned beam is now at the right frequency to excite the transition. When one of these ions absorbs a photon, it receives a momentum kick in the direction that the beam is traveling. Since the ion was originally moving against the beam, this kick decreases its speed. When the ion decays back to its original state through spontaneous emission, it emits a photon in a random direction, and thus gets a kick in the opposite direction, which is of course also random. Since this kick is uncorrelated with the ion’s velocity, it will sometimes increase and sometimes decrease the ion’s speed. In general, the combination of these two kicks – the first of which always decreases the ion’s speed and the second of which can

either increase or decrease it – will tend to lower the temperature of the ions over many iterations.

However, this method has a limit, known as the Doppler temperature. This limit is derived from the discrete size of the momentum changes undergone by the atoms during absorption and emission. This discreteness results in a “random walk” of the particles’ momenta as they absorb and emit photons, which ends up heating the ions at a rate proportional to the linewidth of the transition. This proportionality occurs because the rate at which this random walk occurs – i.e., the rate at which the ions emit a photon – is inversely proportional to the lifetime of the excited state, which is itself inversely related to the linewidth of that state. At the same time, the ions are being cooled at a rate proportional to their velocity (which is itself proportional to their temperature), so we can find an equilibrium temperature – the Doppler temperature – below which the ions cannot be cooled. At optimum conditions, this limit $T_D = \hbar\gamma/2k_B$, where γ is the linewidth of the transition and k_B is the Boltzmann constant.[11] For the transition $2S_{1/2} \leftrightarrow 2P_{3/2}$ in Be^+ , the Doppler temperature $T_D = 466 \mu\text{K}$.

1.4 Resolved Sideband Cooling

To cool the ions below the Doppler limit, we plan to use a technique called “resolved sideband cooling.” This method draws its name from the fact that an ion oscillating at the motional frequency of the trap (f_m) will be resonant not only with a laser tuned to the resonant frequency (f_{res}), but also with two

sideband frequencies modulated by the motional frequency. In theory, a laser tuned to the red sideband, with $f_{red} = f_{res} - f_m$, would be able to remove motional energy from the ions until they were at the ground state. However, for our purposes – Be^+ ions in a trap with a motional frequency of about 1 MHz – this approach is unworkable due to the large linewidth of the relevant transition. This means that if we tried to cool the ions using a laser tuned to f_{red} , we would instead just be pushing the ions to the excited state, which, since excitation is always followed by spontaneous emission, is no different from Doppler cooling.

Instead, to perform resolved sideband cooling we will need another way to manipulate the motional states of the ions. We rely on the fact that if a trapped ion is sufficiently cold, its motional states closely resembles those of a quantum harmonic oscillator – that is, they’re quantized. We assume that we can couple these motional states with the atom’s internal states, which means that we can excite transitions that will simultaneously change the internal and motional states of the ion. Using lasers specifically tuned to these transitions, we can change the ions’ motional states without exciting electronic transitions that would heat the ions back to the Doppler temperature.

To determine the strength of these transitions, we assume that we begin with a thermal ensemble of internal and motional quantum states, where $|M_z\rangle = |\uparrow\rangle$ or $|\downarrow\rangle$ and $|n\rangle = |0\rangle, |1\rangle, |2\rangle, \dots$ represent the time-independent internal and motional eigenstates, respectively. If we ignore any other possible resonances, we see transitions between levels $|\downarrow, n\rangle$ and $|\uparrow, n'\rangle$ – this cycling is known as “Rabi flopping.” Schrödinger’s Equation gives us coupling coeffi-

icients equal to[12]

$$\dot{C}_{\uparrow,n'} = -i^{(1+|n'-n|)} e^{-i(\Delta t - \phi)} \Omega_{n',n} C_{\downarrow,n'} \quad (1.1)$$

$$\dot{C}_{\downarrow,n} = -i^{(1-|n'-n|)} e^{i(\Delta t - \phi)} \Omega_{n',n} C_{\uparrow,n'}. \quad (1.2)$$

If we let $n' = n + s$, then $\Omega_{n',n}$, which is equivalent to $\Omega_{n+s,n}$, is known as the $|s|^{\text{th}}$ blue/red sideband Rabi frequency for $s > 0/s < 0$. It is given by[13]

$$\Omega_{n+s,n} = \Omega_0 e^{-\eta^2/s} \eta^{|s|} \sqrt{\frac{n_{<}!}{n_{>}!}} L_{n_{<}}^{|s|}(\eta^2), \quad (1.3)$$

where $n_{<}$ and $n_{>}$ are equal to $\min(n, n + s)$ and $\max(n, n + s)$ respectively, η is the Lamb-Dicke parameter,¹ Ω_0 is the carrier Rabi frequency, and L_n^α is the generalized Laguerre polynomial

$$L_n^\alpha(X) = \sum_{m=0}^n (-1)^m \binom{n+\alpha}{n-m} \frac{X^m}{m!}. \quad (1.4)$$

To excite transitions between these two states, we use two lasers, both several gigahertz from resonance in order to suppress spontaneous emission, which would in effect return the atoms to the Doppler temperature. The lasers are then detuned relative to each other by a frequency equal to the hyperfine structure of ${}^9\text{Be}^+$ (f_0) plus or minus an integer multiple of the motional frequency; that is, $\Delta f = f_0 + s f_m$, $s \in \mathbb{Z}$. The value of s that we pick determines

¹The Lamb-Dicke parameter quantifies the strength of the coupling between the motional and internal states of the ion. It is defined as $\eta = \frac{2\pi}{\lambda} \sqrt{\frac{\hbar}{2m\omega_m}}$, which equals .475 for Be^+ with a 1 MHz motional frequency.

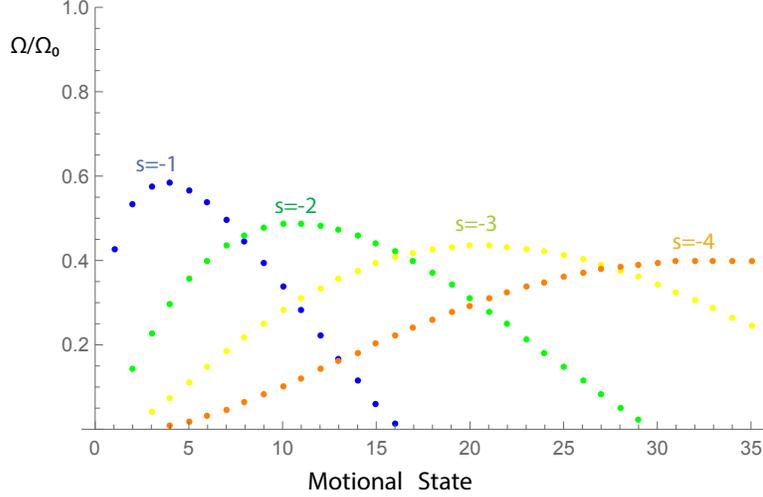


Figure 1.2: Ratios of the first four red sideband frequencies to the carrier Rabi frequency Ω_0 , as derived from Eqs. 1.3 and 1.4.

which sideband frequency we're cooling at: for example, $s = -1$ gives us the first red sideband, whereas $s = 3$ gives us the third blue sideband.

So a pair of lasers set to the first red sideband will cause an atom in the $|\downarrow, n\rangle$ state to move to the $|\uparrow, n - 1\rangle$ state (using specific labeling conventions), and vice versa. This holds true for all motional states n , which means that with this pair of lasers we can cycle atoms between all the pairs of adjacent states – with one important caveat. Controlling the atoms' motional states requires carefully timed laser pulses. With an exact π pulse, we can effectively switch the populations in two coupled states. The length of such a pulse is proportional to the inverse of the sideband frequency for that pair of states, and since the sideband frequency for two coupled states $|\downarrow, n\rangle$ and $|\uparrow, n + s\rangle$ is a function of both n and s [see Eq. 4], any given pulse will be timed properly for only one pair of states.

Moreover, a plot of $\Omega_{n+s,n}$ over a range of values for n and s (see Fig. 1.2) shows that for some values of n and s the sideband frequency approaches zero, meaning that the pulse duration – and thus the time required to cycle between states – approaches infinity. This is clearly impractical for the purpose of cooling ions, so it is necessary to use multiple sideband frequencies over the course of the cooling process. To lower the temperature of the atoms, we want to reduce the motional states of the ions, so we need to change which sideband frequency we use as we decrease the motional states of the ions.

1.4.1 Repumping

Of course, as shown in Fig. 1.5, Rabi flopping only cycles atoms between states but does not cool them – once a π pulse has moved all the ions from the $|\downarrow, n\rangle$ to the $|\uparrow, n-1\rangle$, another pulse will simply move the ions back to the $|\downarrow, n\rangle$, which increases the motional state instead of decreasing it. In order to continue cooling the atoms, we need to be able to change the atoms’ internal states without changing their motional states.

To do this, we use a technique known as repumping. To begin with, we assume that after a π pulse our ions are in the $S_{1/2} |F=1, m_F=1\rangle$ state; we want to transfer them to the $S_{1/2} |F=2, m_F=2\rangle$ state,² where $|F=1\rangle$ is equivalent to the $|\uparrow\rangle$ state and $|F=2\rangle$ is equivalent to the $|\downarrow\rangle$ state in Be^+ (see Fig. 1.3). We use a right circularly polarized 313 nm laser to excite the ions in the $S_{1/2} |F=1, m_F=1\rangle$ up to the $P_{3/2} |F=2, m_F=2\rangle$ state. From there, a given ion spontaneously decays to one of three states:

²We can ignore the motional states of the ions for this discussion – repumping is explicitly designed to leave those states unchanged.

$S_{1/2} |F = 1, m_F = 1\rangle$, $S_{1/2} |F = 2, m_F = 2\rangle$, or $S_{1/2} |F = 2, m_F = 1\rangle$. If it decays back to the first state, which it does with probability $1/2$, it is excited back to the $P_{3/2}$ state again and repeats the process; if it decays to the second state, which it does with probability $1/3$ then we've achieved our goal. But if it reaches the third state, which it does with probability $1/6$, we need a way to return it to the $S_{1/2} |F = 1, m_F = 1\rangle$ state without disturbing the ions in the $S_{1/2} |F = 2, m_F = 2\rangle$ state. The transition from $S_{1/2} |F = 2\rangle$ to $P_{3/2} |F = 3\rangle$ is broad enough ($\Gamma/2\pi = 19.4$ MHz) that if we excite that transition we will inevitably excite ions from the $|m_F = 2\rangle$ as well as the $|m_F = 1\rangle$ states, since they are separated only by Zeeman splitting (on the order of 10 MHz), so we need a different way to move atoms from $S_{1/2} |F = 2, m_F = 1\rangle$ to $S_{1/2} |F = 2, m_F = 1\rangle$.

We do this with the help of a microwave antenna installed near the trap for this purpose. Using this antenna, which emits at a frequency of 1.25 GHz, we are no longer trying to get from $S_{1/2} |F = 2, m_F = 1\rangle$ to $S_{1/2} |F = 1, m_F = 1\rangle$ by going through the $P_{3/2}$ state – a transition path roughly analogous to traveling from Amherst to New York City via the sun – but instead exciting the transition directly from $S_{1/2} |F = 2, m_F = 1\rangle$ to $S_{1/2} |F = 1, m_F = 1\rangle$. This is a far narrower transition, and so we can excite it without having any effect on ions that are already in the $S_{1/2} |F = 2, m_F = 2\rangle$ state. Since we are using the 313 nm laser to continuously excite atoms out of the $S_{1/2} |F = 1, m_F = 1\rangle$ state, if we assume that we begin with n ions in that state then after enough time has passed $2n/3$ ions will be in the $S_{1/2} |F = 2, m_F = 2\rangle$ state and $n/3$ will be in the $S_{1/2} |F = 2, m_F = 1\rangle$ state. If we can use the microwave an-

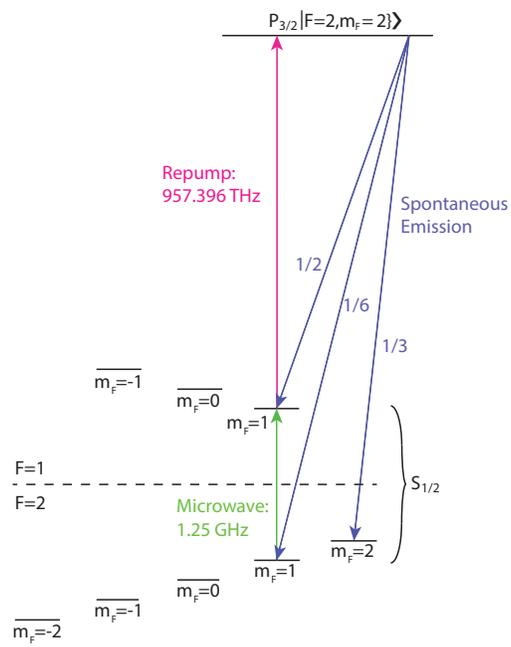


Figure 1.3: The energy levels and transitions involved in repumping.

