Radiofrequency circuit design for ion trapping of O_2^+ molecules

Will Henshon

Advisor: Professor David Hanneke
May 10, 2023

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

Abstract

The Standard Model forms the basis for our understanding of the universe's interactions. It has proven extremely effective for a variety of predictions, but falls short of explaining key phenomena such as dark matter and matter-antimatter asymmetry. Proposed extensions to the Standard Model that address these issues often involve the introduction of dark matter particles, which couple to regular matter; these may cause oscillations and drifts in fundamental constants such as μ, the proton-to-electron mass ratio. These variations can be measured using an atomic or molecular clock. In particular, molecular clocks using vibrational and/or rotational transitions are sensitive to variations in μ. We seek to create a molecular clock using a vibrational transition in O_2^+ molecules to probe for oscillations in μ. To do this, we use an ion trap inside a vacuum chamber to create a system in which oxygen molecules can be examined for extended periods of time. This thesis describes the design and construction of a circuit to control the voltages on ion-trap electrodes. The circuit allows for radiofrequency (RF) driving potentials on electrodes to trap ions and for precise positioning of ions by shifting the RF null point. Additionally, it can support DC offsets of electrodes to constrain ion position along the trap axis and to counteract asymmetries, as well as AC modulation to detect the motional frequencies of the ions. The electronics were used to trap Be^+ ions in crystallized form, and were effective in moving the ions with distances up to 180 microns, over 10% of the ion–electrode distance. Additionally, the axial and radial frequencies of the Be^+ ions were measured using all three spatial axes. This success in trapping ions means that the electronics look to be effective for positioning and maintaining O_2^+ ions for the measurement of variations in μ.

Acknowledgments

I would like to first thank my advisor, Professor Hanneke, for all his help with this process. He has been a central part of my physics experience at Amherst, from my first intro class all the way to now, and I am always reassured by his continual faith in me and my work. He has been a truly amazing advisor, always willing to answer all of my questions, no matter how trivial, to drop whatever he was doing to come help me troubleshoot, and to spend long days in the lab with me over spring break to try to get the trap to work. I could not have completed this project without him, and I am honored to be his thesis student.

Next, I would like to thank my mom and dad, for their love and support during my thesis, and otherwise. Both were always willing to hear me ramble about whatever was going on the lab, no matter how incoherent, and could always make sure that my spirits were lifted, even when it looked like my work might not come together. I am so thankful and lucky to have such great and supportive parents.

Thank you to Brian Crepeau, for all of his help with my electronics this year. Brian took so much of his time to help me troubleshoot and repair parts of my circuit, and made my life so much easier throughout my process

Thank you to Jim Kubasek for his help with in the machine shop. Jim was super understanding and patient with all of my design work, and did an incredible job on all of the parts we needed to put together the vacuum chamber in January.

Thank you to the other students in the lab this year, who made the lab such a fun place to be. I also want to thank the thesis students who came before me the lab, whose work I used frequently and was vital for me to succeed in my own.

Thank you to my class of fellow physics majors, who have made this year and my time at Amherst so much more fun and rewarding. Our game nights gave an excellent break from the time spent in the lab, even if my Wizard skills were fully not up to par this year. Now that my thesis is done, I'll have a lot of time to improve.

Thank you to my brothers, Mark and James, for always being around to talk and take my mind off physics and school in general.

Thank you to all of my family. I'm so fortunate to have so many great and supportive family members who are always there for me whenever I need it.

Thank you to the rest of the physics department, both the faculty and the students, for all their support, mentorship, and friendship throughout my time at Amherst. I am so glad I was able to meet so many great people through physics during my time here.

Finally, thank you to all the friends I've made at Amherst. I have had such an incredible time here, and it would not have been the same without all of the fun and interesting people I've been able to spend time with over the last four years.

This thesis is based upon work supported by the Amherst College Provost and Dean of the Faculty and by the National Science Foundation under grants PHY-2207623 (RUI PM) and PHY-1806223 (RUI).

Contents

List of Figures

Chapter 1

Introduction

1.1 The Standard Model and Beyond

The Standard Model of particle physics is currently the most accurate tool we have for understanding the fundamental interactions that control how the universe works. The model explains how quarks and leptons, which make up matter, and gauge and scalar bosons, which mediate the fundamental forces, build our observable universe. The Standard Model is incredibly precise, and has made many predictions with extreme accuracy. These include the existence of the Higgs boson, which was discovered in 2012, over 50 years after it was theorized [1], and various high-precision measurements, including the electron magnetic dipole moment [2].

Despite these successes, the Standard Model is not perfect. Notably, there are several key mysteries it fails to explain. Gravity currently cannot be explained using quantum field theory, which is the Standard Model's basis for the understanding of each of the other forces and their unification. Since the universe exists, we know that there is more matter than antimatter; however, the Standard Model provides insufficient mechanisms for processes that produce such an imbalance. Additionally, dark matter and dark energy, which together make up over 95% of the mass of the universe based on cosmological observations of gravitational

effects and universal expansion, are made of particles yet to be described by Standard Model physics [3]. Therefore, one of the main goals of modern physics research is to provide explanations for these effects.

Theorists provide a variety of explanations for these puzzles. These include various new particles (which may be classified as dark matter), which induce new fields and interactions that cause measurable effects on regular matter. One of the most well known of the theories is supersymmetry (SUSY), which predicts a series of weakly interacting heavier particles associated with each of the Standard Model particles. Although SUSY was seen by theorists as being an elegant solution for several of the problems and characteristics associated with Standard Model physics, researchers have seen no evidence of supersymmetric particles at the Large Hadron Collider (LHC), despite exploring much of the parameter space in which we would expect to see these particles [4]. Indeed, more generally, weakly-interacting massive particles, or WIMPs, are a very well-supported theory for dark matter. SUSY falls under the broader scope of WIMP dark matter, alongside other models. Many of these models can be referred to as being "natural," a common principle in particle physics theories. Naturalness in particle physics is the rule that certain ratios of constants of nature should not be finely tuned; that is, they should be on the order of magnitude of unity. Although this constraint seems relatively arbitrary, many theorists argue that it is supported by much of our current understanding of the universe. The idea is that the contributions to a parameter should all be on the same order of magnitude, because if they are not, there needs to be very precise tuning to get the desired cancellations [5]. Naturalness creates difficulty in a few instances, including with the mass of the Higgs boson, in what is known as the hierarchy problem. Essentially, the mass of the Higgs boson is on a scale well below what natural theories would suggest, and requires tuning of several coefficients to get cancellations necessary to see the observed properties. Natural theories with WIMPs, including supersymmetry, provide an elegant solution, because the interactions would explain the Higgs's mass scale while still staying with the constraints of naturalness [4]. However, as stated above, these theories

have lost popularity as more and more of the parameter space occupied by natural WIMP theories has been ruled out by experiments at the LHC and elsewhere. This has led some theorists to propose that we are reaching a "post-naturalness era, in which we will have to abandon the elegance of natural theories in favor of less-nice ones that are based on our observations" [6].

If natural WIMPs are ruled out, there are a variety of other extensions to the Standard Model that would explain dark matter and/or quantum gravity. Some, like axions (ALPs), introduce an additional light particle that would couple to Standard Model particles. ALPs could exist as excitations of higher dimensions in string theory, and, similarly to natural WIMPs, would explain the hierarchy problem [7]. Axion dark matter, if it exists, may even facilitate the existence of a multiverse, in which our universe sits in a set of universes, each with different fundamental physical parameters [5]. Moreover, non-natural WIMPs could make up dark matter, leaving the hierarchy problem unsolved; a variety of other theories exist that would solve various issues with the Standard Model in parameter spaces that have not been explored by modern experiments, such as nonthermal, high-mass WIMPzillas [8] and superfluid dark matter [4, 9]. One interesting theory that has recently gained steam is that of primordial black holes (PBHs). PBHs, which would be formed from density fluctuations prior to the Big Bang, could conceivably form a basis for the gravitational observations seen from dark matter without constraints characteristic of ordinary black holes. This theory became more popularized recently because, as argued in ref. [10], LIGO's possible detection of black hole mergers would be consistent with the rate at which we may expect PBHs to merge (it is worth noting that the range of this rate is extremely wide).

1.2 Oscillations and Drifts in Fundamental Constants

Many theories of beyond-the-Standard-Model physics involve additional spatial dimensions or scalar fields that would explain gravitational effects associated with dark matter and

other gaps in the Standard Model. These theories, which include string theories and light dark matter, would mean that fundamental constants, such as α, the electromagnetic fine structure constant, or μ, the proton-to-electron mass ratio, would shift or oscillate over time [11, 12]. Generally, in these theories, some mechanism causes slight changes in the strength of fundamental forces, which would cause slightly differing but correlated effects on α and μ. The models can most generally be separated into two separate categories, as argued in ref. [13]. First, the new field could be massless, which would cause the constants to drift over long timescales associated with cosmological evolution. Alternatively, the new field could be massive, which would cause oscillations in fundamental constants. A massive field would be made up of bosonic matter[1], and we assume it is lightweight enough that it can be treated as a classical wave (using its de Broglie wavelength).[2] The oscillation occurs at the Compton frequency, $f = mc^2/h$, where m is the mass of the theorized particle, and the oscillation amplitude would relate to the energy density of the field as well as the mass of the particle [13, 14]. This energy density is the upper bound on the density of the field, since it cannot exceed the dark matter energy density of 0.4 GeV/cm^3 [12]. Because the dark matter field would have an oscillation amplitude, coupling with Standard Model particles would introduce corresponding oscillations to fundamental constants, with a fixed frequency.

There are wide bounds on the potential masses: in particular, particles with mass well below 1 eV/c^2 are of interest because they can be treated as coherent classical waves, which is important for our measurements [14]. The de Broglie wavelengths for the lower end of the mass range are limited by the need for the particles to fit inside the halos of dwarf galaxies, in order to stay consistent with cosmological observations (these particles would have mass on the order of 10^{-21} eV/c^2). Using these bounds, the frequencies of oscillation in the field can range from 10^{-7} up to 10^{13} Hz [13]. In terms of naturalness, only a small portion of this range is considered to be viable in a natural theory: although specifics depend on the

[1]Fermionic matter would not be possible for this type of field because the Pauli exclusion principle would prevent there from being enough states in galaxies for the dark matter to exist.

[2]This assumption determines how we are able to detect the effects of these particles. At heavier masses, the experiments to find evidence of these particles would look much different than the ones we consider.

coupling parameter in question, particles lighter than 10^{-5} eV/c^2 are forbidden in a natural theory [12]. Particles outside of this naturalness bound then make up a large portion of the parameter space in the oscillation and drift theories. Although these types of particles may lead to "less-nice" theories, it is important that they are probed in experiments like ours, as it seems increasingly likely that new physics lies outside of the natural regime.

1.3 Testing Theories

Effects occurring from various theories of dark matter, while significant, are below the bounds of our current measurements. To test these theories, then, we can improve the bounds on measurements, which will either show that certain theories (or types of theories) are potentially correct, or place new bounds on proposed effects, which may rule theories out as the bounds shrink.

Many measurements of Standard Model physics require extremely high-energy environments. These can be created in places such as particle accelerators (particularly the LHC, referenced above), where particles are crashed together at nearly the speed of light. This can allow for direct detection of new heavy particles, which is how the Higgs boson was discovered. However, this is not the only way to discover evidence of new physics. Atomic, molecular, and optical (AMO) physics uses lower-energy techniques, often with lasers and various types of spectroscopy, in order to make conclusions about general physics effects. The idea of this subfield of physics research is that atoms and molecules provide compact, identical systems within which fundamental particles can be investigated for evidence of new physics. With these systems, we can make precision measurements that will apply to the universe in general, beyond just the atoms and molecules used in the experiments. These types of experiments are especially useful for detecting particles and physics at low energy ranges. Since our technology is not sufficient to directly detect ultralight particles that exist at these low energies, we need to look for indirect effects of or evidence for these particles [14]. These

may include the massive scalar fields that would cause oscillations in fundamental constants.

1.4 Current Limits

As stated above, several theories with significant traction suggest that fundamental constants may oscillate or drift with time. The current limits on these changes come from a variety of experiments. Ref. [12] argues that a massive scalar field has both an oscillating solution, which is characteristic of dark matter, and a Yukawa solution, which is characteristic of ordinary matter, and would therefore look like a fifth fundamental force. Experiments searching for additional fundamental forces are commonplace in AMO physics, and, in general, universal free-fall measurements are powerful for detecting a fifth force of this type. The experiments function by measuring the free-fall acceleration rates of two test masses in identical conditions, either using chemically different test masses that may interact with a scalar field differently, or by varying the spatial position of identical masses. The Einstein Equivalence Principle (EEP) says that objects in free fall should behave exactly the same (when accounting for whether their reference frames are inertial), so any differences in free-fall acceleration must be accounted for by new physics, possibly a new scalar field. It is interesting to note that depending on the desired sensitivity and parameter space, it may be advantageous to make these measurements on the surface of the Earth, or in space, either in a circular or eccentric orbit [12]. While these measurements constrain the strength of a potential fifth fundamental force, they can also provide bounds for oscillations in fundamental constants, since the oscillations would also have an effect on these free-fall measurements. In particular, completed fifth force measurements set strong bounds on the parameter space containing higher-frequency oscillations, so there is more room to improve the current bounds on lower-frequency oscillations (on the order of Hz or mHz).

In our testing of the variation of these fields, we measure dimensionless fundamental constants, such as μ or α, as opposed to other constants with dimensions, because they

are values that depend purely on physics, not arbitrary definitions. As argued in [7], any measurement of a constant with dimensions may describe a new effect in the system being measured, but also may describe the system that defines the dimension, making it difficult to determine where the observed effects lie. Dimensionless constants, on then other hand, are direct measurements of forces and systems, which confines our measurements to those systems. So, using dimensionless constants, if we can improve measurements so that they are sensitive on the order of magnitude of theorized new effects, we may be able to either provide evidence for the existence of the new particles or forces, or rule them out in favor of theories that can support the new results.

Currently, atomic clocks have made the most precise measurements in the drifts of fundamental constants by measuring shifts in spectral lines of atoms; the physics behind these measurements will be explained below. These measurements generally come in the form of $\dot{\mu}/\mu$, which is the fractional rate of change in μ, usually per year. The current best bound on the drift in μ was found using a Yb^+ clock compared with a Cs clock, constraining μ to

$$\dot{\mu}/\mu = (-0.8 \pm 3.6)\text{x}10^{-17}\text{yr}^{-1}, \tag{1.1}$$

while the best current bounds on oscillations of μ have been found by using data from universality-of-free-fall experiments [13]. Studies exploring the possibility of using atomic and molecular clocks to probe these bounds have demonstrated the feasibility of such measurements [15], but currently no results approach the EEP bounds. The constraints that the EEP experiments provide change depending on the mass of the theorized scalar field, and are detailed in ref. [13]. It is also worth noting that some theories propose that dark matter would be clumped around and in large masses of ordinary matter, such as the Earth or the Sun. Since this would increase the density of dark matter on Earth, where our experiments take place, this would make it easier to reach more accurate bounds [13].

In addition to atomic clocks, other techniques can be used to constrain drifts based on

evolving cosmological effects. For example, consider the case of Oklo, a town in the Gabon republic with an open-pit uranium mine [11]. About 1.8 billion years ago, a natural nuclear reactor there went critical, causing a nuclear fission chain. By solving the nuclear reaction chain, and then measuring the concentrations of various isotopes of elements in the mine now, we can get a cross section of neutron fluxes, which is important because the capture resonance of a neutron is the primary factor in the cross section of this reaction. Today, this resonance exists at $E_r = 97.3$ MeV; the idea of this investigation is that the energy may have differed at the time of this reaction. E_r can then be related to the electromagnetic fine structure constant, α, by steps detailed in ref. [11].[3] A variety of analyses have been carried out of this data; the best constraint found so far has been

$$\frac{\Delta\alpha}{\alpha} < (2-8)\text{x}10^{-10}. \tag{1.2}$$

The range in this bound is due to various modelling techniques of the strong force in the analysis. It is worth noting that the bounds here will likely continue to improve as physicists develop a stronger understanding of fundamental constant dependence in the nucleus. Although these bounds are significantly weaker than the ones presented by atomic clocks, they are important to consider because they are constrained over geologic time scales, many orders of magnitude greater than the ones in laboratory experiments. This means that on a per year basis, the bound looks closer to $\frac{\dot{\alpha}}{\alpha} <\sim 10^{-19}$, which is more stringent than laboratory results. Geologic data from meteorites with ages similar to that of our solar system has also been used to place bounds on variations in α [11].