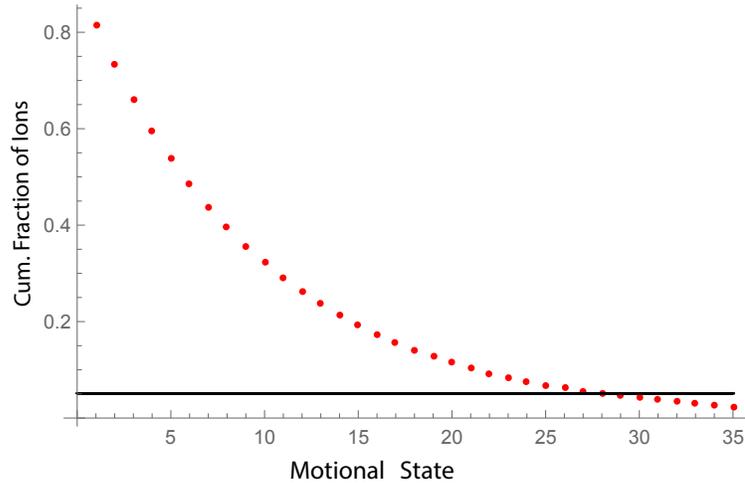


Figure 1.4: A plot of the fraction of the ions that we expect to be at a higher motional state if we have already cooled them to the Doppler limit for ${}^9\text{Be}^+$ in a 1 MHz trap. If we plan to continue cooling all but 5% of the ions, then we should begin resolved sideband cooling on ions in the 27th motional state.

tenna to move all of the $n/3$ ions in the $S_{1/2} |F = 2, m_F = 1\rangle$ state back to the $S_{1/2} |F = 1, m_F = 1\rangle$ state, then after four rounds of microwaving – with continuous repumping – we should have moved more than 80% of the ions into the $S_{1/2} |F = 2, m_F = 2\rangle$ state, and after eight rounds our conversion rate is over 95%.

Now that we can move atoms from one hyperfine state to another with the help of repumping and microwave transitions, it is possible to continually decrease the motional states of our ions and cool them away from the Doppler temperature. If we model the ions' initial state as a Boltzmann distribution around the Doppler limit, then we can determine the proportion of the ions that are in each motional state. The Boltzmann distribution has a long tail, but if we set a threshold for the proportion of ions we want to be able to cool

– say, 95% – then we can solve the equation[14]

$$1 - \frac{\sum_i N_i}{N} = 1 - \frac{\sum_{i=0}^m e^{-E_i/k_b T_D}}{\sum_{j=0}^{\infty} e^{-E_j/k_b T_D}} = .95, \quad (1.5)$$

where N_i is the number of ions in the i^{th} motional state, N is the total number of ions, E_i is the energy of the i^{th} motional state, and m , which we are solving for, is the motional state at which we should begin cooling. In our case, meeting this 95% threshold requires us to begin cooling at the 27th motional state, which, as shown in Fig. 1.2, is most efficient if we use the 4th red sideband frequency.

We begin with a repump sequence, so that virtually all the ions are in the $|\downarrow\rangle$ state, and then use a π pulse to move all of the ions in the $|\downarrow, 27\rangle$ to the $|\uparrow, 26\rangle$ state. This also has the effect of cycling ions in lower motional states, and since the duration of a π pulse depends on the motional state these states will not cycle completely, leaving a mixture of $|\uparrow\rangle$ and $|\downarrow\rangle$ states. However, we then perform another round of repumping, which again restores nearly all the ions to the $|\downarrow\rangle$ state. We continue this process with a π pulse at the 4th red sideband frequency, this one timed specifically to cycle the ions in the 26th motional state, followed by yet another repump sequence. This continues until we have moved all the ions to the 21st motional state, at which point we switch to the 3rd red sideband frequency but otherwise continue with the same process. In theory, repeating this process all the way through the 2nd and 1st red sideband frequencies should cool the ions all the way to the ground state, but it requires a very sophisticated apparatus in order to do so.

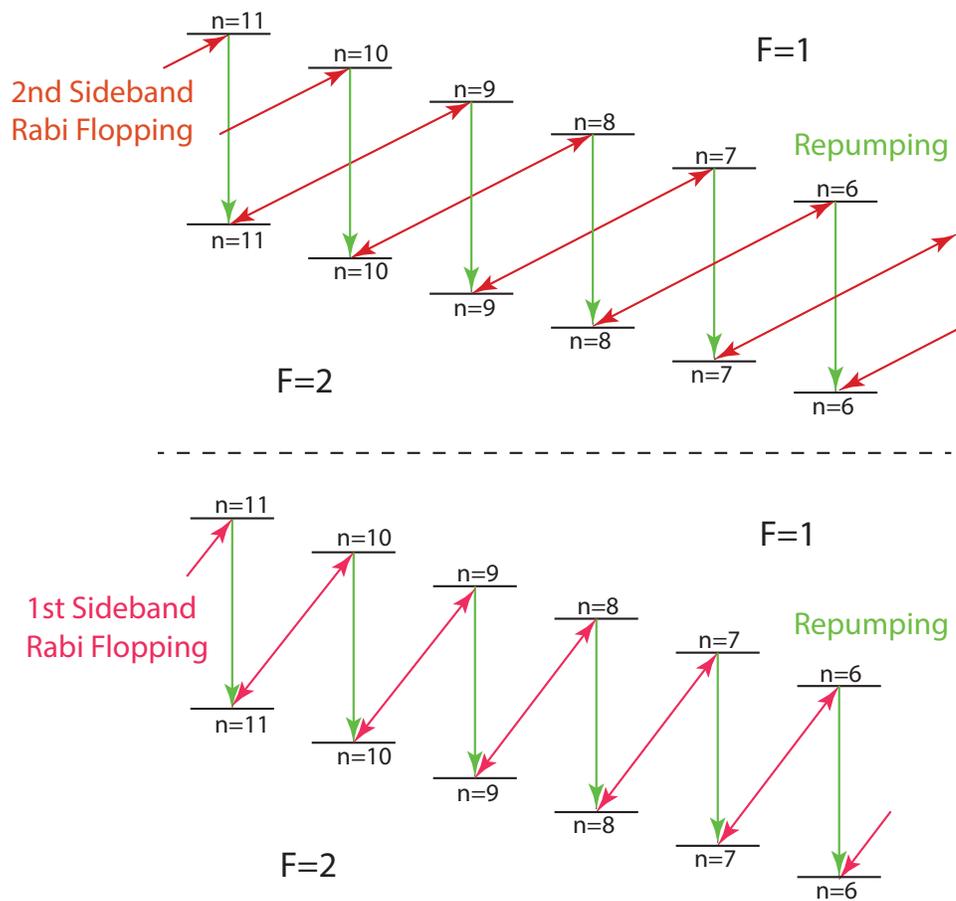


Figure 1.5: A cartoon illustrating transitions in Rabi flopping and repumping for the 11th to 6th motional states, showing both the 1st and 2nd red sideband frequencies.

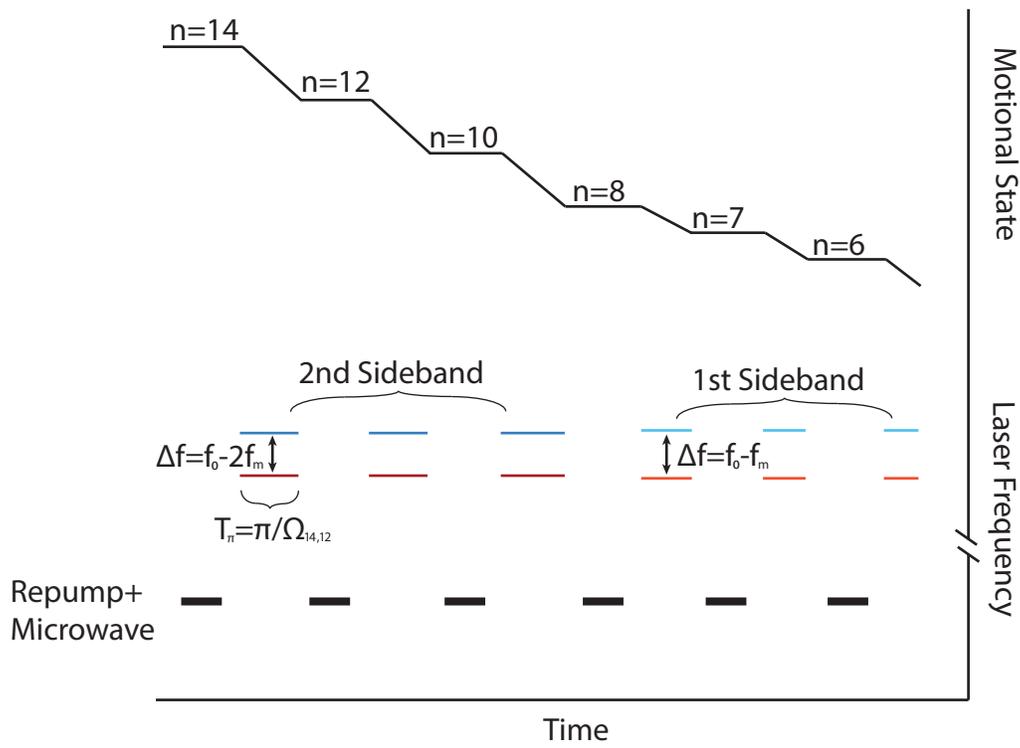


Figure 1.6: A cartoon of the appropriate pulse sequence for cooling a given ion from the 14th to the 6th motional state.

Chapter 2

Lab Overview

The theoretical procedure for trapping and cooling Be^+ ions presented in the last chapter requires a complicated system to be put into practice, but, at the risk of sounding simplistic, these procedures and apparatuses are all designed and constructed primarily to fulfill one of two goals: either to trap the ions or to manipulate their quantum states.

In the first category, the most important pieces of equipment are the vacuum chamber[15] and the ion trap[16]. Using the procedures outlined below, we were able to produce a vacuum that is capable of reaching pressures below $0.1 \mu\text{Pa}$. By running current through a tungsten wire wrapped in beryllium that is inside the vacuum chamber, we can heat the wire and disperse ^9Be atoms into the system. A fraction of these will be hit by the beam from an electron gun, which, somewhat ironically, actually knocks off an electron to produce $^9\text{Be}^+$. These ions are then captured by our linear Paul trap so that we can cool them.