Astronomical data also provide bounds for drifts in fundamental constants. This data extends even further back than the geologic time scales in the Oklo and meteorite data, and functions by tracking the redshift of spectral lines from atoms and molecules. In particular,

[3]While this measurement (and the next one) are made in terms of α instead of μ, since effects on both constants are often correlated, the measurements provide bounds on similar parameter spaces.

the cosmological redshift, z, when measuring light from the universe's past, is defined as

$$z = \frac{\lambda_{\text{lab}} - \lambda}{\lambda}, \tag{1.3}$$

where λ is the actual wavelength of the emitted light and λ_{lab} is the measured value on Earth [7]. This allows us to compare spectral lines of atoms and molecules in the current universe to the past universe. In particular, looking at molecular spectral lines allows us to probe for μ. Currently, the best bound on μ using this method was found using H_2 molecules, and is

$$|\frac{\Delta\mu}{\mu}| \leq 5\text{x}10^{-6}, \tag{1.4}$$

where the time scales are in the range of 10-12.4 GYr; this gives a per year drift constraint on the order of 10^{-16}. Molecules considered for use in current and future probes include NH_3, methanol, and others [7].

This is not an exhaustive list of the techniques used to measure possible drifts and oscillations in constants. Refs. [7] and [11] detail several other types of experiments and observations used to place bounds.

1.5 Atomic and Molecular Clocks

One tool used to measure several fundamental constants is an atomic or molecular clock. Atomic clocks function by using atomic transitions (changes in electron energy levels) to measure the exact frequency of electromagnetic radiation (often in the form of a laser). By converting this frequency into a time, we can create an extremely consistent clock. Cesium is the most common atom used in atomic clocks for a variety of reasons, including because of its angular momentum properties, which make analysis of the system easier because small magnetic fields do not cause shifts in the transition. Therefore, it is used to define the second. In particular, the transition in cesium that is used to form the atomic clock is called

a hyperfine transition. The hyperfine structure occurs because of interactions between the magnetic dipole moments in the electron and the nucleus. This interaction causes a split in the energy levels, which forms the basis for the transition used to define the clock (since cesium has a single valence electron, this is the electron in use). This hyperfine interaction is desirable because the frequency range is accessible by exciting the cesium with high-precision microwaves.

Generally, we want to measure frequency ratios with comparisons of atomic clocks, since the ratios are pure numbers that do not depend on units. Therefore, these types of measurements also use a reference clock, allowing for a comparison of two frequencies [13]. Cesium is often a useful clock because of its high precision, and because it allows for the classification of the data in terms of Hertz, based on the definition of a second. Additionally, this use of frequency ratios allows for the cancellation of some systematic effects, since some forces will affect the atoms in both clocks in the same way.

Each type of transition in an atomic clock can be characterized as a microwave or optical transition, depending on the photon wavelength associated with the transition. Microwave and optical transitions display slightly different behavior in terms of the fundamental constants on which the clock frequencies depend. Microwave transitions are typically hyperfine, so they are related to electron and nuclear spin, and therefore dependent on α. Optical transitions are usually purely electronic, so they are dependent on α and μ. The g factors of the nuclei are also important for microwave transitions, but can be reduced into more fundamental constants and calculated using theoretical models. Based on these relations, the ratio of two microwave clock frequencies depends on α as well as the g factors of the two nuclei; the ratio of two atomic optical clocks depends only on α; and the ratio of a microwave clock to an optical clock depends on the g factor of the nucleus of the microwave clock as well as α and μ [7]. We can then use these relations to make a measurement depending on the types of transitions being compared. Since some ratios of microwave atomic clocks depend on both α and μ, many experiments find a constraint on one of the constants in terms of

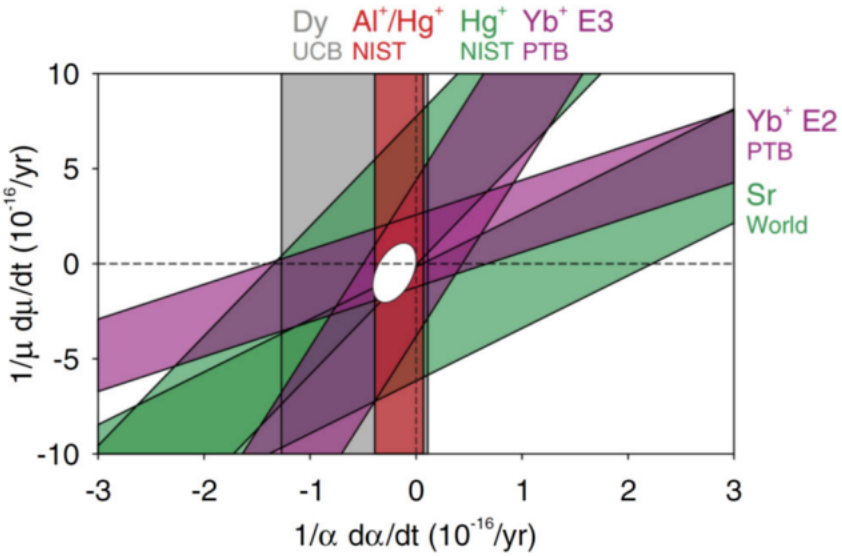

Figure 1.1: Bounds on variation in α and μ from various experiments using optical and microwave atomic clocks. Each stripe represents the results of an experiment (to one standard deviation of uncertainty), and the white region in the center is obtained by combining the bounds from each of the experiments. Note that current bounds on μ are consistent with a lack of variation. Figure reproduced with permission from [7].

the other one; that is, the experiments draw a line with some error on a plot comparing the changes in α vs μ. By combining several of these lines, we can get a very good bound on both μ and α, as seen in fig. 1.1.

Microwave clocks such as cesium have produced a variety of very accurate measurements in drifts in fundamental constants. But over time, progress on improving these atomic clocks has slowed, since technological and systematic limits have prevented significant improvements on their accuracy. Therefore, many current experiments seek to compare optical transitions, including in molecules [7]. These clocks function similarly to atomic clocks, but use vibrational (or rotational) states instead of electronic states as in atoms. In general terms, these transitions are sensitive to μ because they depend both on the mass of the nuclei in the atoms, and on the chemical bond between the atoms, which is dependent on the mass of the electron. In fact, molecular optical transitions are dependent only on μ. These transitions are sensitive to several possible new physics effects because the electron

and proton get their masses from different mechanisms. Protons are made up of quarks, so they get their mass primarily from the strong force and thus have masses proportional to the quantum chromodynamics energy scale, while electrons are fundamental particles, so they get their mass from the Higgs field [13].

In general, there are several different choices for molecules for current or future use in testing for oscillations or drifts in fundamental constants. One early measurement resulting in a good constraint on μ using molecules was found using a rovibrational (combined rotational and vibrational) transition in SF_6, with a cesium microwave clock as a reference [16]. KRb was used to measure the current best limit on drift by using a near-degeneracy between the ground state and a second state; hydrogen molecules, namely H_2^+ and HD^+, are also used because these molecules are extremely well-understood, making the theory and systematic effects well-quantifiable [13]. In our experiment, we plan to use O_2^+ molecules, which have multiple transitions that can be used to investigate μ variation. Oxygen's most common isotope has no nuclear spin, so the molecule has no hyperfine structure, making it easier to drive the correct transitions. Additionally, oxygen's transitions have narrow-lines, which means that the transitions have intrinsically small frequency uncertainty, allowing for long coherent probe times [13]. This means that there is more data produced during a given measurement of a molecule, and therefore the statistics are averaged more quickly, allowing for better uncertainties. Several other types of molecules are being used in other experiments, and it remains to be seen which will be the most useful for setting bounds [7].

1.6 Sensitivity

In order to find time variation of μ, we need two quantum states that have different dependencies on μ. This is possible due to the fact that different atomic transitions energies may

have different relations with fundamental constants. In atoms, this looks like

$$E \propto (\frac{\sigma_A}{\sigma_B})^{\zeta}(\alpha)^{\gamma+2}. \tag{1.5}$$

In this equation, σ_A and σ_B are the magnetic moment of the nucleus and the Bohr magneton, respectively; this ratio is directly proportional to μ. ζ is a fixed exponent, 1 for hyperfine transitions and 0 for optical transitions, and γ is a fixed exponent for each transition that depends on the transition's spin-orbit couplings and electron cloud effects [14]. By considering the coupling between a theorized scalar field and μ and α for two different atomic transitions, we can get a frequency ratio between two transitions with sensitivity to beyond-the-Standard-Model physics [14].

In molecules, which we use, this looks slightly different. When considering the energy levels of atoms, we only need to consider the electronic transitions; molecules also have vibrational and rotational degrees of freedom that contribute additional energy levels. In general, the energy of a diatomic molecule, written in the units of wavenumbers, can be written as

$$E/(hc) = T_e + \omega_e(v + \frac{1}{2}) - \omega_e x_e(v + \frac{1}{2})^2 + B_e J(J+1). \tag{1.6}$$

T_e, the electronic energy, does not depend on μ, so we do not consider it. The vibrational energy is represented by both ω_e, the vibrational coefficient, and $\omega_e x_e$, the anharmonicity coefficient, which scale as $1/\sqrt{\mu}$ and $1/\mu$, respectively [13]. The rotational energy is represented by B_e, which also scales as $1/\mu$. To see the dependence in the harmonic term, we can start a derivation, reproduced from ref. [17], with the energy eigenstates associated with the vibrational quantum number, v:

$$E_v = \omega_e(v + \frac{1}{2})hc. \tag{1.7}$$

We also have

$$\omega_e = \frac{1}{2\pi c}\sqrt{\frac{k}{m}}, \tag{1.8}$$

where k is the spring constant of the chemical bond, and m is the reduced mass of the two nuclei. We can think of this reduced mass as being proportional to the mass of the proton since the proton and neutron get their over 99% of their masses from the strong interaction, so any changes in the proton mass due to a scalar field will apply similarly to neutrons and therefore to m [13]. The spring constant is approximately proportional to the mass of the electron: conceptually, this comes from the chemical bond, which is made up of the electrons, behaving as a spring. Therefore, we get the proportionality relations

$$E_v \propto \omega_e \propto \sqrt{\frac{m_e}{m_p}} = \frac{1}{\sqrt{\mu}}, \tag{1.9}$$

giving the harmonic relation stated above. Similar derivations yield the dependence on μ for both $\omega_e x_e$ and B_e. The vibrational transitions have the best sensitivity in general, since the vibrational coefficient is much larger than the rotational coefficient, and it is difficult to drive transitions with large changes in J because of the conservation of angular momentum.

We can write an absolute sensitivity factor, q_μ, based on a small change in the energy of the state, ΔE, using the fractional change in μ, $\frac{\Delta\mu}{\mu}$, as

$$\Delta E = q_\mu \frac{\Delta\mu}{\mu}. \tag{1.10}$$

This relation allows us to convert from an experimentally measured shift into a variation in μ [13]. To maximize our ability to measure a fractional change in μ, we want as large a value for q_μ as possible.

Generally, for small v (vibrational quantum number), the q_μ scales linearly with v. At higher values of v, the anharmonicity term begins to dominate, taking the sensitivity back towards 0. This can be though of conceptually in that we get better sensitivity from larger

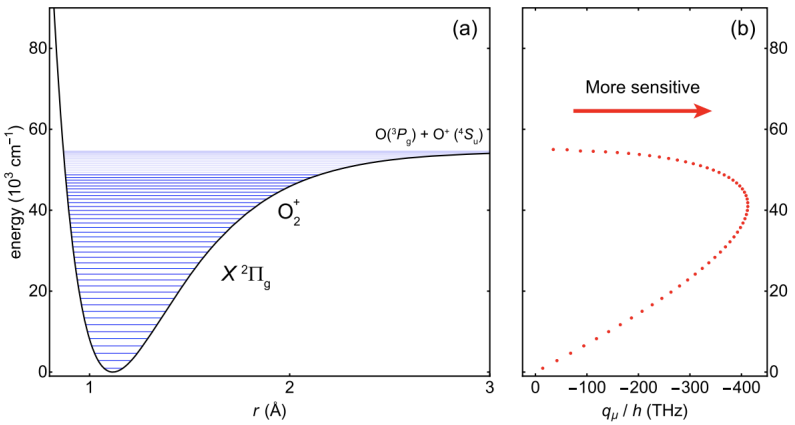

Figure 1.2: Vibrational energy levels and absolute sensitivity for the O_2^+ ground state (labelled $^2\Pi_g$). (a) Potential curve, with energy as a function of bond length. Energy levels are indicated by horizontal lines. (b) Absolute sensitivity for each vibrational state. Note the largest magnitude at about 3/4 of the dissocation energy. Figure from [13], used under the Creative Commons license.

vibrational transitions, but as the transitions get too large, the molecule starts to behave as two individual atoms, meaning that the sensitivity that depends on the molecular structure disappears. The maximum sensitivity for molecules generally appears at about 3/4 of the dissociation energy, which can be seen in fig. 1.2 [13].

1.7 Our Setup

To make a measurement of energy levels in oxygen ions, we need a system that can hold ions in place for long enough that we can drive the desired optical transitions, which allows us to make the measurement of a fractional change in μ. To be able to do this, we employ a radiofrequency ion trap inside of a vacuum chamber. With laser cooling techniques, the ion trap keeps the ions confined in a tight "crystal," an example of which can be seen in fig. 1.4. The scales of these crystals are on the order of tens to hundreds of microns. When they are not interacting with outside particles, the ions can stay trapped for extended periods of time.[4] Indeed, this is why we need a vacuum chamber in our setup. Having a vacuum

[4]We found that ions could stay in the trap in excess of 90 minutes (this limit was determined by the amount of time we looked, not because ions left the trap after 90 minutes), which is more than sufficient for

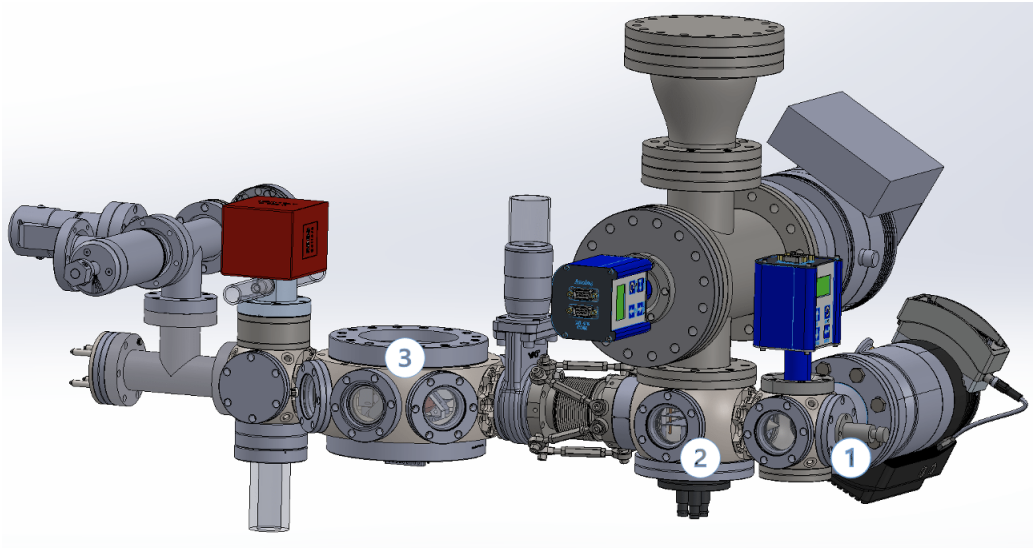

Figure 1.3: SolidWorks model of the full vacuum setup used for the experiment. (1) is the molecular beam source chamber, (2) is the differential pumping chamber, and (3) is the ion trap chamber. Other shown components include various pumps, valves, and gauges, described in detail in chapter 2. The length of the setup is about 80 cm.

ensures that the ions will very rarely interact with any unwanted particles, and therefore will be uninterrupted in the ion trap. If there are additional gas particles near the trap, then the particles may strike the ions and knock them out of the trap's potential well. A detailed discussion of the design and construction of the vacuum chamber can be found in chapter 2.

The second major portion of the experimental setup is the molecular beam apparatus. This is discussed in detail in refs. [17] and [18]. The molecular beam source sends a beam of oxygen through a set of skimmers and into the trap chamber, where it reaches the center of the ion trap. These skimmers, which are 1 mm and 0.5 mm in diameter, select the coldest molecules in the beam, ensure the beam is collimated, and reduce the amount of particles that reach the trap chamber, since excess molecules would worsen the vacuum. In between the two skimmers is a differential pumping chamber, which protects the vacuum conditions in the trap chamber from the relatively high pressure in the molecular beam source chamber. Once the oxygen reaches the center of the ion trap, lasers photoionize it and drive the vibrational transition, which allows us to make our measurement of the molecular energy

our purposes.

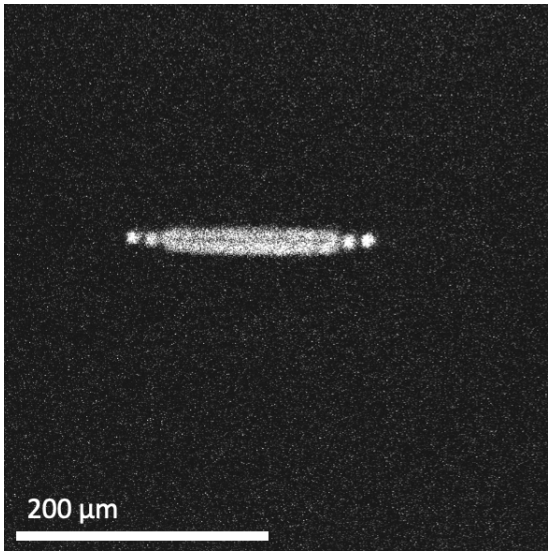

Figure 1.4: Photo of crystallized Be^+ ions. The field of view displayed is 436 microns x 436 microns.

levels. Then, a final laser dissociates the oxygen molecule into a neutral atom and an atomic ion, which we use to confirm that the correct transition has occurred. These lasers enter the trap chamber through the windows on the outside of the chamber. The spectroscopy laser must be very high intensity, so it needs a Fabry Perot cavity inside the vacuum chamber to build up power.

As stated above, one of the key elements of the trap is its ability to confine ions in crystals when used alongside laser cooling. This process requires the introduction of Be^+ ions, which are cotrapped alongside the O_2^+ ions and can be cooled using a 313 nm ultraviolet laser; this causes the ions to appear individually, such as in fig. 1.4, instead of in a cloud. Oxygen ions cannot be laser cooled on their own because of properties of the molecule's electronic and vibrational transitions [19]. Since this cooling reduces the energy of the cotrapped ions, it means that they are less likely to escape the trapping potential and therefore stay trapped for longer. The beryllium ions are also important because they allow us to image the system. The photons from the UV laser induce a cycling transition in the beryllium ions, which causes them to release photons, which can then be detected by an imager. This process is discussed further in section 4.2.1. In order to add the beryllium ions to the system, the trap

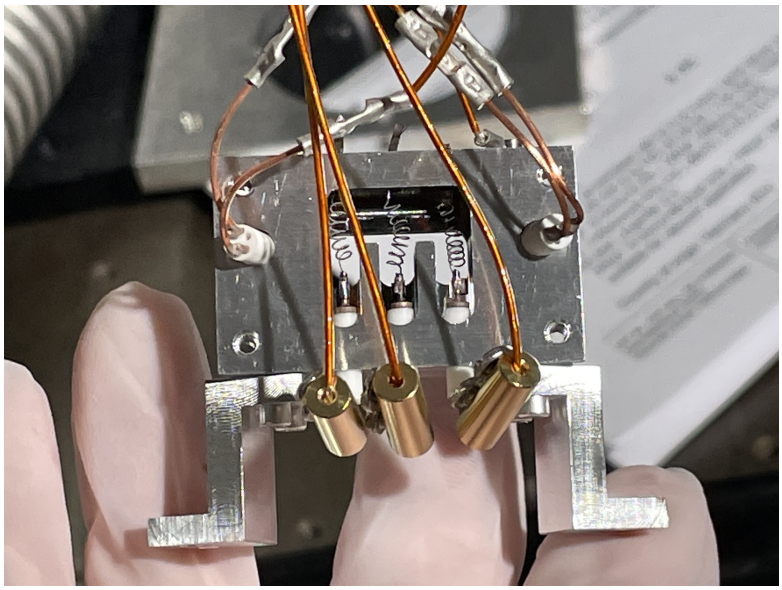

Figure 1.5: Photo of the beryllium oven used in the trap. The coils in the center are coated in beryllium, which is released when heated. When activated, the coils glow orange. The center coil was removed, since it would have blocked the O_2^+ molecular beam. Note the hand for scale.

chamber contains a beryllium oven (see fig. 1.5), which releases beryllium atoms into the vacuum chamber, and an electron gun, which emits electrons that ionize them.

The key upgrade we made to the apparatus this year was linking together the ion trap chamber and the molecular beam apparatus, which were previously fully separate components. The construction of this system was very challenging; significant care was taken to ensure that the beam line aligned with the trap and the vacuum side of the experiment. This was achieved by using springs to suspend the beam apparatus slightly above the optical table containing the other equipment. This new setup introduces additional difficulties in calibrating all of the experimental components. The molecular beam, the lasers for the photoionization and spectroscopy of the oxygen, and the UV laser for the beryllium cooling must all converge at the center of the trap. Some of these beams are easy to move: the cooling laser and the photoionization laser, for example, are focused through lenses on translational stages outside of the trap, so their positions can be easily tuned, and the molecular beam can be tuned with a hexapod bellows, which, while more difficult to adjust than a translational

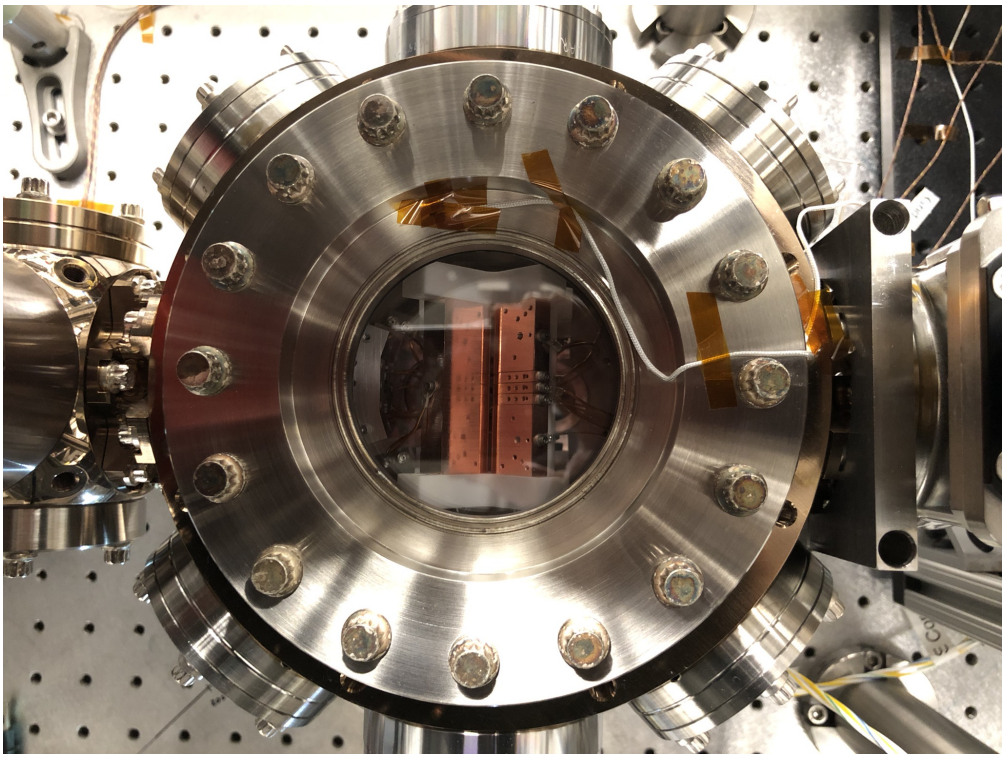

Figure 1.6: Photo of the ion trap inside the vacuum chamber. The window has a diameter of 7 cm.

stage, is still usable. The spectroscopy laser, however, is locked into place by the location of a Fabry Perot cavity inside of the vacuum, so once we seal the vacuum chamber, it cannot be moved. The laser mode is on the order of tens of microns wide, so since it is nearly impossible to tune it to the exact center of the trap (because we can't check our work with the ions when the vacuum is not yet sealed), we require a mechanism that can change the trapping location from the geometric center of the trap to a different point, in line with the beams. To do this, I designed an electronics system that allows for versatile control of both DC and radiofrequency potentials on the four trap electrodes. With these controls, I can drive pairs of in-phase RF signals on trap electrodes that have differences of up to a factor of five on their relative amplitudes, put an arbitrary DC voltage on any of the electrodes, and use AC modulation to test properties of the ions and trap geometry. The ability to vary the RF potential allows for precise tuning of the ion positions in the trap, which allows us to move the ions to the necessary position. The principles and design of these electronics

are discussed in chapter 3, and detailed parameters of the circuit, including trapping results, can be found in chapter 4. Finally, I wrap up with some closing thoughts, including possible improvements and the future of the experiment, in chapter 5.

Chapter 2

The Vacuum Chamber

While the ion trap is the main feature of the experimental setup that allows us control over the ions, it needs certain conditions to function. Namely, in a system at atmospheric pressure, collisions with background gas molecules may either transfer enough energy to the trapped ions that they leave the trap's potential well, or knock them to a physical location outside of the trap center; if ions are not in the trap, we cannot measure their behavior. Therefore, in order to create a system that will maintain the trap for as long as possible and allow us to use our series of laser pulses to make the required measurements, we create an ultrahigh vacuum (UHV) in the trap chamber. In this chapter, I discuss our vacuum setup and the theory on which we base our design choices. Then, I discuss the materials, pumps, and gauges present in our vacuum system before concluding with a summary of our pumpdown.

2.1 Our Vacuum Setup

Our vacuum setup consists of three chambers, each of which is pumped separately. The first is the molecular beam source chamber, and is separated from the second by a 1 mm diameter skimmer. The second chamber is a differential pumping chamber, and is connected to the third, the ion-trap chamber, by a second skimmer, which has a 0.5 mm diameter. The

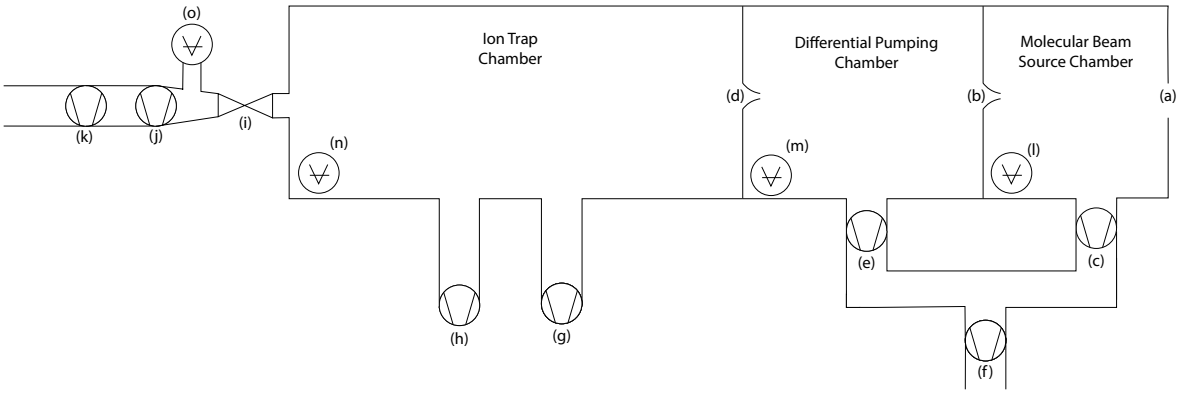

Figure 2.1: Schematic of the vacuum chamber setup. a is the molecular beam source, and b and d are 1.0 mm and 0.5 mm diameter skimmers, respectively. c, e, and j are turbo pumps for each of the three chambers, and f and k are both rough backing pumps with outputs at atmospheric pressure. g is a getter pump and h is an ion pump. i is a valve to allow for removal of pump j. l, m, and n are ion gauges, and o is a residual gas analyzer.

pumping scheme is slightly different in each of these segments, and is described in section 2.4.

Fig. 2.1 shows a schematic diagram of the vacuum chamber setup. The technologies used in each of the pumps and gauges are described below. When the experiment is running, we want the ion trap chamber to be at a pressure around $1\text{x}10^{-10}$ torr, the molecular beam chamber at about $1\text{x}10^{-6}$ torr (shortly after the beam is fired), and the differential pumping chamber in the middle, roughly at $1\text{x}10^{-8}$ torr, protecting the trap chamber from exposure to the higher-pressure molecular beam chamber. Since an ultrahigh vacuum is defined as roughly 10^{-7} torr, we want the ion-trap chamber to be firmly in the UHV regime [20, preface].

2.2 Vacuum Chamber Theory

Before describing each of the specific components and the processes needed to create this vacuum chamber, it is first important to understand the underlying principles of vacuum-chamber design. Generally, gas is modelled using the kinetic theory of gases. This makes several assumptions about the behavior of the gas: the molecules have negligible volume, have negligible intermolecular forces, they move continuously and randomly, and all of the

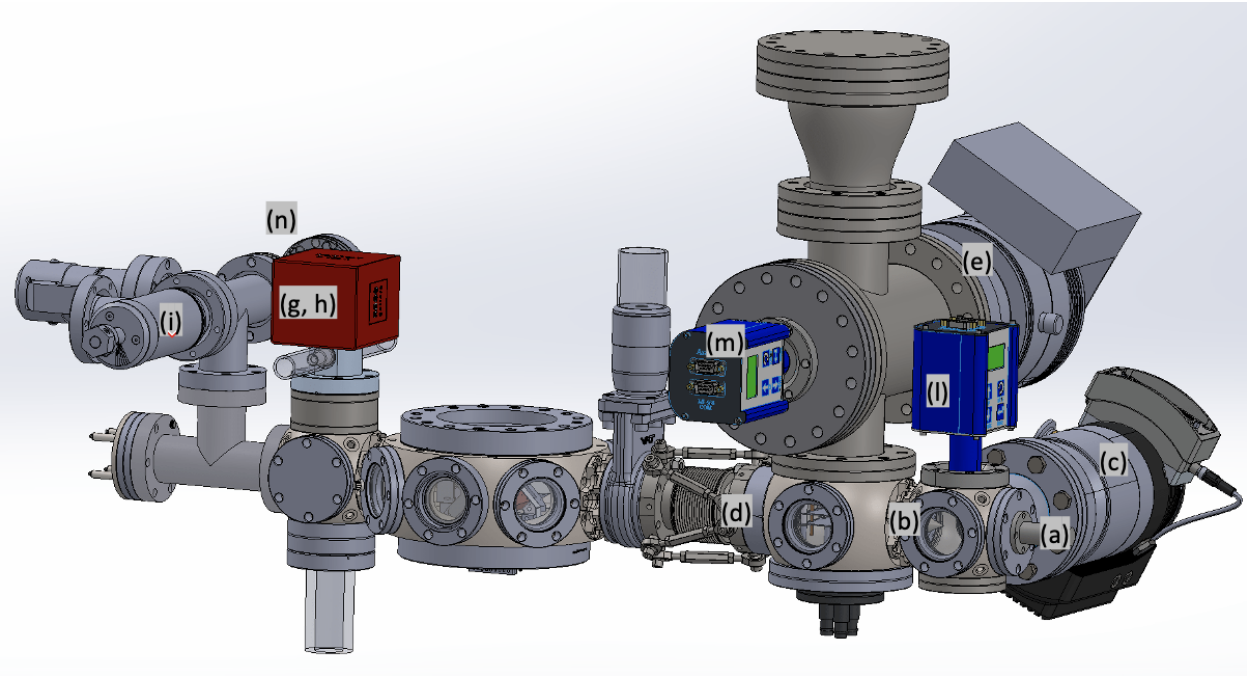

Figure 2.2: SolidWorks model of the full vacuum setup. Components are indexed as in fig. 2.1: a is the molecular beam source, b and d are 1.0 mm and 0.5 mm diameter skimmers. c and e are turbo pumps; the third turbo pump, and the backing pumps for these, are not included in the model. g and h are the getter and ion pumps, which are both in the same part in the setup. i is a valve, and l, m, and n are ion gauges. The residual gas analyzer is also not included. The full model is roughly 80 cm in length.

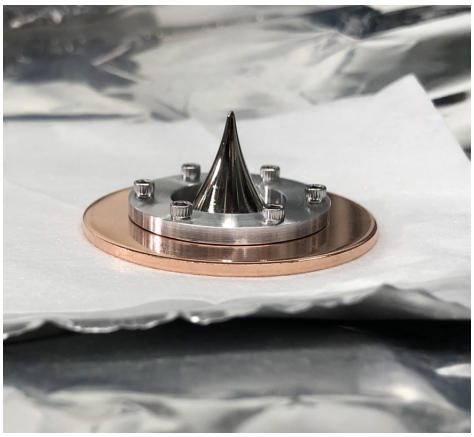

Figure 2.3: Photograph of a skimmer used to limit the molecular beam flow into the trap chamber mounted onto a flange. The copper gasket surrounding the skimmer has a diameter of 2.75 inches. Since the opening has very low conductance, gas must be pumped from the chamber on both sides of the skimmer to maintain a sufficient vacuum.

collisions are perfectly elastic. This can be applied to the case of a vacuum chamber; the main point to consider is that as we approach the limit of a perfect vacuum, there is very little interaction between gas molecules. We can evaluate this property numerically using the mean free path; this is the average distance gas particles travel between interactions. This can be estimated using

$$\lambda = \frac{1}{n\pi d^2\sqrt{2}}, \tag{2.1}$$

where d is the diameter of the molecule (represented as a sphere) and n is the number of molecules per unit volume [20, eq. 1.10]. For atmospheric pressure and room temperature (for nitrogen molecules), λ is about 6x10^{-6} m, while at our intended conditions, about room temperature and 1x10^{-10} torr, λ is about 6x10^{7} m [20, table 1.1]. Since this is much larger than the size of the chamber, an overwhelming majority of the interactions are with the chamber walls. This means that gas particles in the trap chamber are very unlikely to interact with ions in the trap.