The state manipulation – which is primarily intended to decrease the temperature of the ions – is achieved with a sophisticated laser system designed to perform both Doppler and resolved sideband cooling. The system begins with a 940 nm external cavity diode laser, which, at the most basic level, consists of a semiconductor diode with feedback from a diffraction grating. Current is run through the diode, causing emission of photons with a wavelength of 940 nm. These photons exit the diode and then hit the grating, at which point the first-order mode is reflected back into the diode in order to cause stimulated emission and produce more 940 nm photons, while the zeroth order mode is reflected off as the output. We then use a technique known as frequency summing to turn these 940 nm photons into 313 nm photons[17] suitable for Doppler cooling the Be^+ ions using the $2S_{1/2} \leftrightarrow 2P_{3/2}$ transition. The 940 nm laser is stabilized using a confocal Fabry-Perot cavity[18] in order to maintain a constant frequency over time, while the frequency-summing stage is stabilized by a Pound-Drever-Hall device in order to keep it transmitting at maximum power. To perform resolved sideband cooling, the lasers are controlled using a pulse generation system[19] in tandem with acousto-optic modulators.

2.1 Ultra-High Vacuum

The setting for all of our ion trapping is the vacuum system. It is impossible to keep ions cold and trapped if they are undergoing constant collisions with warm gas molecules, so to reduce these collisions we place the trap in as rarified a vacuum as we can produce. There are two important principles for

maintaining a good vacuum: maximizing the flow of gas out of the vacuum and minimizing the flow into the vacuum. Achieving the first goal is relatively straightforward, as there exist a variety of commercial options designed to maintain what is known as an ultra-high vacuum (UHV) system.

The second goal turns out to be much more complicated. This is not because gas is entering from outside the vacuum system (well-sealed valves and windows have virtually zero leakage) but instead because of outgassing from the metals and ceramics inside the vacuum. This outgassing occurs in part because even the seemingly solid materials that make up the vacuum system have a nonzero vapor pressure – in other words, all the metal components of the enclosure and apparatus have gas molecules within their chemical structure which will evaporate back into the vacuum.¹ We reduced this type of outgassing with careful choice of materials, as outlined by Phyto Kyaw in his thesis. For example, our ion trap is made of special oxygen-free copper, while the vacuum chamber itself is constructed out of stainless steel.[15]

Phyto’s thesis also describes the careful methods necessary to reduce a second variety of outgassing, which comes from gases that have adsorbed onto surfaces inside the vacuum system. These gases will adsorb onto any surface that is exposed to air, but when exposed to UHV at room temperature the gases will evaporate quickly enough to have an appreciable impact on equilibrium pressure but too slowly to have an impact on the total reservoir of adsorbed gas. One way to reduce this method of outgassing is to decrease

¹In reality, gas molecules can enter the vacuum via multiple different means, including evaporation, sublimation, and desorption. We use “evaporation” to refer to all of these processes, since all of these types of outgassing have a similar effect on our ability to achieve UHV conditions and since similar methods are used to minimize them.

the overall surface area exposed to the vacuum, but this can only be taken so far before there isn't enough room in the chamber for all the necessary components. To further decrease the amount of adsorbed gases, we performed a procedure known as a "bake-out," which consists of heating the vacuum chamber up to almost 200° C. This heating drives the gas films, which are composed primarily of hydrogen and water vapor, to evaporate more quickly off of the chamber walls so that the pumps can remove them. We also created a cleaning procedure wherein we washed components with a soap (e.g., Alconox) and then a solvent (e.g., methanol)² to remove any foreign substances that could evaporate into the vacuum – the oil from even a single fingerprint will outgas enough to make reaching UHV conditions impossible.

2.2 The Trap

Our lab employs a linear Paul trap (see Section 1.1) designed by Shenglan Qiao, which can trap both atomic and molecular ions. Since we are interested in trapping specific species of ions (Be^+ and O_2^+) it is important to know how the type of ion we want to confine constrains our design choices. To keep the ions trapped, the trap produces an electric field that confines the ion in one direction while forcing it outwards in another direction; these fields oscillate at a high frequency in order to produce an average confining force on the ion. This results in a stable trapping potential only if the ions can't travel too far in an anti-confining direction before the fields in that direction become confining

²The type of soap and solvent we used depended on what material the component was made of.

and force the ion back toward the center of the trap. Because the magnitude of the force on an ion is related to its charge Q and the response of the ion to that force is inversely proportional to its mass m , the charge-to-mass ratio of the ions we want to trap will set the acceptable values of our trap parameters.

There are of course many possible ways to quantify the geometric and radiofrequency properties of the trap, but by convention we measure the trap parameters as follows:

- r_0 is the distance from the ion to the RF electrodes at their closest point.
- z_0 is half the width of the central DC electrode.
- κ is a geometric property of the trap related to the influence of the trap's endcaps.
- V_0 is the amplitude of the RF drive.
- Ω is the angular frequency of the RF drive.
- U_0 is the voltage difference between the endcaps and the central DC electrodes.

To solve for the values of these parameters necessary for trapping Be^+ , we can analyze the equation of motion of the particle using Mathieu's Equation:[12]

$$\frac{d^2x}{d\xi^2} + (a_x - 2q_x \cos 2\xi)x = 0 \quad (2.1)$$

where

$$\xi = \frac{\Omega t}{2}, \quad a_x = -\frac{4Q\kappa U_0}{\Omega^2 m z_0^2}, \quad \text{and} \quad q_x = -\frac{2QV_0}{\Omega^2 m r_0^2}. \quad (2.2)$$

This applies to motion in the x-direction – motion in the y-direction is governed by virtually identical equations – and only certain values of a_x and q_x (and a_y and q_y) will result in the particle being trapped. Not only that, but we need values that agree with the “pseudopotential approximation,” which is another way of saying that the RF needs to be oscillating fast enough that we can consider the time-averaged potential to be a valid approximation. This is a “pseudopotential” because the potential depth depends on the mass of the particle, and it is valid when $a \ll q^2 \ll 1$. Of these parameters, V_0 , Ω , and U_0 are still variable; we can change them by adjusting the signal we send in to the trap. The geometric parameters, on the other hand, were set by the design choices: $r_0 = 1.25$ mm and $z_0 = 1.5$ mm.[16] κ is also fixed by the design of the trap, but since it is based on the complicated relationship between multiple conductors we were not able to determine its value without observations of trapped ions.

These trap parameters, along with determining whether the ion will stay trapped, also define the oscillation frequencies of the ions in the trap. If we ignore the micromotion of the ions, we have the equations[20]

$$\omega_z = \sqrt{\frac{2\kappa QU_0}{mz_0^2}}, \omega_r = \frac{QV_0}{\sqrt{2}m\Omega r_0^2}, \quad (2.3)$$

where ω_z is the axial angular frequency of the ions – that is, 2π times the frequency of oscillation in the elongate direction of the trap – and ω_r is the radial angular frequency of the ions.

In order to investigate the values of ω_z and κ , we photographed trapped

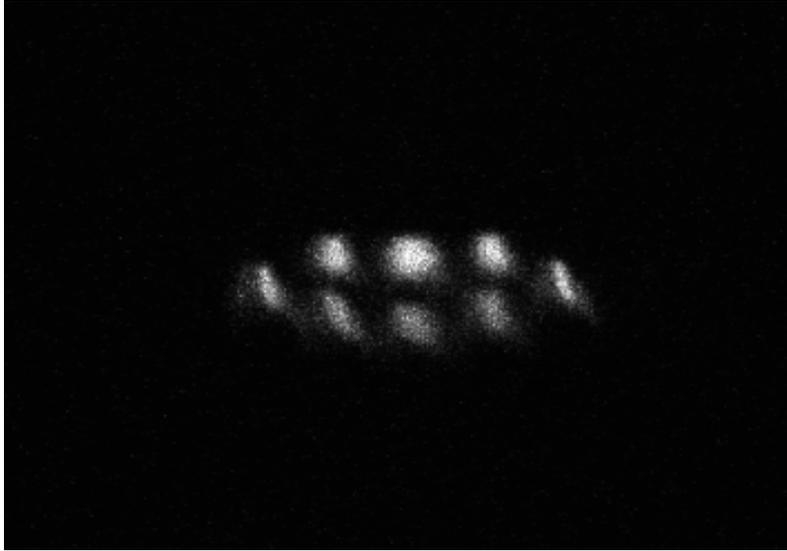


Figure 2.1: A ring of Be^+ ions crystallized in the trap. The ions on either end are separated by $50 \mu\text{m}$.

ions at multiple values of U_0 as shown in Fig. 2.2. By varying U_0 , which is easy to change since it is simply a DC voltage, we can add redundancy to our calculation as well as estimate the uncertainty of our calculations. Since we can measure the N ions' relative positions from the photographs, we can determine ω_z by calculating

$$E_i = E_{quad} + E_{coul} = \frac{m\omega_z^2 z_i^2}{2} + \sum_{j=1, j \neq i}^N \frac{e^2}{4\pi\epsilon_0 |z_i - z_j|}, \quad (2.4)$$

where E_{quad} is the quadratic potential of the trap, E_{coul} is the Coulomb potential of the other ions, and z_i (z_j) is the distance from the i^{th} (j^{th}) ion to the center of the trap. If we minimize the potential energy of the system by setting $\frac{dE_i}{dz_i} = 0$ for all i , we can derive i values for ω_z , which, for a given value of U_0 , should theoretically all be the same. If we then substitute in for ω_z

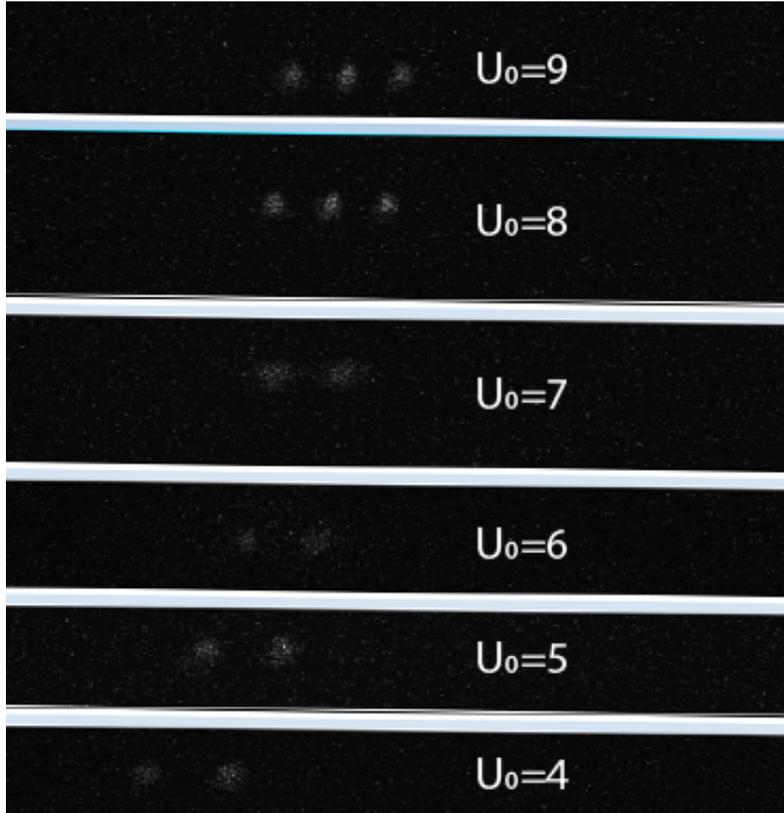


Figure 2.2: Be^+ ions crystallized in the trap as we change the endcap potential from 9 to 4 V. One ion drops out partway through the succession. The distance between the two ions in the lowest panel is about $18 \mu\text{m}$.

using Eq. 2.3, we can back out an experimental value for κ .

This effort is complicated by the fact that the trapped Be^+ ions can pick up a stray hydrogen atom that is floating in the vacuum to form a beryllium hydride ion (BeH^+). These ions remain trapped and continue to influence the others via the Coulomb interaction, but they are not resonant at the same laser frequencies as the Be^+ ions, meaning that they are invisible to the camera. For the calculations to work, we need to estimate the positions of these invisible hydride ions, sometimes without even knowing how many there are.

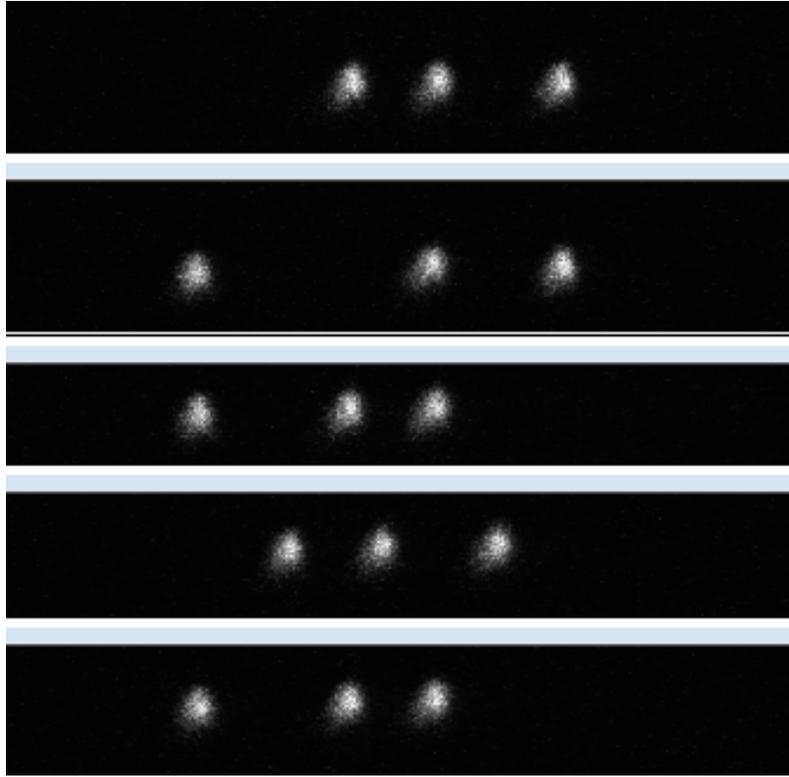


Figure 2.3: Three Be^+ ions crystallized in the trap along with an unknown number of BeH^+ ions. These pictures are taken one after another; the ions are switching positions as a result of collisions with other particles, which gives us a better sense of the positions of the invisible BeH^+ ions. The distance between the endmost visible ions in the lowest panel is about $54 \mu\text{m}$.

We perform these calculations for the two sets of three ions shown at the top of Fig. 2.2. We choose these because they seem to be the clearest examples of ion crystals that do not include any BeH^+ ions. We can measure from this picture that the ions are symmetric around the central atom, which is what we would expect for a simple harmonic potential. This means that it only makes sense to perform the calculation for the first ion, because the positions of the other two ions don't give us any new information about κ . The second ion is at the center of the potential, and therefore E_2 is independent of ω_z and κ , while the third ion will give us identical results to the first, since when we set up Eq. 2.4 to solve for E_3 we will be using identical values of z_i and z_j . We then determine the z_i 's and z_j 's by measuring the relative positions of the ions compared to the total width of the image, which, given that we know the size of the imager (8 mm \times 8 mm) and its magnification (30 \times), represents 267 μm on each side. Using $U_0 = 9$ V, and combining Eqs. 2.3 and 2.4, we have this equation:

$$E_1 = \frac{2e\kappa U_0 z^2}{2z_0^2} + \frac{e^2}{4\pi\epsilon_0 z} + \frac{e^2}{4\pi\epsilon_0(2z)} = 6.4 \times 10^{-13} \text{ J/m}^2 * \kappa z^2 + \frac{3.45 \times 10^{-28} \text{ J m}}{z}, \quad (2.5)$$

where z is the distance from each of the outer ions to the central ion. We then solve the equation $\left(\frac{dE_1}{dz}\right)_{z=1.19 \times 10^{-5} \text{ m}} = 0$ for κ , yielding $\kappa = .1345$. We can do the same thing for the second set of ions, with $U_0 = 8$ V, and derive the value $\kappa = .1136$, which is relatively close to our first answer. We can also find ω_z , which is equal to $3.395 \times 10^6 \text{ rad/sec} = 2\pi \times (.54 \text{ MHz})$ for $U_0 = 9$ and $2.941 \times 10^6 \text{ rad/sec} = 2\pi \times (.47 \text{ MHz})$ for $U_0 = 8$. These values make sense, at

least relative to each other – we would expect the trap to relax somewhat as we decrease the voltage on the endcaps, meaning that the oscillation frequency would decrease.

It is interesting to note in Fig. 2.2 that as we decrease the strength of the endcaps, the ions not only move farther apart as expected but also move across the trap in the same direction. This indicates that some sort of asymmetry is appearing in the trap as we relax the potential that is pushing the center of the well further and further to the left. Unfortunately, the loss of an ion means that we cannot evaluate whether the axial potential itself is developing an asymmetry or whether it remains symmetric about its central point while that point shifts across the trap.

2.2.1 Helical Resonator

As stated above, we can only use the time-average of the potential as a good approximation if we have $a \ll q^2 \ll 1$. The definitions of a and q are shown in Eq. 2.2, and in theory there are many ways that we can adjust the relative values of a and q^2 . However, the values of r_0 , z_0 , and κ are set by the geometry of the trap, and m and Q are invariant properties of the ions. With all those values fixed, we can really only change V_0 and Ω , which are the parameters for the RF potential $V_0 \cos \Omega t$. In particular, to ensure that $a \ll q^2$ we need a value of V_0 that is much greater than U_0 .

To produce this high-voltage RF signal, we originally used a helical resonator designed by Phyto Kyaw that I helped to complete.[15] This resonator is made up of two two copper coils in a copper can. The larger of the two,

called the main coil, is attached to the trap, while the smaller antenna coil is attached to a function generator that serves to drive the resonator. When we supply power through the antenna coil, the entire system resonates within a very narrow band of frequencies. The main coil does so at a higher voltage than the antenna coil, similar to the idea behind a transformer, giving us a high-precision, high-voltage signal to drive the trap. When connected to the trap, our resonator outputs a signal with $\Omega = 2\pi(15.24 \text{ MHz})$ and $V_0 = 1.3 \text{ kV}$. However, Professor Hanneke has recently begun work on a new resonator, since the output frequency from the original resonator was later determined to be too low for our purposes. A higher Ω has multiple benefits, including the ability to perform logical operations on the ions more quickly and the fact that the motional sidebands will be further from the carrier frequency.

2.3 Frequency Tripling

To perform Doppler cooling on ${}^9\text{Be}^+$ ions we need laser light tuned to the frequency of one of the ion's transitions, which in our case is ultraviolet light with a wavelength of 313 nm. However, there are no commercially available lasers that operate at 313 nm; nor are there any that operate at 626 nm, which could be frequency-doubled to get light at 313 nm. Instead, we use a 940 nm laser (in the near infrared) which we frequency-triple using an apparatus designed by Celia Ou as part of her senior thesis.[17]

Frequency tripling is much more difficult than frequency doubling – or even frequency quadrupling, which could be achieved by performing two doublings

– because it requires us to frequency double the 940 nm laser and then sum the resulting 470 nm light with the original 940 nm light to get to 313 nm. Both of these procedures involve sending the light through crystals with a highly nonlinear electric susceptibility χ_e , which relates the polarization of the material to the electric field it is experiencing. This nonlinear susceptibility means that two photons entering the crystal can combine to form a single photon. Since energy is conserved during the process, the resultant photon has frequency equal to the sum of the frequencies of the two original photons. For phase-matching purposes, it is useful for the crystal to also exhibit birefringence, which means that the crystal has different indices of refraction along its crystallographic axes.

Frequency doubling is a special case of frequency summing where the two photons have the same wavelength – in our case, 940 nm – when they strike the crystal and so combine into a photon with wavelength 470 nm. We use a crystal of bismuth borate (BiBO) cut at Brewster’s angle to minimize reflection and maximize power output and placed in a ring cavity to allow the 940 nm light to pass through the crystal many times.

The frequency summing of 470 nm with 940 nm light is somewhat more difficult, since the optical components of the summing apparatus must be able to function with both of the input frequencies as well as the output frequency. Since BiBO is not transparent at 313 nm, we instead use β -barium borate (BBO).

The end result of all this summing is that our 940 nm infrared laser is transformed into a 313 nm light ultraviolet one, but there is a significant cost.

The efficiency of this process is quite low – only about 1% of the input 940 nm light comes out at 313 nm. This is the motivation for the amplification setup that I constructed, since the output of a diode laser produces an order of magnitude less light at 940 nm than is necessary to produce enough light at 313 nm.

2.3.1 Pound-Drever-Hall Lock

For our purposes, it is not enough to merely keep the diode laser at a constant 940 nm, we also need to maintain the stability of the frequency tripling setup – in this case, to make sure that we remain at the highest possible transmission through both the doubling and the summing cavities. We perform this stabilization with a Pound-Drever-Hall (PDH) device, which I constructed during the summer of 2013. In order to lock the cavities on the point of maximum transmission, we need a signal that will allow us to determine what that point is. To produce this signal, we modulate the laser at a specific frequency f_{mod} and then measure the light reflected off of each cavity. With the help of non-linear circuitry, we can turn this reflected light into a signal that has a width in frequency space equal to $2f_{mod}$ as well as distinctive features around the point of maximum transmission to help with cavity stabilization.

For laser modulation purposes, the device produces a very precise $f_{mod} = 14.7456 \text{ MHz}$ ³ using the circuit diagrammed in Fig. 2.4. This circuit is modeled off of a design known as a Colpitts Oscillator, which creates a feedback loop

³A frequency which is in fact somewhat arbitrary – we needed a frequency between 10 and 20 MHz in order to have a signal of the right width, and a 14.7456 MHz crystal oscillator was easy to find.

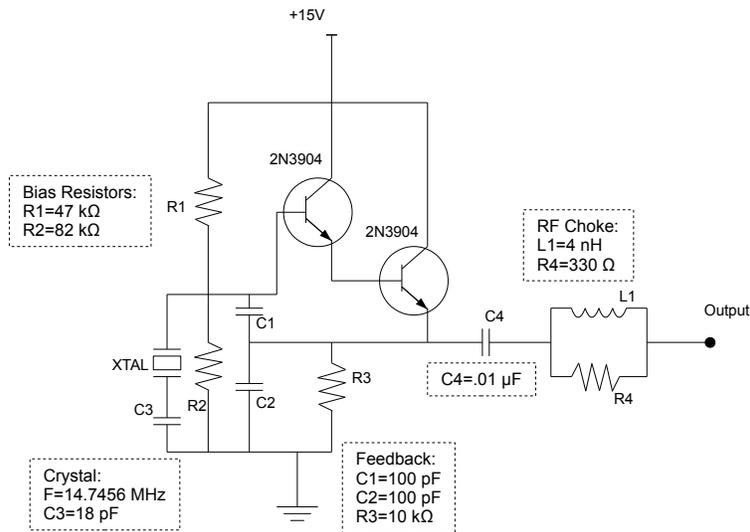


Figure 2.4: A diagram of the oscillator circuit used to provide the 14.7456 MHz signal.

between the base and the emitter of a bipolar junction transistor (BJT), using a pair of capacitors that functions as a bandpass filter. In our case, we use a pair of capacitors in parallel with a crystal designed to oscillate at precisely 14.7456 MHz, which allows us to construct an oscillator with an extremely narrow frequency peak. We also use a pair of BJT's (instead of just one) in a configuration called a Darlington Stage, where the emitter of one transistor feeds into the base of another to provide higher gain. Finally, the output of the the second transistor passes through a radiofrequency choke, which in our case is just a resistor and an inductor connected in parallel to cut off high-frequency signals that are a result of the crystal vibrating at higher order harmonics.