In order to create a system with these types of conditions, gas must be removed from the vacuum chamber. As stated above, this is achieved by using various types of pumps, which reduce the pressure in the vacuum chamber in various ways. The types of pumps used in our vacuum chamber are detailed in section 2.4.

We describe the evacuation of the gas from the vacuum using mass flow, which is simply the rate of removal of the gas from the system, measured in terms of a mass per time. The gas flow rate Q, called the gas throughput, can be written as

$$Q = C(P_1 - P_2), \tag{2.2}$$

where P_1 and P_2 are the pressures of the two chambers connected by a gap, tube or other structure in the vacuum chamber, and C is the conductance of the structure, which refers to the flow of gas through the structure in units of volume per second [20, eq. 1.17].

From relatively simple derivations in ref. [20, ch. 1], we find that the conductance of an

orifice can be expressed as

$$C_0 = A\left(\frac{k}{2\pi}\right)^{1/2}\left(\frac{T}{m}\right)^{1/2}, \tag{2.3}$$

where A is the area of the orifice, k is the Boltzmann constant, T is the temperature of the system, and m is the mass of a molecule. Thus, the conductance depends proportionally on the area of the orifice and on the square root of the temperature, and inversely on the square root of the mass of the gas. So, since $v_{rms} \propto (T/m)^{1/2}$, the conductance is proportional to the root-mean-square velocity of the molecules. Tubes are slightly more complex; the conductance of a long tube, such that the effect of the ends is negligible, is

$$C_{LT} = \frac{4}{3}\frac{r^3}{L}(2\pi k)^{1/2}\left(\frac{T}{m}\right)^{1/2}, \tag{2.4}$$

where r is the radius of the tube and L is the length of the tube, and for a short tube, we get

$$C_{ST} = \frac{C_0}{1+\frac{3}{8}\frac{L}{r}}. \tag{2.5}$$

Most of the components on our apparatus can be approximated as short tubes, since their diameters are significant in comparison to their length, although the tube responsible for pumping the gas out of system can be approximated as a long tube.

Conductance and throughput are both useful quantities to track because they determine the rate at which we can reach the ultrahigh vacuum regime. For most of the components of the vacuum setup, the conductance is relatively high: there is a significant amount of room for gas particles to exit the chamber. One major exception to this is the skimmers (b and d in figs. 2.1 and fig. 2.2), discussed in section 1.7, which are 1 mm and 0.5 mm across, respectively, so that they only allow a portion of the molecular beam through. For this reason, instead of relying on flow through the skimmers as we pump, we pump independently from each of the sub-chambers that are separated by the skimmers.

Still, even though we pump separately from each of the three chambers in the setup, it is important to consider gas flow through them as a possible limiting factor on the minimum

pressure in the trap chamber, since we want the trap chamber to have lower pressure than the other two chambers. Therefore, I reproduce the calculations completed by Professor Hanneke confirming that this does not limit our vacuum. In general, from eq. 2.3, we see that the conductance of a circular hole for molecules is

$$C = (0.111L * s^{-1} * mm^{-2})A, \tag{2.6}$$

where A is the area of the hole in square millimeters. Using equation 2.2, we let P_1 be the pressure of the differential pressure chamber and P_2 be the pressure in the trap chamber. In this case, the pressure in the trap chamber is much lower than in the differential pressure chamber, so we approximate $P_1 - P_2 \approx P_1$. Since the skimmer has a 0.5 mm diameter, it has conductance of 0.022 L/s, and the differential pressure chamber has a pressure of about 3x10^{-8} torr, since it is being pumped with a turbo pump (see section 2.4.1). Then the gas throughput, Q, into the trap chamber, is about 2.6x10^{-9} torr*L/s. This gas is then pumped out by the getter pump in the trap chamber (see section 2.4.2). Since the overwhelming majority of this gas is oxygen from the portion of the beam that was rejected by the skimmer, this is the relevant gas for pumping; the NexTorr pump has a pumping speed for oxygen of 65 L/s when accounting for conductance through the trap chamber, which allows for a final pressure of about 1x10^{-11} torr. This value is well within our acceptable bounds for the trap to work as expected. The full system is slightly more complicated than this, since we have an additional skimmer between the molecular beam and the differential pumping chambers (the molecular beam chamber will be at a higher pressure than the differential pumping chamber), but the result does not change significantly.

Since any exposures to the atmosphere of the room allow gas into the vacuum chamber, it is necessary to avoid any kinds of leaks in the chamber. In general, the chamber is made up of a series of components, which have stainless steel bodies, with many also containing glass viewports to allow lasers to pass into the trap to strike the ions, and to allow for

calibration of the lasers and molecular beam. Most of these parts are manufactured to maintain ultrahigh vacuum conditions, so they contain no significant leaks (types of minor leaks are discussed below). In order to connect these parts, a copper seal is used. Each of the connectors between the parts contains a groove in which a copper gasket can be placed; when two of the parts are bolted together, knife edges on each of them partially cut into the copper gasket, creating a seal that allows no air inside of the knife edge. This method means that loose bolts may cause leaks in the vacuum, so all of the connections were tested for leaks at several points throughout the pumping process. To do this, we sprayed liquid methanol onto the seals. Methanol is easily identified by the residual gas analyzer (discussed in section 2.5.2), so when any methanol made it through the seals, it could be determined exactly where it entered the system, and the relevant bolts could then be tightened.

Virtual leaks are another possible problem to achieving ultrahigh vacuum. Virtual leaks occur when small openings, such as screw threads (see fig. 2.4), trap gas inside of the vacuum chamber. Although this does not cause the final pressure of the system to be reduced (since no outside particles are introduced), the time to reach the desired ultrahigh vacuum is increased drastically, since the conductance through such a small opening is tiny. To combat this, relief holes are drilled to all possible pockets in the vacuum chamber, often down the center of the screw, or directly through the material to the pocket.

Vaporization may also cause a higher pressure inside the vacuum chamber. In the vacuum chamber, many solid and liquid molecules are bound to the sides of the chamber through intermolecular forces. In particular, water vapor is commonly found here, because of its strong intermolecular forces due to the polarity of its molecules. Since particles exist at a distribution of kinetic energies, occasionally one of these solid or liquid particles vaporizes and introduces a gas particle to the vacuum chamber, which increases the pressure of the chamber. This process does not depend on the pressure of the vacuum chamber, but on the temperature of the chamber and on the type of particle in question and its intermolecular force strength with the walls of the chamber [20, section 1.5.3]. To combat this, we perform

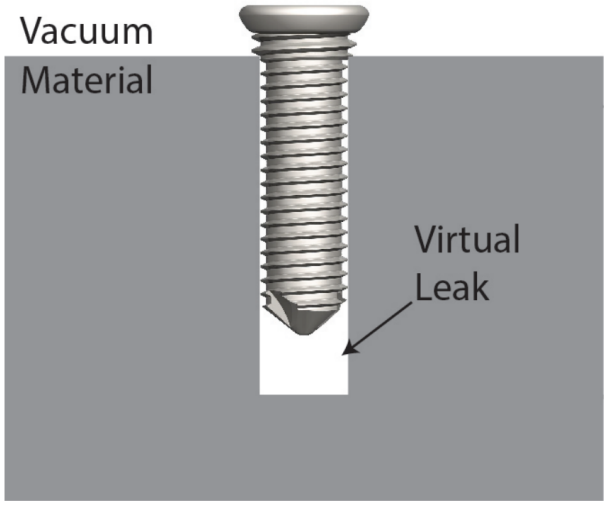

Figure 2.4: Example of a virtual leak in the vacuum chamber. Image from [21], used under the Creative Commons license.

a "bakeout" of the vacuum chamber following its construction. This introduces enough kinetic energy that nearly all of the relevant solid and liquid molecules on the walls of the chamber (again, mainly water molecules) become gas molecules, which can be easily pumped out of the chamber. This bake takes about two weeks; since various parts of the experimental apparatus are sensitive to different levels of heat, the parts are held at different temperatures throughout the process, ranging from roughly 80°C up to about 160°C. Our bakeout is discussed below in section 2.6.

Surface and volume outgassing are two final problems to consider. These occur when gas particles are absorbed or adsorbed by the metal volumes in the vacuum chamber, and leak out over time into the vacuum chamber. Combating these is similar to vaporization; the metal components are baked ahead of construction, at temperatures of up to 430°C, which causes a lot of the gas to leave their volumes, securing a better vacuum. Additionally, the bakeout following the construction of the system allows for additional hydrogen outgassing from the metal components. These cleaning procedures are considered standard, and therefore allow us to use tabulated values for outgassing; assuming that the stainless steel (which makes up almost all of the vacuum chamber) provides the vast majority of the outgassing, we get a

gas from outgassing of 2x10^{-11} W [21]. When accounting for pumping speed, this gives us a limit on the lowest possible vacuum we can reach, estimated for our setup to be on the order of 1x10^{-11} torr by ref. [21]. This value is sufficiently low that surface and volume outgassing are not an issue for the level of UHV we hope to reach.

2.3 Apparatus Materials

In order to create a system that can reach the ultrahigh vacuum regime, we need to choose materials that can meet the necessary conditions. Indeed, the set of constraints on these materials is demanding. Vacuum components must have a relatively small amount of outgassing, and must be able to withstand the bakeout temperatures, which in many cases are over 100°C. Additionally, any materials forming the walls of the vacuum system must be able to withstand the pressure difference between the vacuum (less than 10^{-3} torr, although for mechanical purposes, the vacuum is essentially perfect), and atmospheric pressure. For this reason, we are very deliberate about the makeup of the components used inside the vacuum chamber.

2.3.1 Metals

There are several types of metal in use in the vacuum chamber. Ref. [20, section 2.3] argues that there are two main constraints on the metal in an ultrahigh vacuum environment: the vapor pressure at the desired temperature must be below the intended pressure of the system, and gas from the atmosphere must not be able to permeate the metal. This second point is only relevant for the stainless steel components and copper gaskets in our system since these are the only metals forming a boundary between the vacuum and the atmosphere. Gas permeation through these metals is unavoidable completely, but under the conditions in our assembly, this is not an important factor in determining the eventual pressure of the vacuum.

To address vapor pressure, we can look at each of the types of metal in the vacuum chamber. The main metal in our system is stainless steel, and in addition, there are a variety of other metals in smaller quantities. In addition to the gaskets, wires and the ion trap itself are made of copper. The various screws in the vacuum chamber are made up of titanium or silver-plated stainless steel, and the baseplate holding the trap is made of aluminum. There are also other metals in use, such as in the eyelets attached to the trap electrodes. The vapor pressures of most of these types of metal are tabulated and easy to refer to, such as in ref. [20, table 2.4], and when doing so, we see that none of these metals have a vapor pressure above $1\text{x}10^{-11}$ torr at room temperature, so this is not an issue in our system.[1]

Outgassing is also a significant concern when using metal components in a vacuum chamber. There is detailed information on the solubility of various gases inside metal, but in our case, we primarily consider hydrogen. As stated above, we use a two-step baking process for most of the metal components (described generally in section 2.2 and for our system in section 2.6), but hydrogen outgassing may cause the ultimate vacuum limit in the system. Still, as argued in section 2.2, this is unlikely to reach a point where the vacuum is not suitable for ion trapping.

2.3.2 Glass

The glass in use in the vacuum chamber is fused silica, which is pure SiO_2. Chosen for its optical properties (it allows UV light through, as opposed to standard glass, which does not), this glass makes up the viewports throughout the vacuum chamber. These include the large one that allows for the imaging of ions in the trap, and the smaller ones that allow laser beams through to perform the experiment. For the purposes of the experiment, it is

[1]The exception here is the beryllium oven, which heats beryllium coils in order to increase the vapor pressure and send the beryllium into the vacuum system. This has the side effect of outgassing hydrogen into the vacuum chamber, but this is pumped out by the getter pump once the heating is complete. Although the pressure in the trap chamber increases when the beryllium oven is activated for 1 minute to about $4.67\text{x}10^{-9}$ torr, after 5 minutes of pumping, the chamber returns to $3.4\text{x}10^{-10}$ torr, only a factor of 1.25 above the initial pressure of $2.7\text{x}10^{-10}$ torr.

important that the glass does not become warped or deformed during any bakeouts, since the laser light must enter the chamber unimpinged; fused silica is ideal for this, since it has a very high softening point (over 1500°C) compared to other types of glass, and has a very small expansion coefficient [20, table 2.1].[2] The small expansion coefficient means that as the glass expands with temperature, there is not significant stress from the glass pressing against the metal seals. This is important since any stress can change the laser behavior through the glass by causing birefringence, so a small expansion coefficient ensures that this is not a problem during the bakeout. Glass is not a perfect vacuum material in the sense that it does allow the permeation of some outside gases. However, for our purposes, these effects are negligible. There is also an effect on pressure due to outgassing in the glass. Just as with the metal components, the bakeout speeds up this process significantly, and allow us to reach a stable pressure within a reasonable timeframe without having to account for the continued outgassing from the glass.

2.3.3 Ceramic

Ceramics are another suitable material for use in vacuum systems, with the restriction that only non-porous ones are effective. For this reason, we use MACOR ceramic for the trap supports in our vacuum chamber. MACOR is made up of about 55% fluorophlogopite mica and 45% borosilicate gas; it maintains its structural integrity at over 1000°C, and can be machined up to the tolerances needed to serve as supports for the trap electrodes (other ceramics can be too brittle and crack when machined) [22]. Generally, ref. [20, section 2.4] advises that ceramics should not be heated to within 400°C of their softening temperature when the tolerances are significant; this is then not an issue for our bakeout. Although permeation through ceramics occurs similarly to glass, this is not significant in our case, since there are almost no ceramic surfaces forming a border between the vacuum and the atmosphere (the only possibility is that some of the electrical feedthroughs have ceramic

[2]The softening point refers to the point at which the glass reaches a certain degree of viscosity. In our bakeout, we want to avoid approaching this value.

insulators). In general, ceramics outgas in a similar manner to glass. Bakeout temperatures higher than glass are in general required, but because the MACOR material has almost no porosity, the bakeout of the rest of the components is sufficient.

2.3.4 Plastics

Finally, there are several different plastics in use in the vacuum system. Many plastics are extremely porous and can outgas almost indefinitely, and bakeouts cannot be used to speed up the outgassing since the plastic will melt [20, section 2.5]. However, there are a few kinds of plastic that can be used in ultrahigh vacuum environments. PEEK plastic is used as part of the electrical connector to the oven in order to hold the wires in place. The wires inside the vacuum are all coated in Kapton insulation, and the gate valve, which provides a border between the molecular beam and the trap chamber, has a seal made of Viton rubber. All of these plastics share the property that they do not have significant issues with outgassing, and that they maintain their structural integrity at temperatures up to and beyond 150° C, which is necessary for the bakeout. The one possible exception here is Viton: it has an issue with outgassing unless it is vacuum baked (ours is), but even so, it still could be one of the largest gas loads in the system.

2.4 Types of Pumps

Once we have a system that has materials and conditions suitable to create an ultrahigh vacuum environment, we need to pump the gas out of the system to reach the desired pressure. Many types of pumps exist for this purpose, and we include several in our setup. As stated above, our system consists of three chambers, the molecular beam, the differential pumping chamber, and the trap chamber. For the first two chambers, we use turbo pumps, which are kept in operation throughout the entire experimental process (c and e in figs. 2.1 and 2.2), with a scroll backing pump (f). The third chamber, the trap chamber, is first

pumped down very roughly with a diaphragm backing pump (k in fig. 2.1), then a turbo pump (j) is used to bring the chamber into the UHV regime; finally, the turbo pump is removed, and getter (g) and ion (h) pumps are used to maintain the vacuum. These types of pumps are described below.

2.4.1 Turbo Pump

The first type of pump in the vacuum system is the turbo pump. Inside the pump, a cylinder with blades attached, similar to a propeller, spins extremely quickly. The blades strike gas molecules and direct them out of the vacuum chamber. The process assumes that the momentum transferred to the molecules by the blade is much greater than the starting momentum of the molecule, and that the mean free path length of the molecules is long enough that the interactions between the molecules in the pump is insignificant; both of these are true in a system like ours. The ratio of the pressure on each side of the pump can be written as

$$\frac{P_{out}}{P_{in}} = e^{Av}, \tag{2.7}$$

where v is the angular velocity of the cylinder, and A is a constant based on the geometry of the pump and the properties of the gas [20, section 3.3]. To maximize A, the rotating surface should have a high surface area, and should have very small gaps between it and the surrounding wall. The turbo pump in our trap chamber, the Leybold TurboVac 90 i, reaches up to 72,000 revolutions per minute [23]. Since the turbo pumps rely on low pressure to work, and small particles may damage them, the turbo pump is activated after the backing pump in the trap chamber. The molecular beam source chamber also contains a Leybold TurboVac 90 i, while the differential pumping chamber contains an Agilent 304 FS turbo pump, which reaches 60,000 revolutions per minute [24].