This 14.7 MHz signal is then amplified using an off-the-shelf amplifier from Digikey before being sent into a four-way splitter. One output of the splitter

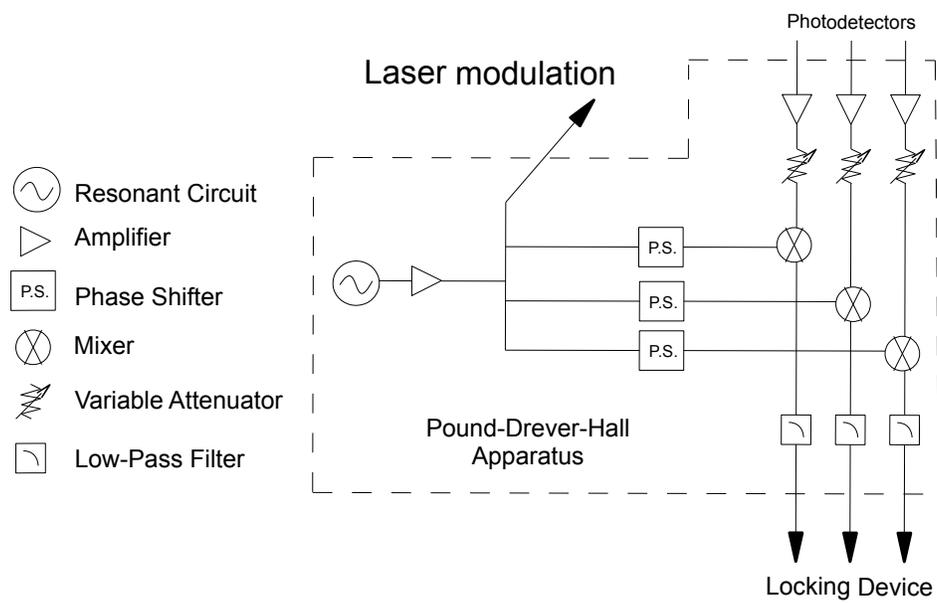


Figure 2.5: A block diagram of the components used in each channel of the PDH.

goes to frequency modulate the laser, which, in frequency-space, has the effect of adding small sidebands at 14.7 MHz. If we use a piezoelectric to scan the length of one of the cavities, and then use a photodetector to measure the beam after it is reflected off of the cavity, we can see a sharp minimum when we scan the cavity across peak transmission, as well as two smaller minima from the sidebands.

Each of the other three outputs of the splitter connects to its own adjustable phase shifter (one of these channels is extra, since the PDH is only being used to stabilize two cavities). These phase shifters compensate for the fact that the distance traveled by the 14.7 MHz signal is different from the distance traveled by the laser – if we didn't correct for this, then the signals would be out of sync and we would just get a mess. Meanwhile, the light reflected off of each of the two cavities (frequency doubling and frequency summing) is captured by two photodetectors, the signals from which each go through an amplifier as well as a variable attenuator, which gives us the ability to control the degree of amplification for each signal.

We then use a mixer to combine each pair of signals – one 14.7 MHz signal with one photodetector signal – to create a new signal, which has two main frequencies: a signal at 29.4 MHz ($14.7+14.7$) and a DC signal ($14-14$). We then put the signal through a low-pass filter so that we can isolate the low-frequency component. If we adjust the phase shifters so that the phases are in sync, and the attenuators so that the signal strengths are the same, then as we scan across the cavity the DC output changes in accordance with the characteristic PDH shape[21] seen in Fig. 2.6. This signal has three useful

properties for locking the cavity:

1. At maximum transmittance the DC signal goes to zero, which is an easy value to lock onto.
2. The slope across maximum transmittance is very steep, so a small deviation leads to a relatively large non-zero output.
3. The saddle-like features are fairly wide – $2f_{mod} = 29.4$ MHz – and on opposite sides of the y-axis, so that the locking mechanism can find its way back to maximum transmission even after a large perturbation.

We feed the two outputs – one for the doubling and one for the summing cavity – into a locking device, which is itself connected to two piezoelectrics that each control the length of one cavity. The locking device adjusts the piezoelectrics in order to maintain a DC signal of zero from each output, and therefore to lock the cavity at maximum transmission.

2.4 Laser Stabilization and Tuning

In order to perform the cooling and state preparation for Be^+ , we need precise control over the frequency of our laser. The first step in achieving this control is ensuring that the laser is stable at certain frequencies over long periods of time, which is necessary since we plan to be measuring minute changes in transition frequencies over the span of many months. This stabilization is achieved with the help of a confocal Fabry-Perot cavity designed and built by Cheyenne Teng.[18] This Fabry-Perot cavity is designed with two highly

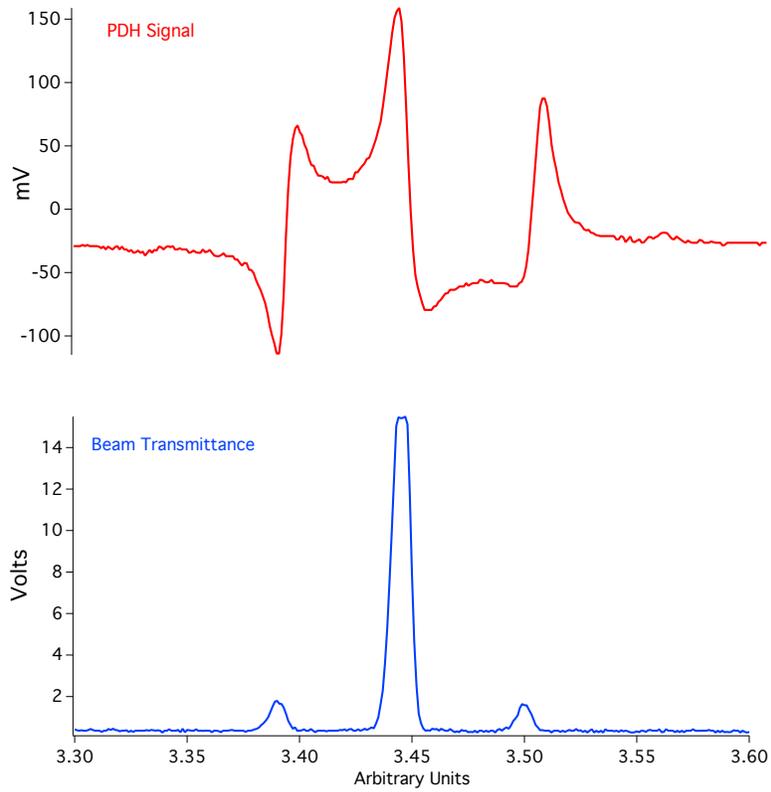


Figure 2.6: The light transmitted through a Fabry-Perot cavity and the resulting Pound-Drever-Hall signal as we scan the length of the cavity using a piezoelectric device. Note that the PDH signal crosses zero at a steep slope at the point of maximum transmission through the cavity.

reflective mirrors ($\sim 97\%$ reflectivity) so that a beam is only transmitted if the cavity length is an integer multiple⁴ of the beam's wavelength.

As a result, if we use a piezoelectric device to change the length of the cavity at a linear rate we should see sharp transmission peaks at regular intervals, which, in theory, should be enough to detect any frequency drift in the laser. However, since the piezoelectric does not have a linear response over a large range, this approach is impractical. Instead, we use a helium-neon (HeNe) laser, which has a frequency that is stable to within 3 MHz over an eight-hour period, as a reference point. If we send the beams from both the HeNe and the 940 nm diode laser through the cavity while simultaneously scanning the piezo over only a few transmission peaks, the gap between the peaks of the HeNe and the diode laser should remain constant if the frequency of the diode laser isn't drifting. So by measuring this gap and adjusting the current in order to keep it constant, we should be able to keep the diode laser as stable as the HeNe – in fact, by using a Thallium Fluoride transition in Professor Hunter's lab as an external reference Cheyenne was able to confirm that her system was able to keep frequency drift under 200 kHz over a three hour period.

Using this Fabry-Perot cavity allows us to stabilize the 940 nm laser, but it doesn't give us the ability to tune its frequency. To do that, we use a wavemeter borrowed from Professor Hunter's lab, which uses a traveling Michelson interferometer to compare the frequency of our beam to that of a HeNe laser. Using the precise readings – good to within about 100 MHz at 940 nm – we can

⁴For the confocal cavity, transmission is actually achieved if $n\lambda = 4L$, where n is an integer and L is the cavity length, since the beam travels from mirror to mirror four times on one pass through the cavity.

shift the frequency of the diode laser (within the very narrow gain spectrum of the diode), which in turn shifts the frequency of the 313 nm output of the frequency-tripling cavity so that it can excite the right transitions in our ions.

This is itself only enough to perform the Doppler cooling, which merely requires a stable laser beam very slightly detuned from the 313.132 nm transition. To perform resolved sideband cooling, we need a laser system that can perform a series of pulses at varying frequencies, all of which are 6.6 GHz away from 313 nm.⁵ This system was partially designed by Jiajun Shi as part of his senior thesis, and it has two main components. The first component, which is a mixture of hardware and software, provides the signals necessary to drive the pulse sequences; these require frequency-switching capabilities on the order of microseconds. The second component consists largely of devices known as acousto-optic modulators (AOM's), which, at the most basic level, are just pieces of quartz attached to piezoelectrics. These piezoelectrics are programmed to oscillate at a specific frequency, driving sound waves through the quartz crystal. A laser beam sent through such an oscillating crystal undergoes a process very similar to Bragg diffraction, resulting in a beam that is shifted in both angle and frequency. The frequency change is important for obvious reasons: using these AOM's, we can use a single laser to create pulses at a variety of different frequencies. It turns out that the angular change is quite useful as well since, with some clever placement of optics, turning the AOM on sends the frequency-shifted beam into the trap while turning the AOM off sends the unshifted beam somewhere else entirely. Thus, from the

⁵The 6.6 GHz detuning can be achieved using the Fabry-Perot cavity and the wavemeter, since these allow us to detune our beam to a stable and precise frequency.

perspective of the ions, we can turn the laser off between pulses by switching off the AOM's without having to switch off the diode itself – a relatively slow process that would cause a host of problems. By using the pulse-generation system to control the AOM's, we should be able to create the pulse sequence shown in Fig. 1.6.

Chapter 3

Laser Amplification

In order to reach higher Rabi frequencies, we need about 20 mW of 313 nm laser power incident out of the frequency tripling stage. Due to the lossy nature of the frequency tripling cavity, this requires a great deal of power at 940 nm. We designed and assembled an amplifying apparatus to provide this power.

3.1 Design

The apparatus was based partially on Professor Hanneke's model of our previous amplifier setup. However, Professor Hanneke suggested a few key ways in which my setup should differ:

- The entire apparatus should be self-contained inside an aluminum box.

We are operating at 940 nm, which is in the near infrared and is therefore quite dangerous – it is outside the visible spectrum, and so can cause retinal

damage without triggering the body's blink reflex. The previous version is kept within an opaque enclosure, but an aluminum box would allow us to keep this version on a mobile breadboard while still allowing people to be in the lab while the lasers are on without having to wear thick orange goggles.

- It should be fiber-coupled on both the input and output sides.

The input of the amplifier is coming from a 940 nm diode laser, and the output of the amplifier will go to the frequency tripling stage – we don't have any actual use for light at 940 nm. Rather than having my apparatus right next to the tripling stage and the diode laser, or having the beams travel long distances across the table, the most practical option is to connect the disparate apparatuses with optical fibers.

- There should be no mirrors on the input or output side, both to reduce the distance traveled by the beam and to reduce the overall degrees of freedom.

Professor Hanneke's previous tapered amplifier design had suffered from low coupling rates into the fiber, perhaps because it had many excess degrees of freedom that made it difficult to find the optimum configuration. To avoid this problem, my original design had no mirrors and used only the controls on two Thorlabs fiberports to control the beam position. Because tilted/tipped mirrors are the easiest way to control the beam path, building the apparatus without mirrors made it absolutely crucial that the optical components all lie along a single axis, since other than the adjustment screws on the fiberport we

have almost no control over the direction of the beam. During the design process, a number of features were included to make sure that all the components would remain along this central axis, including dowel pins on the amplifier mount and cage rods for the lenses.

The entire assembly was configured in SolidWorks – the parts that we needed to machine as well as those that we were purchasing. As a result of the third criterion, virtually all of the parts were centered around a central axis that was supposed to be the beam path.

These SolidWorks files were given to Jim Kubasek, who machined the necessary pieces for us, including:

- Five 1/4-inch-thick aluminum sheets to serve as the sides and top of the box, and one 3/8-inch-thick aluminum baseplate, which screws into the optical table. The front and back sides include holes for the fiberport and for mounting cage rods to hold some of the lenses necessary for beam collimation, while the left side has holes for 9-pin D-Sub connectors, which we use to provide power to the amplifier and to other devices in the box.
- Four brass adjusters, two coarsely-threaded and two finely-threaded. Each finely-threaded adjuster was designed to hold an aspheric lens very close to either the input or output facet of the amplifier.
- Copper holders for the brass adjusters as well as a copper baseplate on which the tapered amplifier rests. These pieces were made out of copper due to its higher thermal conductivity, which allows us to more easily

keep the tapered amplifier at a stable temperature.

3.2 Apparatus

3.2.1 Tapered Amplifier

One common device for laser amplification is the tapered amplifier, which, in essence, is just a wedge-shaped prism of semiconductor material filled with a gain medium. With the application of current, the active region of the semiconductor diode becomes populated both with free electrons and with “holes” – atoms that are missing electrons from their valence shells.[22] The amplifier is seeded with a low-powered beam (in our case, around 10 mW), which enters through the narrow input facet, and, as a result of the presence of electron-hole pairs, when the seed beam passes through the amplifier it causes free electrons to occupy the holes, emitting a photon in the process and increasing the beam’s power. This is the same principle on which diode lasers work, with the notable difference that the tapered amplifier doesn’t employ an optical cavity – in fact, if enough light is reflected back into the amplifier it could permanently damage it, which is an important consideration in its design and operation.

This increase in beam power is one of the motivating factors behind the amplifier’s “taper.” If it had an output facet similar in size to its input facet, the power flux would be nearly 20 MW/m² – potentially enough to melt the glass. However, for the purposes of coupling into a fiber, it is important for the beam to be single-mode. It is helpful to picture this with an analogy

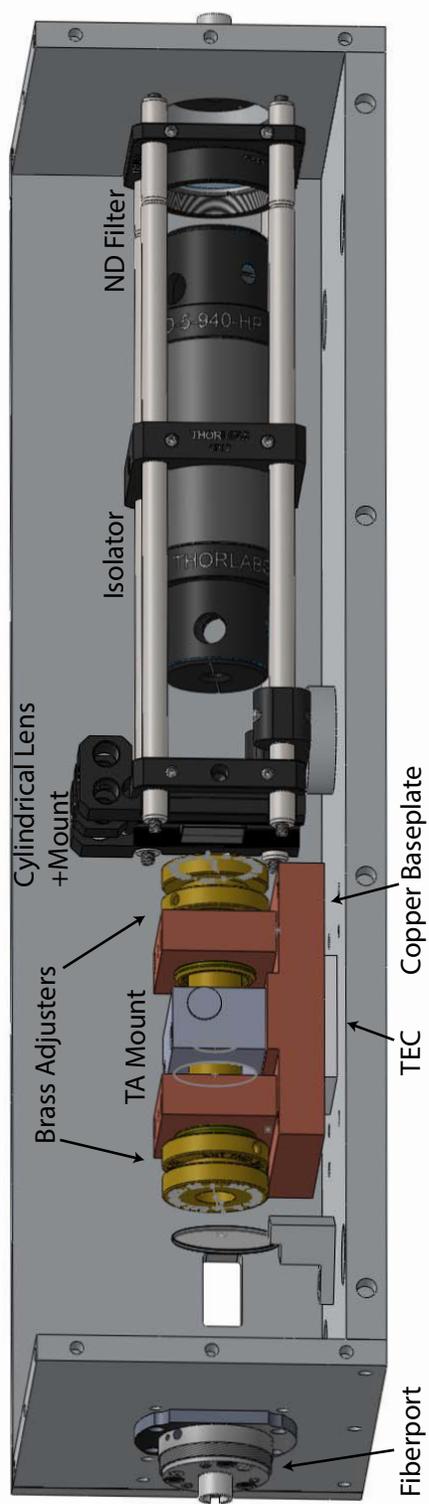


Figure 3.1: SolidWorks design of the original tapered amplifier apparatus. The box is 13.0 inches end-to-end (with ~ 1.5 inches extra from the fiberports).

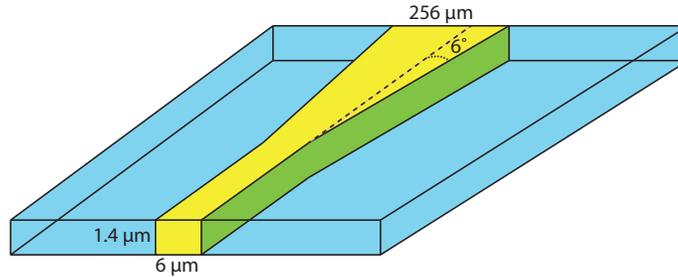


Figure 3.2: A cartoon (not to scale) of the tapered amplifier itself; semiconductor chip is in yellow.

to a particle in a finite well – when the beam enters the narrow part of the amplifier, it is as if it is a particle in a very narrow well and, at low energies, can only occupy the lowest mode. If the amplifier simply stepped out from $6\ \mu\text{m}$ in width to $300\ \mu\text{m}$, this would be the equivalent of suddenly expanding the walls of the well, introducing many higher frequency vibrations. Instead, the amplifier expands in the horizontal direction at an angle of about 6° , which, in principle, is gradual enough not to introduce higher-order modes. Since we are trying to couple into a single-mode fiber, any power in higher-order modes will be reflected off of the fiberport, decreasing our coupling efficiency.

The tapered amplifier itself is just a small chip, but it is housed in a copper mount which measures about $1.61 \times 0.87 \times 0.59$ inches.[23] The mount is fixed to the copper baseplate both with screws and with small dowel rods, since even a tiny rotation would put the beam out of alignment.

3.2.2 Lenses

Our goal in designing this amplification apparatus is to end up with a high-powered laser beam. The tapered amplifier itself takes care of the “high-powered” part, but other measures are necessary to make it a beam. The light exiting the amplifier is constrained to the dimensions of the facet in both the x and y-directions, and so, as a result of diffraction, it diverges in both directions. However, since the facet is significantly wider horizontally than it is vertically, it diverges less in that direction, leading to a phenomenon known as astigmatism. Due to this astigmatism, the vertical beams appear to have a waist at the output facet of the tapered amplifier while the horizontal beams seem to be diverging from a waist inside the chip itself. Thus, while it requires only one lens to focus our seed beam onto the tapered amplifier, we use three different lenses to collimate the output beam.

Laser light enters the apparatus through a 780-HP fiber from our lab’s other tapered amplifier setup. It begins to diverge very quickly from there, but is collimated through the use of a fiberport with a 4 mm focal length lens. The seed beam exits the fiberport as a collimated beam with a radius of about $500\ \mu\text{m}$. It then enters a 3.1 mm focal length aspheric lens, which focuses it down to the $6\ \mu\text{m} \times 1.3\ \mu\text{m}$ input facet.[24] This lens is held in one of the pairs of brass adjusters that can be screwed in or out to better focus the beam. It then enters the amplifier, where it increases in power by up to two orders of magnitude. As it exits the amplifier, it is collimated in the vertical by another 3.1 mm focal length aspheric lens, but as a result of the astigmatism this lens fails to collimate the beam in the horizontal; the beam contracts to a focus in

the horizontal direction soon after the aspheric lens and then begins to expand again. Therefore, we need a lens that only converges the light in one direction: a cylindrical lens.

We originally attempted to collimate the beam in the horizontal using a single 13.7 mm focal length cylindrical lens, but it appears that the amplifier's astigmatism is greater than we expected and as a result we were unable to get the cylindrical lens close enough to the aspheric lens to collimate the beam in the horizontal; under this configuration the beam instead comes to a focus several centimeters in front of the fiberport.

To remedy this, we switched to a 20 mm focal length cylindrical lens. However, if we place this lens 20 mm from the object (which is actually the image of the aspheric lens), then the beam has expanded considerably from where it would be collimated by the 13.7 mm lens, and therefore its waist would be larger than we anticipated. In the abstract, this doesn't sound like the worst problem to have, since we can choose a collimating lens at the output to focus the beam into the fiber. But because this would only affect the beam in the horizontal direction, the end result would be an elliptical beam, causing a decrease in coupling efficiency. As a result, we place the lens more than 20 mm from the objective – 47 mm, to be specific – so that instead of collimating the beam it will begin to focus it. We then place a –20 mm focal length (i.e., concave) cylindrical lens 22 mm behind it in order to collimate the beam with a smaller waist.

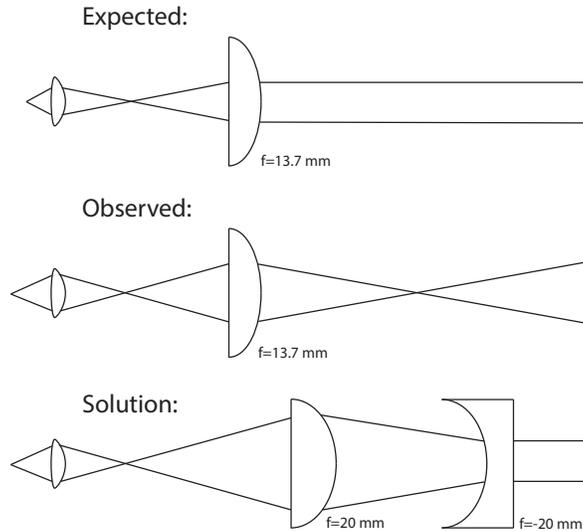


Figure 3.3: A cartoon of the lenses used to collimate the beam in the horizontal.

3.2.3 Other Components

After the lenses, we insert an optical isolator. This device is intended to achieve high transmission of light in one direction while blocking all light in the other direction. In essence, it consists of a tube with a polarizer at each end and a Faraday rotator in the middle. The axis of polarization of the input polarizer is 45° inclined from that of the output, and the magnets in the middle are designed to use the Faraday effect to rotate the polarization of the light by exactly 45° . Thus, correctly polarized light can enter and exit the isolator without any reflection if it is traveling in the right direction, but light that enters through the output of the isolator will rotate 45° in the wrong direction, and as a result will be completely blocked by the input polarizer.

The mirrors are used to capture light reflected off of each polarizer and send it out perpendicular to the beam path. This isolator is necessary in order to prevent reflections off of the fiberport from entering the tapered amplifier in the wrong direction and damaging it. In our setup, we were able to achieve just over 90% transmission of light through the isolator.

Another way to damage the tapered amplifier is to run it at a high current (>2 A) without a seed beam. To prevent this, we use a device known as a “dead seed switch,” which consists of a photodiode, a few op-amps, a switch, and a potentiometer. By placing a window between the input fiberport and the tapered amplifier, we can reflect a small fraction of the seed beam onto the photodiode. When the photodiode is activated to a sufficient degree (what counts as “sufficient” is set by adjusting the potentiometer), the switch is closed. We connect the dead seed switch to the interlock on the current controller, and the result is that it is impossible to turn on the tapered amplifier if there isn’t a seed beam.

During the process of coupling the output beam into a fiber, it is possible to burn the fiber cladding if too much power is sent in improperly. We could avoid this by only coupling the beam at low power, but, inconveniently, the astigmatism of the output beam from the tapered amplifier is correlated to the current setting. So instead, we insert a neutral density filter with an optical depth of 2, meaning that only 1% of the beam’s power passes through on the way to the fiber. This way, the tapered amplifier is outputting more than 750 mW of power, but only 8 mW or so are incident on the fiberport – not enough to burn it. We then removed this filter once we achieved high enough coupling

that we could safely send the entire beam into the fiber.