It is worth noting that based on equation 2.7, turbo pumps create a ratio of pressures between the input and output, rather than reaching an absolute pressure on the input side.

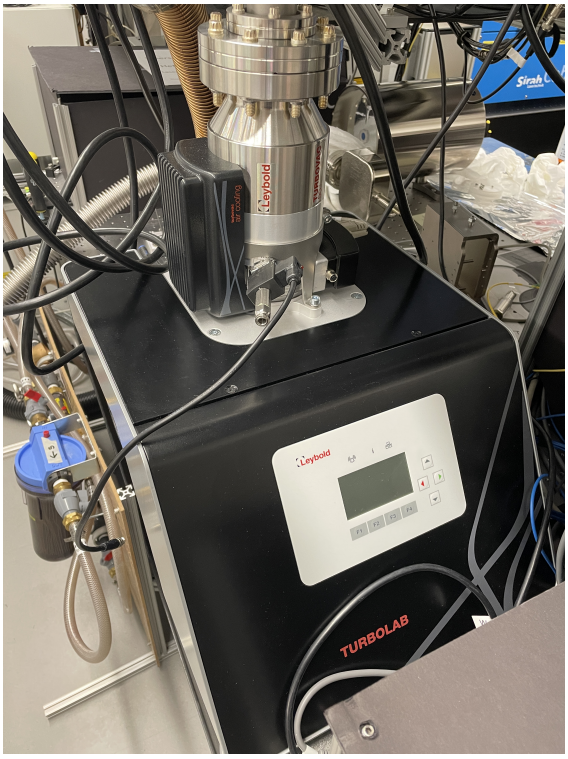

Figure 2.5: Turbo pump used in the vacuum chamber. The model is a Leybold TurboVac 90i. Additional turbo pumps are used in the molecular beam chambers because of low conductance through the skimmers. The base of the pump is 46 cm x 37 cm.

This means that by reducing the output pressure of the turbo pump, we can reduce the pressure on the input end. Therefore, we use less powerful pumps connected to the outputs of all of the turbo pumps. The turbo pumps in the molecular beam chamber and the differential pressure chamber are connected to the same scroll pump, which achieves a pressure of about 10^{-3} torr, and the turbo pump in the ion trap chamber is connected to a diaphragm pump, which achieves a pressure of about 10^{-1} torr. These backing pumps are also necessary because the turbo pumps cannot spin up to full speed at atmospheric pressure, so they bring the system to a lower pressure for the spin up.

2.4.2 Getter Pump

A second type of pump used in our vacuum system is a non-evaporable getter pump. A getter pump functions by the process of chemisorption, in which a chemical reaction occurs that binds gases in the vacuum to the getter surfaces. Titanium, for example, is a common piece of a getter alloy; the binding energy of oxygen chemisorbed on titanium is 1000 kJ/mol, compared to about 15 kJ/mol for physical adsorption [20, section 3.6]. Getter pumps rely on using high amounts of surface area for gas molecules to react with, since the binding happens almost entirely at the surface. The rate of pumping depends entirely on the rate at which the gas reaches the getter surfaces in the pump. In our case, we use the NexTorr D 100 pump, which contains both a getter pump and an ion pump, and contains a series of stacked disks of getter material designed to maximize the surface area [25].

Since the getter material relies on having clean surface area to react with gas molecules, once it is exposed to the atmosphere, the pump becomes nonfunctional: oxygen and hydrogen in the atmosphere form films on the surface, making it impossible for the getter to react with other gas molecules. Fortunately, in order to reactivate the pump, we can raise the temperature of the system. The hydrogen can be removed in what is called a "regeneration," where a temperature of 150°C causes the hydrogen bonds with the getter material to be broken and enter the vacuum system. At that point, the hydrogen is then removed from

Figure 2.6: Getter and ion pump used for the vacuum chamber with the getter surfaces exposed. The getter materials are the disks within the hexagonal cage, while the ion pump is located out of view inside the red box. The dimensions of the red box are 6.5 cm x 7.5 cm x 8.5 cm.

the vacuum system by the turbo pump. Similarly, an "activation" temperature of 500°C breaks the remaining bonds, primarily with oxygen, diffusing the surface layer into the getter bulk [25]. We perform both of these processes during the bakeout of the vacuum chamber. Since the getter pump is reliant on not reaching saturation, as once the surface area is used up it stops pumping, we do not activate it until the turbo pump has taken the pressure of the chamber down to under 10^{-8} torr. It is worth noting that the pump in use is a non-evaporable getter: this means that it contains a bulk getter material that is activated by heat, as compared to an evaporated getter, which uses an evaporated metal surface on tube walls. In our case, the NexTorr D 100 pump has a maximum pumping speed of 100 L/s for H_2 and O_2 [25].

2.4.3 Ion Pump

A third type of pump used in our vacuum system is the ion pump. An ion pump functions by ionizing the gas in the vacuum, which it achieves by introducing a large number of electrons

into the system and then using electric and magnetic fields to cause the electrons to spiral around, ionizing nearby gas molecules. These ions are then attracted to cathodes in the ion pump, removing them from the system. Our ion pump takes advantage of this attraction by also using getter surfaces to absorb as much of the incoming gas as possible. There are difficulties associated with an ion pump: because the ions hit the titanium cathode surface with relatively high energies, in a process known as sputtering, titanium particles and ions deposited on the surface can be ejected [20, section 3.7]. In the case of the NexTorr D 100 ion pump, used in our setup, the getter surfaces in the pump are designed to react with as many of these particles as possible in order to prevent backflow from the pump into the trap system [25]. It is worth noting that the primary purpose of the ion pump in our setup is for pumping noble gases. Once the turbo pump in the main experimental chamber is disconnected, the getter pump is the primary mode of pumping in this section. Since gas may still leak into the system via any of the leak types described above, it is important that we are still able to remove it. The getter pump is very effective at removing most types of gas from the system; however, since it depends on chemical reactivity to remove gases, noble gases, which are highly nonreactive, are not affected. Therefore, the ion pump ionizes and pulls these molecules out of the trap system, allowing us to reach lower pressures. The NexTorr ion pump reaches pumping speeds of 6 L/s for Ar, the primary gas of concern in our system [25].

2.5 Gauges and Detectors

Finally, while we make sure to consider all of these pumping techniques and material properties when designing the vacuum, it is also important that we have a way to measure the pressure inside the chamber, so that we test the effectiveness of the real vacuum. In our case, we use ion gauges for continuous monitoring throughout the course of the experiment, and a residual gas analyzer during the initial pumpdown to get further information on the

contents of the vacuum.

2.5.1 Ion Gauges

We include three different ion gauges in our setup, one in each of the three different chambers (l, m, and n in figs. 2.1 and 2.2). These gauges work very similarly to the ion pump: the gauges release a stream of electrons, which ionize gas atoms and molecules in the volume of the detector. The molecules are then pulled to a charged collector wire, which measures the net current of the ions, allowing it to make an accurate measurement of the number of ions. This can then be converted into a pressure. In our setup, we have Lesker 354 series gauges on the molecular beam and differential pumping chambers, and an NR-F-UHV SRS ion gauge with an SRC IGC100 ion gauge controller on the trap chamber. Additionally, the ion pump in the trap chamber can also be used as an ion gauge. The Lesker gauges have lower bounds of $1\text{x}10^{-9}$ torr [26], the SRS gauge has a lower bound of $2\text{x}10^{-11}$ torr [27], and the ion pump's gauge has a lower bound of $1\text{x}10^{-10}$ torr [25].[3] Although currently the Lesker gauges are not effective because the molecular beam source chamber and differential pumping chambers are at pressures below their lower bounds, once the molecular beam is activated, the pressures will increase and we will be able to measure accurate pressures using the gauges.

2.5.2 Residual Gas Analyzer

During our initial pumpdown, we also use a residual gas analyzer (RGA, o in fig. 2.1) to measure the contents of the trap chamber. Our RGA, an SRS RGA 100, contains both an ion gauge and a quadrupole mass filter. Since we already have ion gauges in each of the three vacuum chambers, the ion gauge on the RGA is redundant and simply confirms our readings from the other trap chamber gauge. The interior of the quadrupole mass filter

[3]The ion pump's gauge can still be used to directly measure the current from charged particles below the lower limit, though this is more useful as a relative measurement than an absolute one.

looks relatively similar to a linear Paul trap (discussed in section 3.1): it has four electrodes in a square, with an RF AC potential on each of the electrodes, with the opposite ones in phase and the adjacent ones 180° out of phase. By tuning these RF voltages (along with DC offsets), the two sets of in-phase electrodes serve as low-pass and high-pass mass filters, respectively [28]. This allows the RGA to measure charge-to-mass ratio for ions in the trap chamber. By measuring the charge-to-mass ratio of ions and the frequency of each mass, we can measure the partial pressures of each type of gas in the vacuum. This is useful for identifying leaks, as described in section 2.2, and for determining whether the bakeout has removed the desired water vapor from the system.

2.6 Pumpdown and Bakeout

Over the course of the early spring, we pumped down the vacuum chamber from atmospheric pressure to under $1\text{x}10^{-10}$ torr. In January, we completed several upgrades to the trap chamber and connected the trap chamber and molecular beam apparatus; following this, we began to pump down the trap chamber first using the scroll pump, and then using the turbo pump. Over the course of about 10 days, the turbo pumps brought the pressures down to $2.7\text{x}10^{-8}$ torr in the trap chamber, $2\text{x}10^{-8}$ torr in the differential pumping chamber, and $3\text{x}10^{-9}$ torr in the molecular beam chamber. At this point, the chamber was prepared for bakeout. To do this, we attached thermocouples to various surfaces to help ensure that we would not damage any components by exceeding their temperature limits, and then wrapped the delicate components in aluminum foil to avoid damage from uneven heating. These included the glass viewports, electrical feedthroughs, and the welded hexapod bellows. In particular, it is important to be careful with the glass because different temperatures on different parts of the viewports may also cause varying degrees of stress on the glass and therefore optical asymmetries, which lead to birefringence. Therefore, we made sure to cover each viewport with aluminum foil without making contact with the glass to ensure

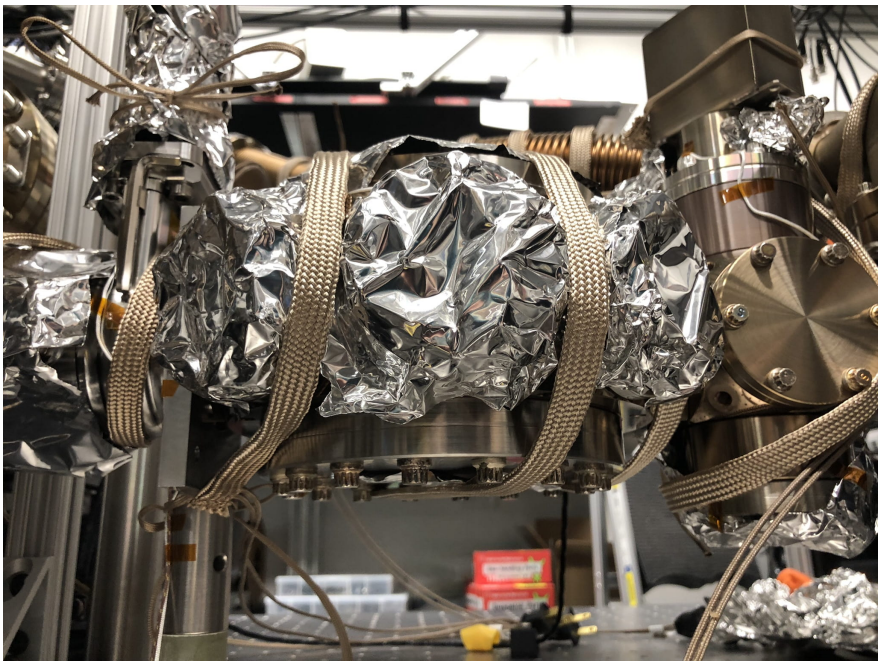

Figure 2.7: Photo of the ion trap chamber in the process of preparing for the bakeout. Windows were covered in aluminum foil to avoid thermal contact with the glass, and then the chamber was wrapped in heating tape. The diameter of the trap chamber is 15 cm.

the windows were heated evenly. We then wrapped the chamber in heating tape, and then aluminum foil and fiberglass insulation to make sure the heating was even and efficient (seen in figs. 2.7 and 2.8).

The bakeout occurred over the course of about 16 days. During the initial heating, the temperatures were raised by up to 5-10°C per hour, up to between 90° and 160°C, depending on the specifications of the components being heated (some of the higher temperature regions had their temperatures raised over the course of several days). The turbo pump in the differential pumping chamber, for example, was only raised to 80°C, while the gate valve between the differential pumping chamber and the trap chamber was raised to just over 150°C. During the bakeout, partial pressures of each of the relevant gases in the trap chamber were carefully tracked to determine when they had reached acceptable levels; in particular, we focused on the water vapor because of the properties discussed above in section 2.2. The partial pressures of some of these gases, measured using our RGA, can be seen in fig. 2.9. Over the course of the bakeout, there were several spikes in these partial pressures: these

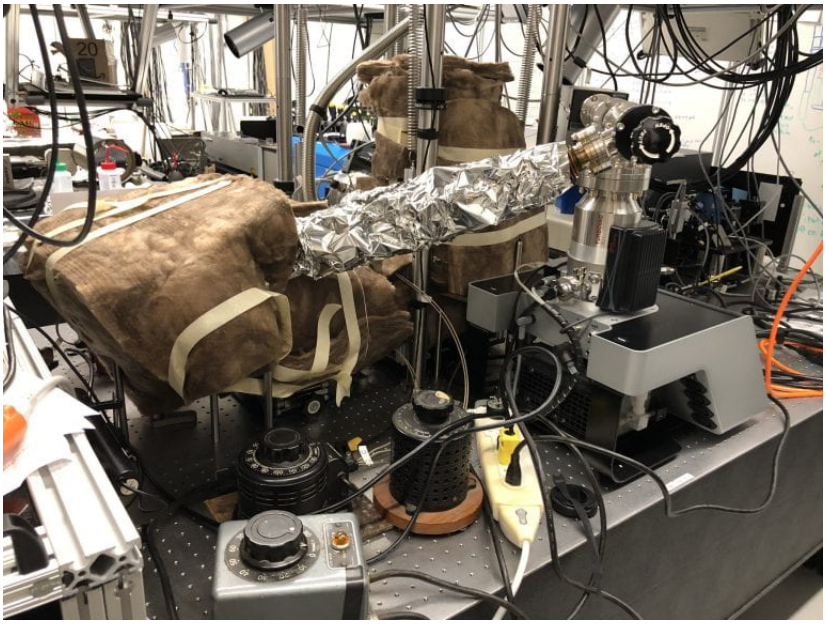

Figure 2.8: Photo of the apparatus fully insulated for the bakeout. Most of the apparatus was covered in aluminum foil, and then in fiberglass insulation to ensure efficient heating. The length of the tube wrapped in aluminum foil is 60 cm.

occurred due to regeneration of the non-evaporable getter, degassing of the ion gauge, and activation of the non-evaporable getter and ion pump. At the conclusion of the bakeout, the pressure in the trap chamber reached 3x10^{-10} torr, in the differential pumping chamber it reached 1x10^{-9} torr, and in the molecular beam chamber it reached 9x10^{-10} torr, although the pressures in the differential pumping chamber and the molecular beam chamber are less reliable because they are close to the lower bound of the ion gauges. When the experiment runs, the molecular beam will introduce more oxygen to the system, so the lower bounds of these detectors will not be limiting to our data collection.

Because the pressure in the trap chamber is on the order of our desired 1x10^{-10} torr, the vacuum conditions are are sufficient to trap ions using the upgraded trap electronics. This means that the vacuum pumping and bakeout was a success, and the current design of the vacuum chamber will be effective for ion trapping in the full experiment.

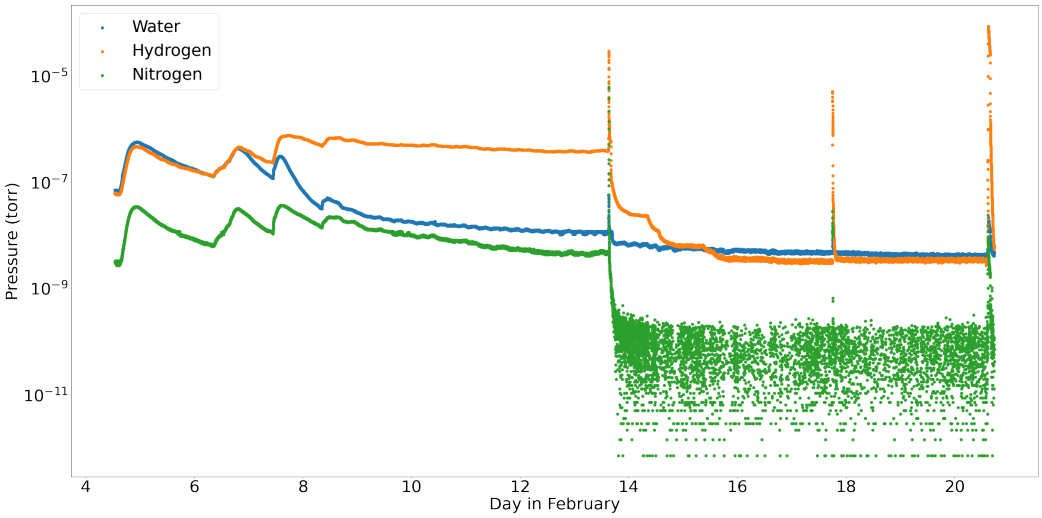

Figure 2.9: Partial pressures for water vapor, hydrogen, and nitrogen during the bakeout. The narrow peaks in the graph, from left to right, occur due to the getter pump regeneration, ion gauge degassing, and getter pump activation.

Chapter 3

The Ion Trap

With a functional vacuum chamber that can provide the conditions necessary for ion trapping to occur, the ion trap itself is the next component required for our experiment. While the most straightforward way to create a trap may seem to be using fixed DC potentials to hold the ions in place, this type of equilibrium is inherently unstable, so it would not be effective for us to use to make our measurements. So, a more complicated radiofrequency (RF) technique is required. In this chapter, I will discuss the behavior of ions inside an RF trap, and then the theory of the electronics for setting up this type of trap. Finally, I will enumerate the specifications for the trap circuit that I built, and how I implemented each of them.

3.1 Paul Trap

In order for ions to be trapped, they need to be inside a potential well, which confines them at a given point in space. To find what this looks like, we can consider a potential Φ at an equilibrium point, using a generic Taylor expansion, reproducing the derivation in ref. [29]. The Taylor expansion contains a term for both a static potential and an oscillating potential in the trap. In this way, the potential at a point in the center of the trap can be written:

$$\Phi(x,y,z,t) = V\frac{1}{2}(\alpha x^2 + \beta y^2 + \gamma z^2) + U\cos(\Omega t)\frac{1}{2}(\alpha' x^2 + \beta' y^2 + \gamma' z^2). \tag{3.1}$$

In this equation, V is a static potential, and U is an oscillating potential. α, β, γ, α', β', and γ' are parameters corresponding to the geometry of the trap. This oscillation comes at frequency Ω, which will usually be in the radiofrequency region.

This potential equation has to fulfill the Laplace condition, which says that $\nabla^2\Phi = 0$ for all time, leading to restrictions on the geometric factors:

$$\alpha + \beta + \gamma = 0 \tag{3.2}$$

$$\alpha' + \beta' + \gamma' = 0 \tag{3.3}$$

These conditions mean that there is no point at which there exists a 3-D potential well for any given time, so a static trap is not possible. Therefore, we will need to trap charges dynamically, hence the RF trap. There are several possible solutions to the above conditions, leading to a variety of different trapping conditions. In one possible solution, we design the trap so that the factors are as follows:

$$-(\alpha + \beta) = \gamma > 0 \tag{3.4}$$

$$\alpha' = -\beta', \gamma' = 0 \tag{3.5}$$

This combination of geometric factors gives us a linear Paul trap, with, specifically by applying the $\gamma' = 0$ condition, an RF null at $x = y = 0$. The Paul trap consists of four near-hyperbolic, linear rods, evenly spaced in a diamond around the trapping well, seen in fig. 3.1. We assume symmetry in the geometry and voltages of the system in order to simplify the calculations. This turns out to be a reasonable assumption for this system; any

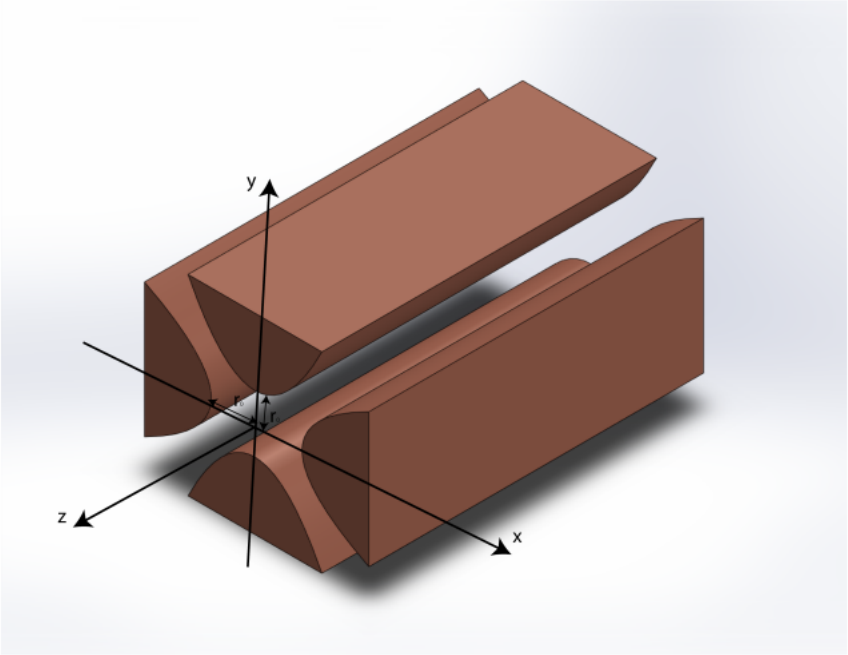

Figure 3.1: Diagram of a linear Paul trap. The trap consists of four near-hyperbolic electrodes, surrounding a trap axis. Each set of opposite trap electrodes is in phase, and the adjacent electrodes are 180° out of phase. Note that our trap is rotated 45° from the image shown, such that the y-axis and x-axis in the image are located in the gaps between the electrodes. Figure from [21], used under the Creative Commons License.

asymmetries can be addressed with some of the later-discussed circuit tuning capabilities.

In our case, each pair of opposite rods is in phase at our radiofrequency, Ω, and the two sets of pairs are 180° out of phase with each other. Equation 3.5 suggests a dynamical confinement in the xy-plane based on the RF frequency, and no RF dependence on the z-direction. In the trap, this means that the particles oscillate in and out of the center of the trap, towards the electrodes. Equation 3.4 implies that the static potential is confining in the z-direction, and anticonfining along the x- and y-axes.

Using the equations of motion for a particle at a given charge $Z|e|$ and mass m at a given potential, we can construct the second order differential equation:

$$\ddot{x} = -\frac{Z|e|}{m}\frac{\partial \Phi}{\partial x}. \tag{3.6}$$

By substituting equation 3.1 and computing the derivative, applying the transformations,

$$\xi = \frac{\Omega t}{2} \tag{3.7}$$

$$a_x = \frac{4Z|e|V\alpha}{m\Omega^2} \tag{3.8}$$

$$q_x = \frac{2Z|e|U\alpha'}{m\Omega^2} \tag{3.9}$$

and isolating the motion in the x-direction, we get the equation

$$\frac{d^2x}{d\xi^2} + [a_x - 2q_x\cos(2\xi)]x = 0. \tag{3.10}$$

This can be solved using standard methods for solving the Mathieu equation. It is worth noting that the solution requires a degree of approximation, but the higher-order terms are insignificant for the values of a_x and q_x relevant to our experiment. The solution gives rise to a region of stability for a_x and q_x that create a trap, which can be seen in fig. 3.2. (We could complete this same process for the y-direction, which gives the same result but rotated by 90°.)

Graphically, the region appears as a kite shape. From a qualitative perspective, q_x can be thought of as a reflection of the potential set up by the driven RF voltage. This induces secular motion with similar behavior to a harmonic oscillator. In the limit of low q_x, this frequency is given by

$$\omega_r = \frac{q_x\Omega}{2\sqrt{2}}. \tag{3.11}$$

These oscillations occur in the radial direction, meaning that the ions move in and out of the center of the trap. Using the definition in ref. [30], we can then write an effective radial

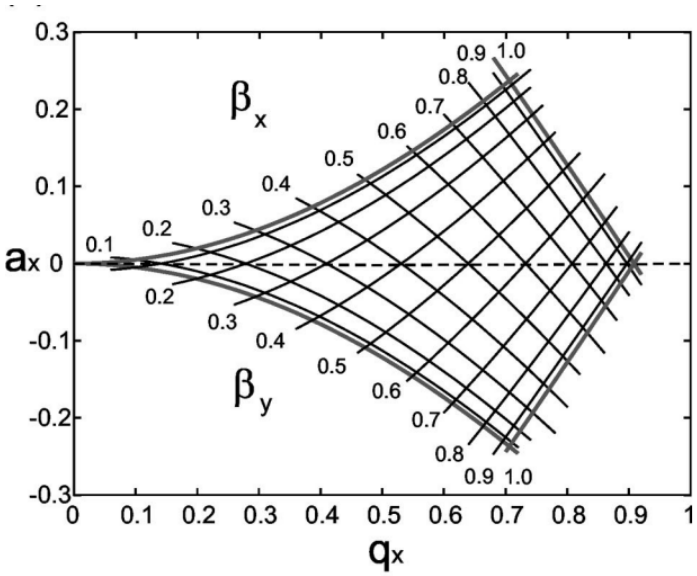

Figure 3.2: Stability diagram for the a_x and q_x quantities in a linear Paul trap. a_x and q_x are both dimensionless factors corresponding to ion motion in the trap. β_x is relevant in the secular frequency terms, and can be written as $\beta_x = \sqrt{a_x + q_x^{1/2}}$. Generally, we want to keep $\beta << 1$ to keep the secular frequency much lower than the RF frequency. Figure reproduced with permission from [29].

potential for the trap as

$$\Phi_r = \frac{m}{2Z|e|}\omega_r^2(x^2 + y^2), \tag{3.12}$$

allowing us to describe the potential at a point in or near the center of the trap.

The a_x parameter can be thought of as a reflection of the potential generated by a static voltage. Referring back to fig. 3.2, we can draw a few conclusions about the types of motion in the trap. At high values of q_x, we lose the harmonic behavior, and, at very high values, the ions may even be ejected from the trap since the q_x value is proportional to the driving amplitude. Therefore, we use $q_x < 0.3$ in most cases; at this value, we need to be careful to keep a_x much lower than q_x. Conceptually, this comes from the fact that if a_x is too large, the anticonfining static potential may be too strong, which causes the ions to fall out of the trap. At low q_x values (where the harmonic secular behavior exists), we also can have micromotion, which occurs as smaller oscillations at the RF frequency when the ions are not exactly at

the RF null; the amplitude of these is a factor of $q_x/2$ smaller than the secular amplitude. Therefore, in order to confine the ions reliably in the trap and to minimize micromotion, we want to stay in the region in the center of the stability diagram, with $q_x < 0.3$ and a_x close to 0. In particular, q_x can be expressed in the equation

$$q_x = \frac{2eU}{m\Omega^2 r_0^2}. \tag{3.13}$$

By plugging in the relevant values for the trap, including the radius of the trap (r_0 =1.25 mm) and the mass of oxygen and beryllium ions, we get that the RF potential must be above about 50-100 V to keep q_x in the region we desire. Additionally, we can see that q_x, and therefore ω_r (eq. 3.11), depends on the charge-to-mass ratio of the particles; this is relevant in section 3.4.3 when including AC modulation in the trap.

Excess micromotion can be caused by several different factors. These can include a phase difference between the driven RF signal at electrodes in the trap or stray fields in the trap that cause asymmetries in its geometry. In general, micromotion makes it so that our simple harmonic motion approximation is inaccurate, so we want to minimize this to the greatest extent possible. Excessive micromotion can add excess energy to the ions that must be removed by laser cooling, and also adds Doppler shifts to measurements. The Doppler shifts are especially significant for the measurement of variations of μ. First-order Doppler shifts are not significant, since the ions are confined tightly in the ion trap. Second-order shifts, however, are important; these occur due to relativistic motion of the ions from micromotion. The ions have a nonzero v_{rms}, and therefore have a time dilation effect on the distance to the laser source [13]. This is anticipated to be the largest source of error in our measurement, which is surprising given the velocity of the ions from micromotion: since the frequency of the micromotion is on the order of 10 MHz, and the distance of the motion is on the order of 10 microns, the velocity of the ions is under 100 m/s. This gives a sense of the precision of the measurement: relativistic effects at speeds less than one millionth of the speed of light are among the most significant sources of error. Therefore, it is important that we can both

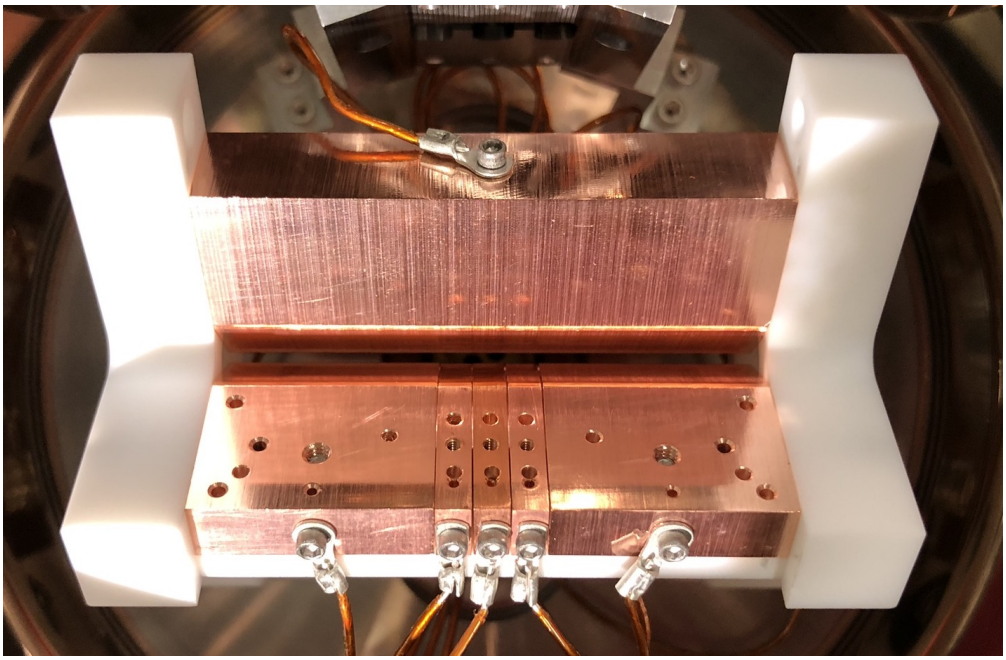

Figure 3.3: Photograph of the ion trap from above, showing a solid electrode (top) and segmented electrode (bottom). On the segmented electrode, the three segments in the center are 3 mm each, and the two longer segments on either side are 20 mm.

minimize and quantitatively parameterize this effect for the trapped ions. In later sections, I will detail the ways in which we safeguard against this micromotion.

The radial motion is not the only type of motion in the trap, though. In order to confine the ions in the axial direction, static potentials are introduced. To do this, we split two of the trap electrodes (positioned diagonally opposite to each other) into five subsections, each with the same RF but at possibly different static potentials (see fig. 3.3). This allows us to have a potential minimum at the center of the trap, which keeps the ions confined in the axial direction. Using the derivation in ref. [30], we can approximate this static potential as

$$\Phi_s = \frac{m}{2Z|e|}\omega_z^2[z^2 - (1/2)(x^2 + y^2)], \tag{3.14}$$

where m is the mass of the ion, Z is the charge of the ion, and ω_z is the axial frequency, which can be written as

$$\omega_z = \sqrt{\frac{2\kappa Z|e|V}{mz_0^2}} \tag{3.15}$$

In this equation, V is a static potential, as in equation 3.1. z_0 is typically taken as half of the length of the central electrode of the trap; in our case, since the central electrode is 3 mm, z_0 is 1.5 mm (see fig. 3.3). Finally, κ is a geometric factor for the trap, which depends on the specific dimensions and symmetries as well as the static potential distribution. In our trap, for a typical distribution of static potentials, κ is simulated to be about 0.170. The distribution may be changed depending on the specific experimental conditions in the trap, so κ is measured or simulated at the time of measurement.

In addition to providing axial confinement, this static potential slightly weakens the strength of the confinement in the center of the trap. This slightly reduces the frequency of the secular oscillations in the radial direction; we can write this new frequency as:

$$\omega_r' = (\omega_r^2 - (1/2)\omega_z^2)^{1/2}. \tag{3.16}$$

This is how the a_x and q_x parameters connect in driving the secular frequency of the ions. Then, returning to the potential in the radial direction, we can rewrite equation 3.12 as

$$\phi_r = \frac{m}{2e}(\omega_r')^2(x^2 + y^2), \tag{3.17}$$

giving us a more accurate picture of the potential at a given point in space near the center of the trap. As in equation 3.12, this depends on a frequency and the mass-to-charge ratio of the trapped ions.

In conclusion, we are able to accurately predict the behavior of the ions in the trap by driving an RF signal on the trap electrodes, and introducing static potentials to confine the ions along the trap axis. This means that by controlling these quantities in a circuit, we can precisely manipulate the behavior of ions in the trap to achieve the specific experimental objectives.