In order to maintain a consistent beam shape and power, it is best to run the tapered amplifier at a constant temperature. Without some sort of temperature controller, the tapered amplifier would heat up during operation due to the current running through it. To that end, we place a thermo-electric cooler (TEC) under the tapered amplifier mount. A TEC consists of a number of semiconductors connected in series with a ceramic plate above and below; it functions as a heat pump that uses electrical energy to pull thermal energy out of the “cool” plate and transfer it to the “warm” plate. The cool plate is in thermal contact with the tapered amplifier mount, while the warm one is in thermal contact with the aluminum base plate, which is large enough to absorb the power output of the TEC indefinitely without a large temperature increase. In concert with a PID controller and a small thermistor in the copper baseplate, the TEC keeps the temperature of the tapered amplifier stable while it is running.

3.3 Beam Shape

For coupling purposes, it is important to know the beam waist in both the x- and y-directions when it is collimated. We can measure the beam by using a translation stage to move a razor blade across the beam and recording the power at each step. With this method, we gather the data shown in Figs. 3.4 and 3.5. If we attempt to fit a partially integrated two-dimensional Gaussian function to the beam data, we find quite a good fit in the horizontal direction,

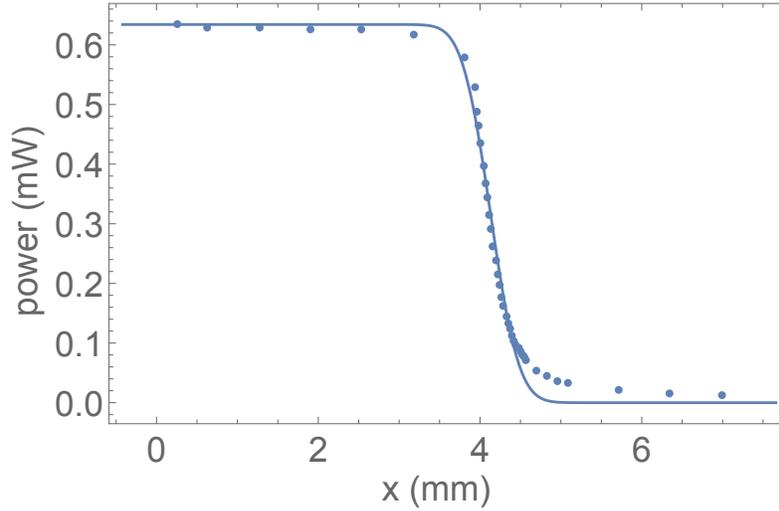


Figure 3.4: Transmitted laser power as we use a razor blade to increasingly block the beam from right to left (from the perspective of one looking along the k -vector of the beam). We fit the data to the function $\frac{P}{2} \left(1 + \text{Erf} \left[-\frac{\sqrt{2}(x-x_0)}{w} \right] \right)$, where P is the full power of the beam, x_0 is the position of the center of the beam, and w is the beam waist.

with a waist of about 0.415 mm. However, in the vertical direction we see more power than expected on the bottom of the beam: this is a phenomenon known as a “coma,” which is a small, lunate tail that travels parallel to the rest of the beam but cannot be focused or collimated along with it. In our setup, we see that the coma represents about 10% of the beam power, which we are unlikely to be able to couple into the fiber since it falls outside the beam’s Gaussian profile. Ignoring the coma, the beam has a waist of about 0.552 mm in the vertical direction, which is fairly close to its horizontal dimension, and the more round the beam the better the expected coupling.

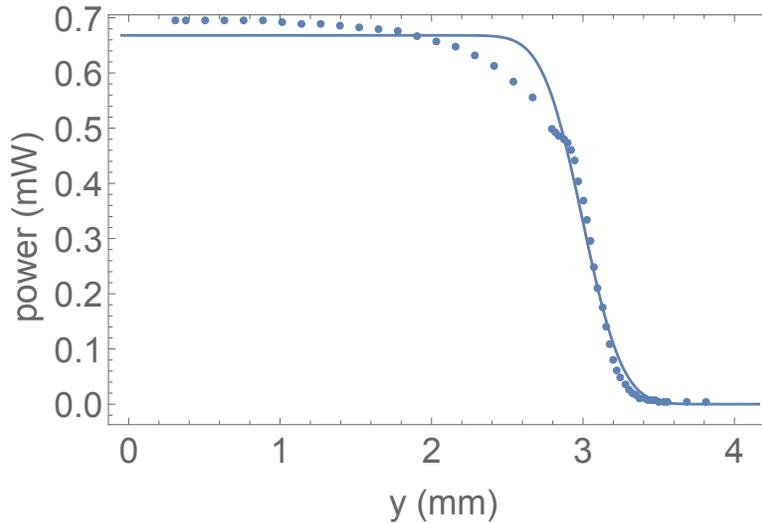


Figure 3.5: Transmitted beam power as we move a razor blade through the beam from bottom to top. Note that we find more power than expected in the bottom portion of the beam.

3.4 Coupling

For the tapered amplifier to provide power to the frequency tripling stage, it has to be coupled on both the input and the output sides. On the front end, this means getting a 10 mW beam focused onto the $6\ \mu\text{m} \times 1.4\ \mu\text{m}$ input facet of the amplifier chip, which requires five things: the beam must have the correct position in two directions, be angled correctly in two directions, and be focused properly. We designed the apparatus with six degrees of freedom, which in theory is more than enough to achieve this: the fiberport has two screws to control the beam position and three screws to control the angle, and we could move the aspheric lens in or out via the brass adjusters.

However, we faced a significant challenge in aligning it due to a complication in the brass adjusters: there was an unexpectedly large amount of slippage

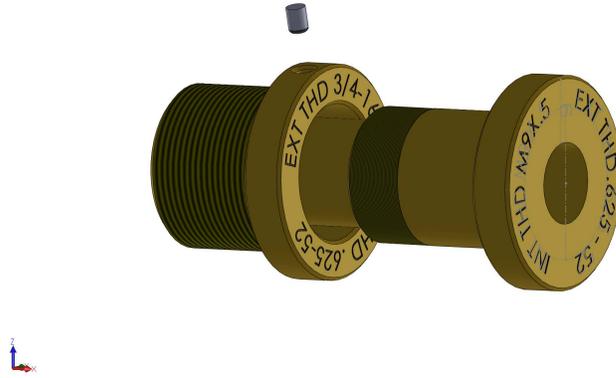


Figure 3.6: Designs for the brass adjusters and set screw.

between the internal threads of the coarse adjuster and the external threads of the fine adjuster, as well as between the external threads of the coarse adjuster and the internal threads of the copper mount. As a result, the input coupling can be changed from very high to virtually zero just by wiggling the internal adjuster (or, for that matter, accidentally bumping it). Ideally, this problem would be solved by tightening down all the set screws, which would prevent the adjusters from moving around. But doing so actually moves the adjusters, so if we get everything aligned with the set screws loose and then tighten them down, the coupling drops to zero. Additionally, for nearly all the configurations of the adjusters, the slippage was significant enough that the angular controls on the fiberport were not able to compensate for the shift in beam angle due to the movement of the aspheric lens.

In the end, we found that the only configuration of the adjusters that allowed for good coupling was with the set screw that coupled the two adjusters screwed in and the coarse adjuster positioned such that the set screw is on the bottom. Since it was impossible to reach the set screw when it was pointed down, it was necessary to screw it in while it was sideways, then turn the adjusters until it was pointed down. This required a great deal of trial and error, since we couldn't know if the alignment would be correct until we rotated the adjusters, and once we did we couldn't make any more adjustments since the set screw was tightened down. In addition, we had to leave the set screw that coupled the coarse adjuster to the copper mount loose, since tightening it would push the system out of alignment while preventing us from adjusting it further. This created vulnerabilities to any accidental bumps, although we were able to mitigate the risk by screwing the coarse adjuster all the way in.

We were able to eventually achieve a reasonable output from the tapered amplifier via this trial and error method. Interestingly, this output did not precisely match the specifications[23] provided to us, as we show in Fig. 3.7. Over a range of coupling efficiencies, we see higher power than expected at low currents and lower power than expected at high currents. We were unable to account for this deviation from the company's specifications, but since we generally ran the tapered amplifier at a current of about 2.5 A, where our output at the highest coupling more or less matched the predicted value, it did not present a significant problem.

When trying to couple the output side, we unfortunately faced the same slippage in the adjusters, and in this case the increased distance from the

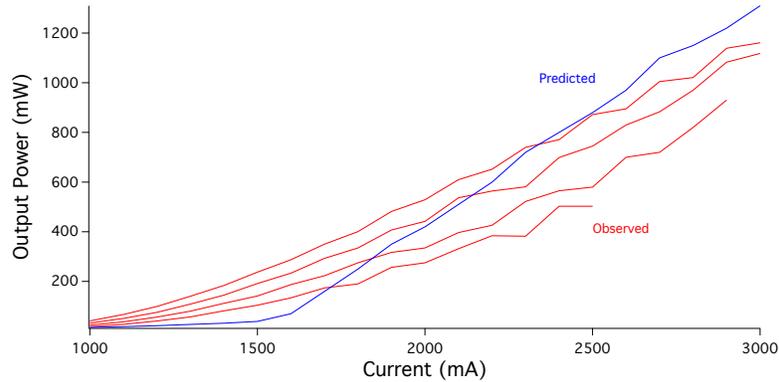


Figure 3.7: Plot of beam power vs. current for a range of observed couplings in red. The values given to us by manufacturer are in blue.

aspheric lens to the fiberport further compounded the problem. Ultimately, we concluded that, despite all our efforts to keep all the components coaxial, it would be impossible to keep the output beam aimed at the fiberport without adding extra degrees of freedom. To do so, we purchased two right-angle mirror mounts which we could add to the cage rod system. These mirrors allowed us to move the beam enough to hit the fiberport, although they unfortunately meant that the apparatus no longer fit inside the aluminum box but instead had to be fixed in some places to the laser table.

At this point, we have achieved about 16% coupling into the output fiber at 2.5 A, which is just over 110 mW. The first measure we plan to take to increase that result is to improve the coupling on the input side of the tapered amplifier. During the alignment process, there were multiple times when we achieved good couplings before one accident or another caused the system to come out of alignment. Since the input coupling is currently only about two thirds of the best value we achieved, we hope that more power out of the

tapered amplifier will lead to more power into the output fiber.

We also plan to experiment with smaller focal length lenses for focusing the output beam into the fiber. Due to incorrect early calculations, we were originally trying to use a fiberport lens designed for a significantly larger beam. With the addition of the mirrors, we are no longer dependent on the fiberport controls and can use a simple collimating lens, of which we have multiple in the lab.

Chapter 4

Conclusion

4.1 Tapered Amplifier Design Notes

By the time of writing this thesis, we were able to get 110 mW out of our fiber-coupled tapered amplifier setup, and we suspect that we will be able to improve this figure by a significant margin. But there are a number of features of the design that could be improved upon in a future iteration.

The most important components to improve are the brass adjusters. The aspheric lenses held in the brass adjusters are crucial for focusing the seed beam onto the input facet of the tapered amplifier and for collimating the output beam in the vertical direction. However, the current design of adjusters have far too much slippage between their threads, with the result that gravity tends to cause them to sag away from the horizontal. In addition, the set screws that we use tend to jostle the adjusters, making it nearly impossible to have the system simultaneously coupled and tightened down.

These two problems were at the root of many of the difficulties I faced in aligning the apparatus, but the difficulties were exacerbated further by our relative lack of control over the beam path. In the original design, a great deal of emphasis was placed on ensuring that all of the optical components could sit along a concentric beam path so that we did not have to include any mirrors for additional degrees of freedom, but the problems in the brass adjusters caused significant deviation away from this path. We had hoped that the Thorlabs fiberports that we purchased would be able to correct for any perturbations away from the central axis of the apparatus, but the fiberports were not as versatile as might be hoped, and were also difficult to work with due to hysteresis in their springs. On the input side, these obstacles forced me to adopt the procedure described in Sec. 3.4 in order to get enough of the seed beam onto the input facet of the chip to stimulate emission. On the output side, we eventually were forced to place mirrors in the beam path just to add extra degrees of control, since even the slight change in angle from the brass adjusters, when multiplied by the longer distance traveled by the output beam, made it impossible to get the beam focused into the fiber using the existing setup.