3.2 Circuit Principles

In order to implement these trapping principles to create an ion trap, a physical circuit is required to drive the RF and static potentials. Put most simply, the trap circuit functions as an LC resonator. Although other capacitors are present, the primary source of capacitance in the circuit is the trap itself. Since the four electrodes in the trap can be looked at as hyperbolas or rods of charge at a given point in time, there is a capacitance between each set of them. In our case, the capacitance between the trap electrodes sum to about 20 pF. Since the four sets of adjacent trap electrodes behave as if they are in parallel, each of the pairs has a capacitance between them of roughly 5 pF. (There is capacitance between opposite electrodes as well, with a reduction factor of $\sqrt{2}$ from the larger capacitance, meaning that the actual capacitance of each electrode gap is slightly less than 5 pF.) This can be seen in fig. 3.4; the circuit on the left shows the general structure of the circuit, with a step-up transformer driving an RF signal across C1, which represents the trap. The more detailed view on the right shows the behavior of the electrodes in the trap; electrodes 1 and 2 are in phase with each other since they are shorted together, and the same is true for electrodes 3 and 4.[1] Each of the gaps between electrodes, represented by C2-C5, can then be thought of as a 5 pF capacitor.

Since we want a q parameter of roughly 0.3 or lower (see eq. 3.13), the voltage across the electrodes must be at least on the order of 50-100 V for the trap to function properly. This would be difficult to drive using a typical 50 Ω source because it would require a very large amount of power. Therefore, we instead use a step-up transformer to reach the necessary voltage, which does not require significant current because it is largely just driving a voltage across capacitors. Since the secondary coil is simply a coil of wire in the RF circuit, it also

[1] As drawn, there is no capacitance between the opposite electrodes, since they are shorted together. The capacitance comes from the circuit concepts introduced in section 3.3.

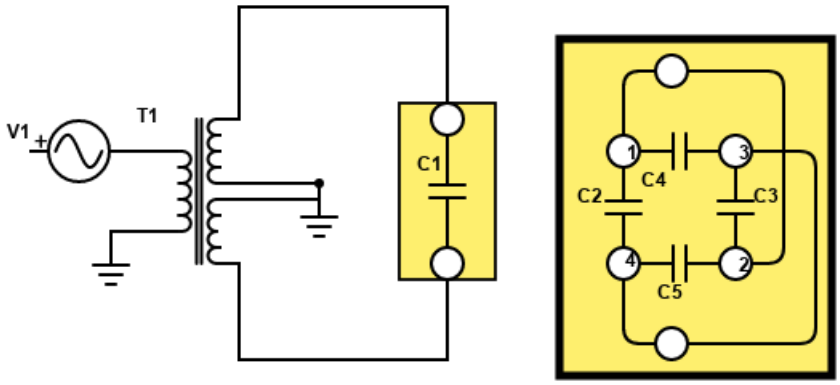

Figure 3.4: Simplified circuit diagram for driving RF potential across a set of two electrodes in the trap. C1 is a representation of the trap, and T1 is a step-up transformer. The portion on the right represents more detailed behavior of the four electrodes in the trap. C2, C3, C4, and C5 all behave like 5 pF capacitors.

behaves as an inductor; the resonance of the circuit, then, can be written as

$$f_0 = \frac{1}{2\pi\sqrt{LC}}, \tag{3.18}$$

where L is the inductance of the coil and C is the total capacitance of the circuit. Ideally this RF resonance is as high as possible (in our case, this means around 15 MHz because of physical constraints), since this corresponds to more tightly trapped ions, and we need the entire crystal to be within the waist of the laser that drives the oxygen vibrational transition. Since we want $q < 0.3$, and the mass and charge of the ions is fixed, the radiofrequency is the only tunable parameter that allows us to raise our radial frequency. An RF above about 10 MHz is suitable for this purpose, but staying closer to 15 MHz will allow us greater flexibility and margin for error, since we can turn down RF voltages later if we want slightly looser trapping. This means that we want to minimize the inductance and capacitance in the circuit.

For a given transformer core, the inductance can be expressed as

$$L = A_L(N_{\text{turns}})^2, \tag{3.19}$$

where N_{turns} is the number of secondary coil windings, and A_L is the inductance constant of the core [31]. The inductance constant is a factor of the core material and geometry. Since we want a lower inductance, we want to use relatively few secondary coils (note the tradeoff here between step-up behavior and inductance), and a low inductance constant. Indeed, the transformer core, a ferrite toroid, was chosen primarily for its relatively small A_L value. Specifically, we used a Micrometals T184-0, which has $A_L = 3$ nH/N^2 [32]. It should be noted that it is conceivable that a lower-A_L coil could be used; however, we were unable to get sufficient coupling when this was tested. We wrapped the transformer 3 times with a primary coil, and 16 times with each of two separate secondary coils, wound in opposite directions. These coil counts were found by impedance matching the circuit using a network analyzer–the primary coil counts were varied, as shown in fig. 3.5, with the goal of finding the value at which the greatest percentage of power from the RF source was passed into the circuit, which corresponds to the least power reflected back to the voltage source. Since 3 primary coil windings produced the least amount of reflection, this value was selected. Therefore, the coil is expected to have an inductance of about 819 nH.

Other transformer cores, wire types, and secondary coil counts were tested, and it is still possible that different combinations may be able to drive the circuit at a resonance closer to the desired 15 MHz frequency; the current coil runs roughly between 10.8 and 12.5 MHz, depending on tunable components in the circuit. The two secondary coils are counterwound in order to get the out-of-phase behavior we want on the adjacent trap electrodes; by winding the transformer coils in opposite directions, we expect the two sides of the electrode to be exactly 180° out of phase with each other.

Therefore, with by taking advantage of the LC resonance of the circuit and introducing

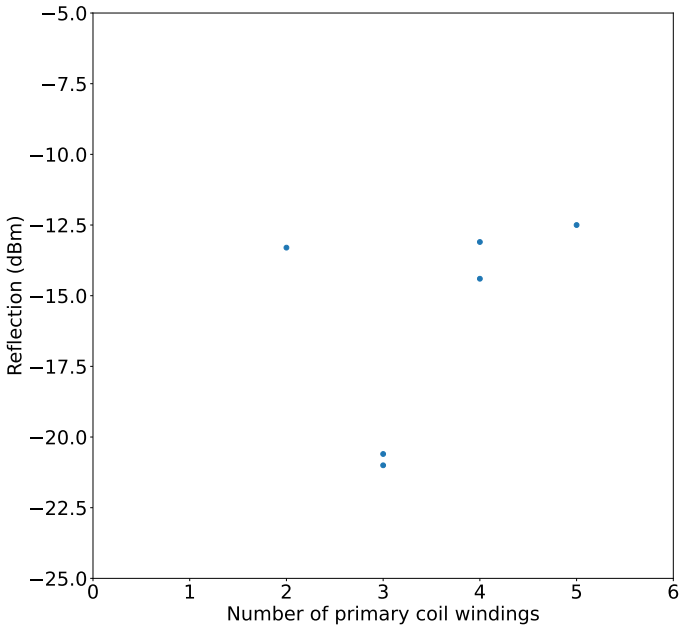

Figure 3.5: Results of impedance matching for primary coil count. Each point was measured at resonance, which varied between 12.56 MHz (for 2 coils) to 12.86 MHz (for 5 coils). Data was collected using a Keysight E5063A network analyzer.

counterwound secondary coils, we can drive high-voltage signals on each of the trap electrodes, with the adjacent electrodes 180° out of phase with each other. From here, we can add more functionality that will address various specifications and allow for fine-tuning of the physical trap circuit.

3.3 Positioning Ions

In our derivations for the behavior of ions inside of the trap, we assume that the motion of the ions is centered in the middle of the trap electrodes, meaning that on average, there is an equal distance from the ions to each of the four trap electrodes. This assumption adds symmetry to these calculations, making the computations significantly simpler. But for the purposes of the experiment, this is not always possible. The ions need to be at the focal point of the laser and molecular beams in the center of the ion trap, and as described in section 1.7, this is incredibly difficult to achieve mechanically, since both the laser waist and the ion crystal width are on the scale of tens of microns. There are a few options for adding the fine tuning necessary to resolve this issue: moving the laser by adjusting the mirrors in the trap chamber, or using electric potentials to move the ions themselves. Moving the mirrors is conceptually simpler than adjusting the potentials, because it is a mechanical process rather than an electronic one. However, the optical cavity that the mirrors form is located inside of the trap's vacuum chamber. Therefore, the parts for moving the mirrors must be able to work at the high-vacuum environment, discussed in chapter 2. Because of this restriction, these parts turn out to make this method cost-prohibitive for our purposes. The best solution, then, is to add tuning methods to the RF circuit that allow us to move the position of the ions.

This tuning can be achieved by changing the relative potential between two in-phase electrodes, which moves the trap's equilibrium point away from the center. To simplify the derivations for this, I will focus on a shift in voltages along one diagonal axis, based on the

derivation in ref. [33]. I label two opposing trap electrodes A and B, and assume that the RF amplitude of one is greater than the other. This means that for some δ, which we usually take to be close to 0, we have

$$U_B = (1 - \delta)U_A. \tag{3.20}$$

We assume $\delta > 0$, which means that $U_A > U_B$. For the purposes of the circuit, we simply can relabel whichever electrode we desire to have higher voltage as A.[2] Using this adjustment, and taking r_0 to be the distance from the trap's midpoint to each of the electrodes and x to be the position of the ions (with $x = 0$ at the midpoint, and the electrodes at $\pm r_0$), we get a new trap potential:

$$\Phi(x) \propto (U_A + U_B) - (U_A - U_B)\frac{x}{r_0} + (U_A + U_B)\frac{x^2}{r_0^2} + ... \tag{3.21}$$

We omit higher-order terms for simplicity; for small x values (relative to r_0), they are negligible.[3] By solving for x values with the minimum potential, and then substituting the above δ relation for the two potentials, we get:

$$x_0 = \frac{U_A - U_B}{U_A + U_B}\frac{r_0}{2} = \frac{\delta}{2 - \delta}\frac{r_0}{2}, \tag{3.22}$$

where x_0 is the updated RF null. This means that, taking the same assumption as above that δ is close to 0, we have that $x_0 = \frac{\delta r_0}{4}$. Since the trap radius is 1.25 mm, we see from fig. 3.6 that we expect to be able to move the ions by over 200 microns with $\delta = 0.5$.

Therefore, by adjusting the voltages on the trap electrodes, we can control the position of the potential saddle point, changing the position of the trap axis. This means that the ions will be able to sit exactly at the position in the trap where the laser is, so we will be able to achieve high efficiency laser contact with the ions.

[2] In ref. [33], δ is defined slightly differently, corresponding to my definition of $1 - \delta$.

[3] We never get a positional shift of more than about 15%, so this approximation is reasonable.

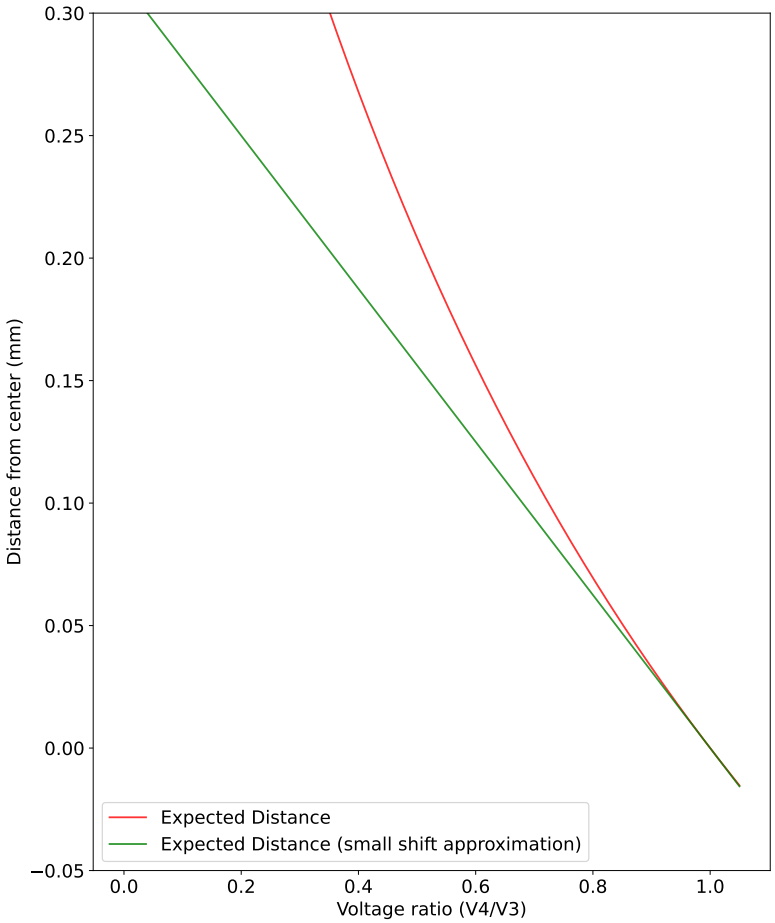

Figure 3.6: Theoretical effects of RF potential changes on ion equilibrium position in our trap (with $r_0 = 1.25$ mm). The green line shows the small-voltage shift approximation (accurate near $\delta = 0$), and the red line shows the full model.

3.4 Circuit Design

Beyond the simple case discussed in section 3.2, the circuit requires several additional components that will address various design constraints, primarily to increase the experimental efficiency. These components include an apparatus to move around ions in the trap without causing excessive micromotion using the strategies described in section 3.3, tools that allow for measurement of the circuit to test its behavior while it traps ions, AC modulation to test properties of the trapped ions, and the addition of static potentials that can correct asymmetries in the trap as well as providing axial confinement. There are also some design specifications for the circuit: it needs to work at potentials at or above 100 V and frequencies close to 15 MHz, as stated above. It also needs to fit into a circular area of roughly 15 cm diameter under the trap chamber.

3.4.1 Ion Positioning

As stated in section 3.3, by reducing the potential across one of the trap electrodes, we can push the saddle potential away from that electrode, allowing for control over the ion position. To physically achieve this, we add a variable capacitor in series with each of the trap electrodes. In this way, the voltage from the RF input is partially dropped over the variable capacitor, leaving less to drop over the trap electrode.

Quantitatively, the potential across a trap electrode looks like

$$U' = \frac{U_{in}}{1 + \frac{C_t}{C_s}}, \tag{3.23}$$

where U' is the updated trap potential, U_{in} is the input RF signal coming from the transformer, C_t is the capacitance of the trap, and C_s is the capacitance that we introduced in series [33]. This directly follows from applying rules for adding capacitances in a circuit.

In practice, we want to have access to a wide range of capacitances, and be able to fine tune the capacitance even once the circuit is running so that we can manage the positioning

of the ions with the most precision possible. This means that we use variable capacitors to adjust the voltage across the electrodes. In general, we want to have our range of capacitance extend from about 5 pF (which corresponds to a displacement value of $0.33r_0$), to about 1000 pF (which corresponds to a displacement of about $0.003r_0$). This level of precision is necessary because of the small cross section between the ion crystal, the molecular beam, and the waist of the lasers.

Since the capacitors need to work at high voltages and frequencies, only a few capacitor types are effective. In particular, ceramic capacitors have been tested to work extremely well at our full range of conditions. Unfortunately, ceramic variable capacitors are not produced for ranges of 5 pF to 500 pF. Therefore, we instead create our own variable capacitor by putting five smaller capacitors in series with the electrode, and in parallel with each other, with clips allowing us to "turn on" different capacitors by attaching the clips.

In our case, the capacitors, which can be seen in fig. 3.7, have the values 500 pF, 200 pF, and 100 pF. The last two capacitors are each 8-40 pF variable capacitors. A smaller capacitor (2-6 pF) was tested to get to the desired 5 pF performance; however, there was no significant difference between this and the 8-40 pF capacitor.

3.4.2 Circuit Testing

Generally, circuits can be monitored and tested by measuring the voltage at desired points in the circuit. However, monitoring an RF circuit is not this simple. The existence of the LC resonance means that the introduction of additional components will add capacitance to the circuit, and therefore push down the resonant frequency. Indeed, any monitor will likely use a 6 foot BNC cable to connect the circuit to a scope, which adds capacitance in parallel–the measured cable had a capacitance of 216 pF. This means we would need to either let the monitor remain in the circuit perpetually, and therefore push down the resonance (which we want to avoid), or only add it when we take measurements, which would force us to change the frequency and therefore not gain an understanding of the circuit behavior at resonance.

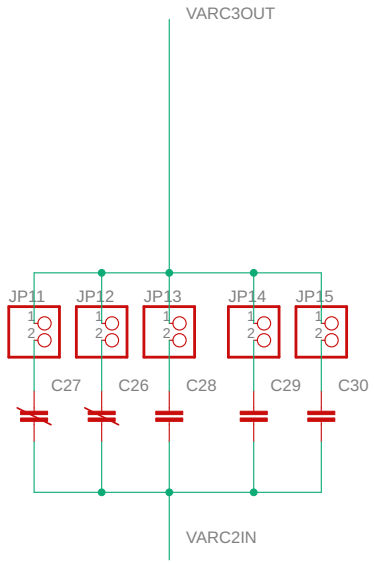

Figure 3.7: Schematic of the variable capacitor used to change the net capacitance of each branch of the circuit. C27 and C26 are 8-40 pF variable capacitors, C28 is a 100 pF capacitor, C29 is a 200 pF capacitor, and C30 is a 500 pF capacitor. JP11-15 are jumpers, allowing for fast adjustment of the capacitance.