When we introduced a second cylindrical lens due to the tapered amplifier's higher-than-expected astigmatism (see Sec. 3.2.2), we discovered a slight problem. The extra lens did allow us to focus the beam in the horizontal, but it caused (or at least exacerbated) an optical phenomenon known as a "coma," which adds a small crescent-shaped tail to the focused beam (see Sec. 3.3). We initially suspected that it was caused by the beam entering the cylindrical

lenses off-axis, but when we repositioned the lenses to make sure the beam was centered and perpendicular to the plane of the lens, the beam was no longer pointing at the fiberport. This led us to discover that the beam was exiting the tapered amplifier at the wrong angle, making it impossible to have the beam enter the lenses on-axis and still reach the fiberport. Since this occurred before the introduction of the two mirrors, we had another end-plate machined with an off-center hole for the fiberport so that the coma could be reduced.

However, even after making all the adjustments mentioned above to ensure that the beam is entering the cylindrical lenses at the appropriate angle, we still observe a coma. Since coma is generally the result of incorrect positioning of lenses with respect to the beam path, the coma likely stems to some degree from the output beam of the tapered amplifier entering the aspheric lens at the wrong angle. This is almost certainly another effect of the brass adjusters sagging due to gravity, which causes the lens to be angled slightly upwards.

If I were to begin again and design another tapered amplifier apparatus, there are three main features that I would change:

1. First and foremost, I would switch to having only one brass adjuster. The intention behind having a coarse and a fine adjuster on each side was to provide more freedom in the position of the aspheric lens, but we never needed to use that added range. Meanwhile, the extra slippage from having two adjusters was responsible for a significant portion of the problems we faced.
2. I would switch to an L-shaped design for the box, with a single mirror on the output side to offer some degree of control over the beam path. In our

design, we first attempted to use no mirrors for the sake of compactness and reduced beam path. Then, when we discovered that there were not enough degrees of freedom, we added two mirrors, but I felt that the extra beam control from the second mirror was often a hindrance, as it allowed the beam to be in the correct position with respect to the fiberport while the angle was significantly off. I think that the presence of at least one mirror is a necessary design component, since the lack of any mirrors greatly increases the design's vulnerability to unexpected problems – like, for example, the fact that the output of our tapered amplifier was at a significantly different angle than we expected. With a single mirror, we could get the beam position right with the mirror and then use the angular controls on the fiberport – which tended to be more reliable than the position controls – to make the final corrections.

3. I would further decrease the distance between the input fiberport and the tapered amplifier. Unlike on the output side, where there is a minimum distance necessary to fit the cylindrical lenses and the isolator, the input side is constrained only by how much space is needed to align the beam. But since aligning the seed beam becomes easier as the distance becomes shorter, it is helpful to keep this distance to a minimum.

4.2 Future Work

As with virtually any experimental physics thesis, there were some things that I was not able to accomplish in the time before the thesis was due. Some

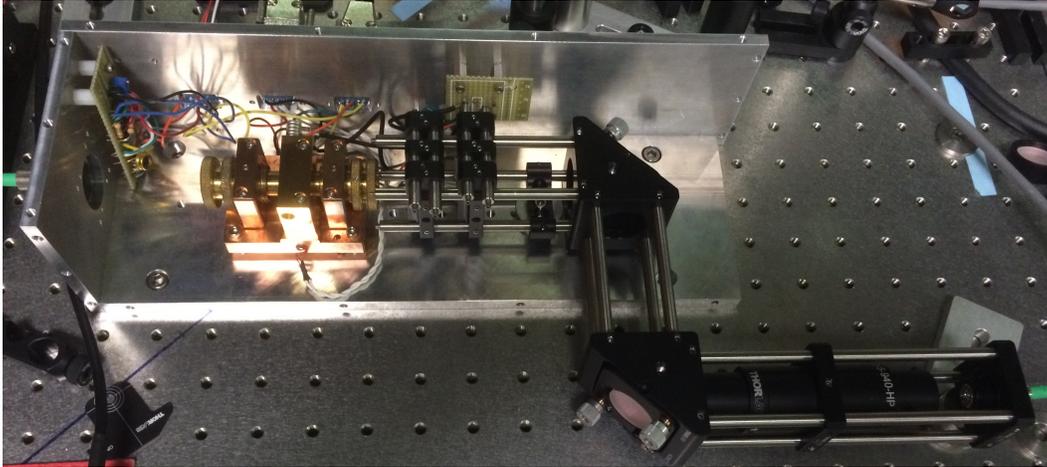


Figure 4.1: A picture of the apparatus in its current form. The holes in the table are 1" apart.

of these I hope to accomplish myself before I graduate, others will almost certainly be left to Professor Hanneke and future students.

4.2.1 Box

As a result of the two added mirrors, we face a new challenge: the apparatus no longer fits in its box (see Fig. 4.1). The extra cage rods mean that none of the original six aluminum sides have the proper dimensions. At some point, a new box will have to be designed and machined, after which all of the components will have to be transferred from the old baseplate to the new. This will also require disconnecting and resoldering many of the wires leading into the box, since they enter through three holes in one of the sideplates. After all this is done, both the input and output sides will need to be realigned, since all the optics will be moved during this process.

4.2.2 Doppler Experiment

Our ability to crystallize ions (see Sec. 2.2) tells us that we are able to use a red-detuned laser to Doppler cool the ions down to a certain temperature. But to actually reach the Doppler limit (see Sec. 1.3), we need to be able to detune the laser by a specific amount: 9.7 MHz, or half of the linewidth of the transition. From a practical standpoint, it is not enough to merely tune the laser to the “correct” frequency using the wavemeter (see Sec. 2.4), since that method is precise to only about 100 MHz. Instead, we need to scan across a range of red-detuned frequencies, taking data on the intensity of emission from the trapped ions at each step. We expect this intensity to steadily increase as we approach 313.132 nm from the red-detuned side, only to suddenly drop as we cross resonance and switch from Doppler cooling to Doppler heating. The point of maximum slope – where the ions’ intensity is increasing the fastest – is 9.7 MHz detuned to the red, exactly where we want to Doppler cool.

4.2.3 Resolved Sideband Cooling

Once we have done all this, we plan to use the additional power from the tapered amplifier device to finally begin to cool the ions below the Doppler temperature. The procedure for this, which is outlined in Sec. 1.3, will be performed using the equipment described in Sec. 2.4. We are still working on the automation of this system, and it has not yet been tested on actual ions, so there is a good deal of work left before it can reliably produce something like the pulse sequence shown in Fig. 1.6. Once we can produce such a sequence, we will theoretically be able to cool ions down to the ground state, which will

pave the way for us to introduce O_2^+ ions into the trap so that we can begin performing QLS experiments.

Appendix A

Frequencies and Detunings

There are a number of different frequencies at which we are running lasers or driving electrodes, as well as detunings from other lasers or transitions.

Frequencies	Value	Region
$2S_{1/2} F = 2\rangle \leftrightarrow 2P_{3/2}$ transition	957.39602(13) THz [10]	UV (313.132 nm)
Doubled light	638.264 THz	Blue (470 nm)
Laser diode & tapered amplifier	319.132 THz	IR (940 nm)
Hyperfine splitting	1.25 GHz	Microwave
Current resonator frequency	15.24 MHz	RF
$2S_{1/2} \leftrightarrow 2P_{3/2}$ Linewidth	19.4 MHz	RF
PDH oscillator frequency	14.7456 MHz	RF
Detunings	Value	Detuned from
Raman transitions (absolute)	6.6 GHz	$2S_{1/2} \leftrightarrow 2P_{3/2}$
Doppler cooling	9.7 MHz	$2S_{1/2} \leftrightarrow 2P_{3/2}$
Raman transitions (relative)	$f_0 + s(f_m)$	red vs. blue beam
Zeeman splitting	$(g_j B \mu_B)/(4h)^1$	adjacent m_F levels

¹where g_j is the Landé g-factor (~ 2), B is the magnetic field, μ_B is the Bohr magneton, and h is the Planck constant. This translates to a splitting of about 700 kHz/gauss.

Bibliography

- [1] M. W. Ray, E. Ruokokoski, K. Tiurev, M. Möttönen, and D. S. Hall, “Observation of Isolated Monopoles in a Quantum Field,” *Science* **505**, 544 (2015).
- [2] M. J. Jensen, T. Hasegawa, J. J. Bollinger, and D. H. E. Dubin, “Rapid Heating of a Strongly Coupled Plasma near the Solid-Liquid Phase Transition,” *Phys. Rev. Lett.* **94**, 025001 (2005).
- [3] S. G. Crane, S. J. Brice, A. Goldschmidt, R. Guckert, A. Hime, J. J. Kitten, D. J. Vieira, and X. Zhao, “Parity Violation Observed in the Beta Decay of Magnetically Trapped ^{82}Rb Atoms,” *Phys. Rev. Lett.* **86**, 2967 (2001).
- [4] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, “14-Qubit Entanglement: Creation and Coherence,” *Phys. Rev. Lett.* **106**, 130506 (2011).
- [5] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, “Spectroscopy Using Quantum Logic,” *Science* **309**, 749 (2005).

- [6] F. Wolf, Y. Wan, J. C. Heip, F. Gebert, C. Shi, and P. O. Schmidt, “Non-Destructive State Detection for Quantum Logic Spectroscopy of Molecular Ions,” *Nature* **530**, 457 (2016).
- [7] D. J. Griffiths, *Introduction to Electrodynamics* (Pearson, 2013).
- [8] F. G. Major, V. N. Gheorghe, and G. Werth, *Charged Particle Traps: Physics and Techniques of Charged Particle Field Confinement* (Springer, 2005).
- [9] P. K. Ghosh, *Ion Traps* (Oxford University Press, 1996).
- [10] J. J. Bollinger, J. S. Wells, D. J. Wineland, and W. M. Itano, “Hyperfine Structure of the $2p^2P_{1/2}$ state in $^9\text{Be}^+$,” *Phys. Rev. A* **31**, 2711 (1985).
- [11] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer-Verlag New York, 1999).
- [12] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, “Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions,” *J. Res. Natl. Inst. Stand. Tech.* **103**, 259 (1998).
- [13] D. Leibfried, “Quantum state preparation and control of single molecular ions,” *New J. Phys.* **14**, 023029 (2012).
- [14] D. V. Schroeder, *An Introduction to Thermal Physics* (Addison Wesley Longman, 2000).

- [15] P. Kyaw, *Constructing an Ultra-High Vacuum Chamber and a Radio Frequency Helical Resonator for Trapping Ions*, Undergraduate honors thesis, Amherst College (2014).
- [16] S. Qiao, *Constructing a Linear Paul Trap System for Measuring Time-variation of the Electron-Proton Mass Ratio*, Undergraduate honors thesis, Amherst College (2013).
- [17] C. Ou, *Third Harmonic Conversion*, Undergraduate honors thesis, Amherst College (2013).
- [18] C. Teng, *Frequency Control and Stabilization of a Laser System*, Undergraduate honors thesis, Amherst College (2013).
- [19] J. Shi, *Radiofrequency Synthesis System for Laser Modulation*, Undergraduate honors thesis, Amherst College (2014).
- [20] M. G. Raizen, J. M. Gilligan, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Ionic Crystals in a Linear Paul Trap," *Phys. Rev. A* **45**, 6493 (1992).
- [21] E. D. Black, "An introduction to Pound-Drever-Hall laser frequency stabilization," *Am. J. Phys.* **69**, 79 (2000).
- [22] W. W. Chow and S. W. Koch, *Semiconductor Laser Fundamentals* (Springer, 1999).
- [23] m2k Laser, *Tapered Amplifier @ 0950 nm for MOPA* (2011).
- [24] A. Schildecker, private communication.