So, we need to add a component that will allow for circuit testing without affecting the performance significantly.

In circuit terms, this looks like the schematic shown in figure 3.8. The test point portion of the circuit is a capacitive voltage divider; the resistor and capacitor together take $\frac{1}{100}$ of the circuit voltage to the test point. C1 is a variable capacitor (1.5-3 pF) because is important to tune this divider to the desired behavior precisely based on the total capacitance of the cable, scope, and electrode circuit branch. It should be noted that the values for R1 and C1 were calculated to also take into account the capacitance of the BNC cable and scope used to measure the circuit. C1 is expected to be about 2.5 pF, and R1 is 100 MΩ.

The rest of the diagram consists of the BNC cable and the scope. The BNC cable, as stated above, has a capacitance of 216 pF, represented by C2. The scope includes a capacitance of 11 pF, and a resistance of 1 MΩ, represented by C3 and R2, respectively. Both of these are in parallel with an ideal voltmeter, which gives us our measurement for

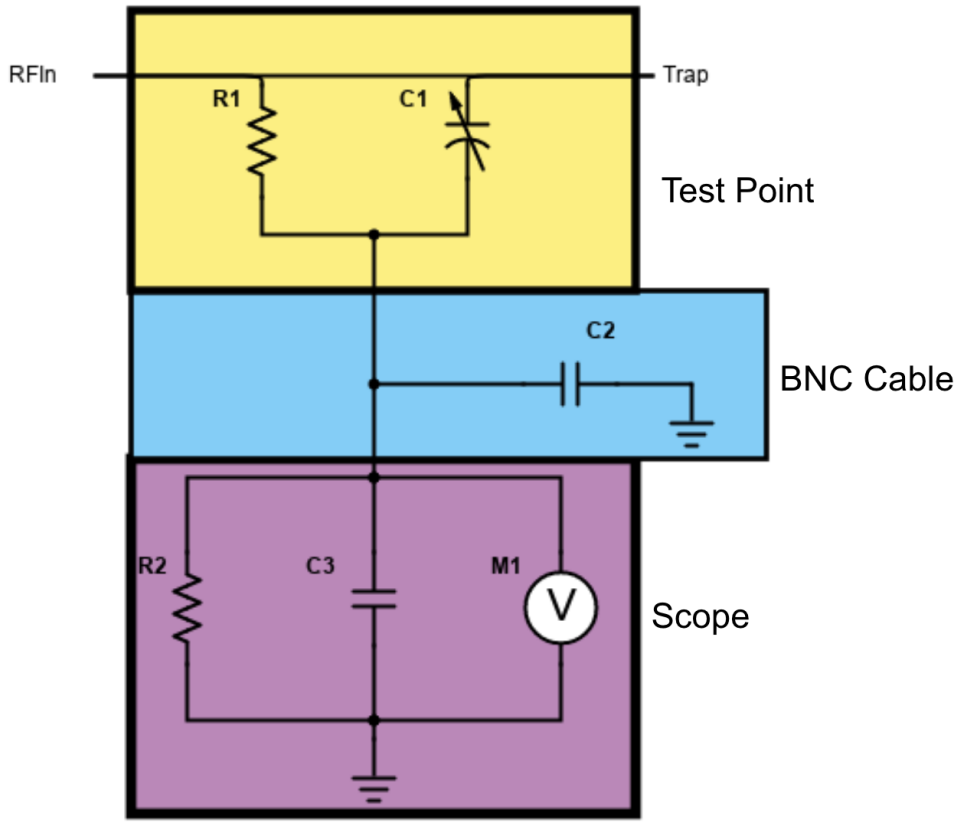

Figure 3.8: Circuit test point and voltage measurement components. Components are split into a test point, which is built into the circuit, a 6 foot BNC cable, and an oscilloscope.

the circuit potential at the test point. This allows us to continually track the behavior of the circuit while ions are trapped.

3.4.3 Voltage Offsets and Modulation

The final notable part of the circuit is the potential offset functionality. The offsets serve several purposes; the first main one is the ability to use static potentials to confine the ions along the trap axis, as described in section 3.1. Additionally, having the ability to add DC offsets to any of the electrodes or electrode segments gives us the ability to correct asymmetries in the trap geometry. Asymmetries may arise due to minor defects in the trap electrodes and/or buildup of charge on insulators, and cause the ions to display micromotion, which we want to minimize. Furthermore, we will use AC modulation, referred to as "tickling." This AC signal is driven at a significantly lower voltage and frequency than the RF signal. Although this modulation can be used to push ions between electrodes, this is not a good way to tune their position because it pushes ions away from the RF saddle point, causing micromotion. The primary purpose of this AC modulation is instead is measure secular and radial frequencies by pushing the ions resonantly at those frequencies. Radial frequency depends on ion mass and charge (see eq. 3.17), so we can use this to measure mass-to-charge ratios of particles in the trap to confirm we are trapping the desired O_2^+ ions. In particular, as stated in section 1.7, once the transition frequency has been measured, an additional laser dissociates the ion into two atoms, one of which is an ion; since this remains in the trap, we can measure its charge-to-mass ratio, which allows us to check that the desired vibrational transition took place (since if it did not, the dissociation laser would not have the correct energy to split the molecule). It is also useful to measure the κ parameter in this manner with the axial modulation. The AC modulation is included for two of the five segments on one of the segmented electrodes, and for one of the solid electrodes, so that we can push along each of the x-, y-, and z-directions.

This can be seen in fig. 3.9, in the branch containing the DC and AC input signals. Note

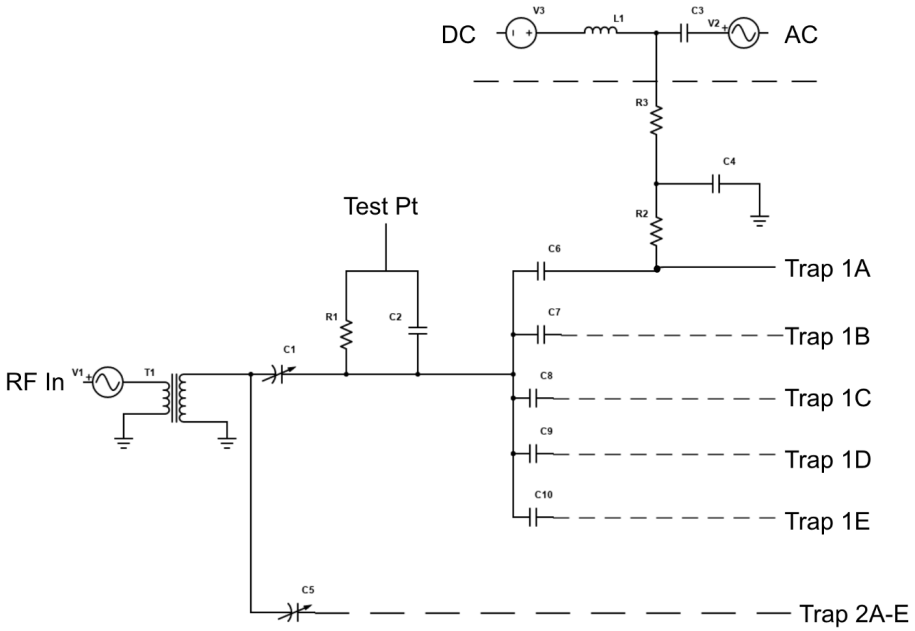

Figure 3.9: Simplified schematic for a pair of in-phase electrodes. Note that the branch containing the DC and AC offsets is replicated in each of Trap 1A-E, and that Trap 2A-E replicates the design from Trap 1A-E.

that this branch is omitted for simplicity in each of the other four trap electrode outputs; the only difference is that only a few of them contain the AC input (used for the modulation) and the corresponding capacitor (C3). This capacitor functions as a block for the DC signal coming from the other input; for the AC signal, the capacitor behaves like a wire, but for the DC signal, it behaves like an open circuit. Similarly, the inductor L1 behaves like a wire for the DC signal, and an open circuit for the AC signal. Together, these force the DC and AC signals down into the circuit, instead of interfering with the other voltage source. All of the DC offsets pass through an inductor to maintain consistent performance, even in the branches without AC modulation. In this case, L1 is 250 μH, and C3 is 10 μF. On the actual circuit, these components are placed on an auxiliary board, away from the rest of the trap circuitry. This is purely due to spatial concerns. We want to put as much of the RF circuit as close to the ion trap as possible to minimize the capacitance introduced by cables, which pushes down the resonant frequency. However, there is limited space under the trap (see

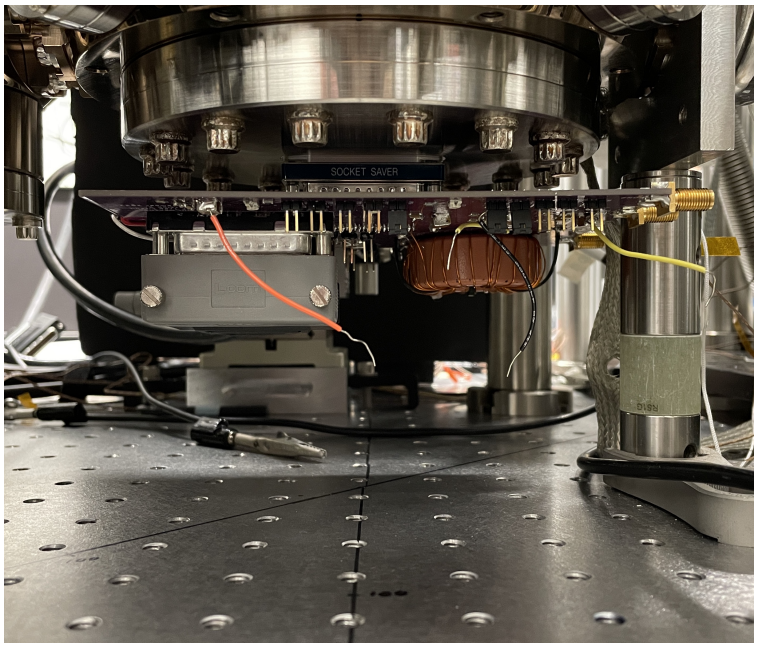

Figure 3.10: Photo of the trap chamber and circuit system from the table level. Note that the total clearance under the trap chamber is 11.0 cm, the clearance under the circuit board plane is 8.8 cm, and the area under the trap chamber is roughly circular with a 15 cm diameter.

fig. 3.10); the area is circular, with a diameter of about 15 cm, and an 11 cm depth. Since the static DC offset and the AC modulation do not affect the circuit resonance, we put these on an auxiliary board, and then connect a cable to the rest of the circuit.

Returning to the main board, we now reach the next set of circuit components, R2, R3, and C4. These form two separate lowpass filters, with one filtering the signal going down into the trap and one filtering the signal going up into the voltage sources. In this case, R3 is a 681 Ω resistor, R2 is a 100 kΩ resistor, and C4 is a 1 nF capacitor.[4] This means that the downwards filter allows frequencies below $f_0 = \frac{1}{2\pi}\frac{1}{RC} = 234$ kHz, and the upwards filter cuts frequencies above $f_0 = 1592$ Hz. The more important of these filters is the downwards filter; this prevents noise from the voltage offsets from reaching the trap and affecting the RF resonance. The upwards filter prevents the RF signal from reaching the voltage sources. This is not integral to the functionality of the circuit, but serves to protect the voltage sources.

[4]To maintain equivalent impedance in the segmented and solid electrode branches of the circuit, R2 is 20 kΩ and C4 is 5 nF in the solid cases. This changes the filter cutoff for the downwards filter, but the change in behavior is not significant for our purposes.

One unintended part of the design here is the inadvertent inclusion of an additional low-pass filter blocking the AC modulation signal. The 100 kΩ resistor (R2) and the variable capacitor for each of the branches (C1) form this filter because the signal through the capacitor is then grounded through the transformer coil. Using the same relation as above, for the largest anticipated value of C1 (about 700 pF) and R2 (100 kΩ), we get a cutoff frequency of 2.27 kHz. Although this would seem to be an issue, since our modulation signals may be over 1 MHz, the attenuation due to this filter turns out to not matter because the ions need so little modulation to display the desired behavior (see sections 4.1.4 and 4.2.4).

Finally, we have the remaining capacitors, which I refer to as the blocking capacitors. In our case, these are 22 μF capacitors, and serve a similar purpose to C3 in this circuit. They block all of the DC signal from reaching the branches corresponding to other parts of the electrode, while allowing through the RF signal. This is important because we want to keep all of the parts of the electrode at the same phase and frequency, but at different static potentials; if we did not have these capacitors, this would be impossible, since the electrodes would be shorted together. These capacitors are labelled C6-C10 in the diagram. The variable capacitors described in section 3.4.1 (C1, C5), and the test point described in section 3.4.2 can also be seen in the diagram.

This is a full summary of the main components of the circuit. A full image of the printed circuit board design is shown in fig. 3.11, and the completed physical board in fig. 3.12. The full schematic can be seen in appendix 5.2. The schematic looks very similar to fig. 3.9, but with branches corresponding to all four electrodes. Two of these contain the same components as fig. 3.9 (except for the slightly changed filter values) and power the long electrodes, but only one branch leading into the trap. These two are powered by one direction of the transformer winding, and the two that power the segmented electrodes look exactly like the one shown in fig. 3.9, and are powered by the winding in the other direction. Additionally, a PCB diagram of the auxiliary board (fig. 3.13) and the completed circuit with components (fig. 3.14) are also included.

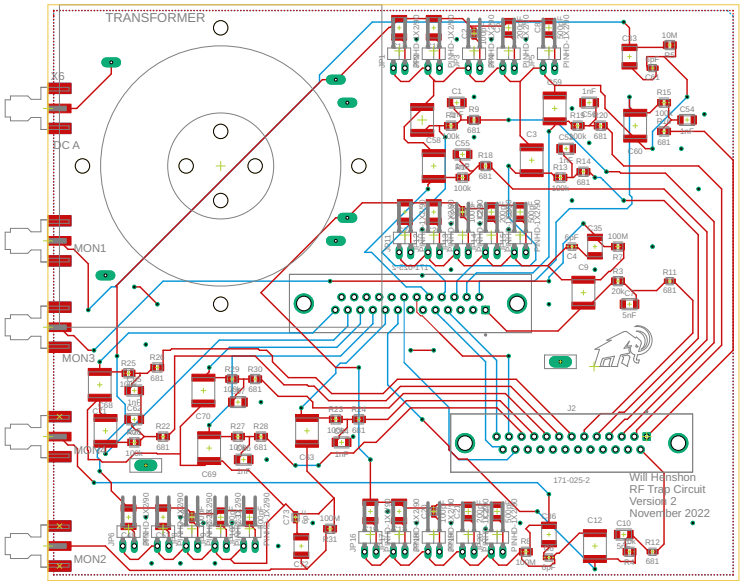

Figure 3.11: Design of the completed printed circuit board for the trap circuit. The design was created in Eagle. Dimensions of the board are 15.24 cm x 12.7 cm.

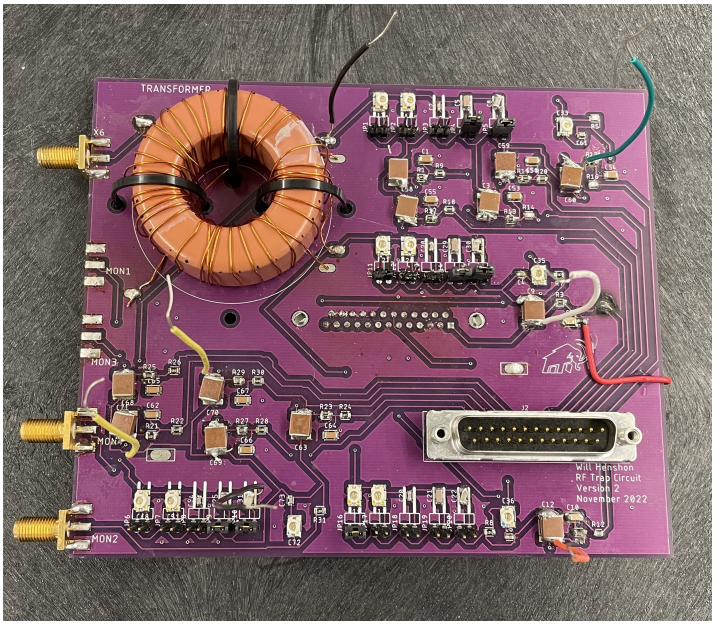

Figure 3.12: Fully built printed circuit board with components for the trap circuit. Dimensions of the board are 15.24 cm x 12.7 cm.

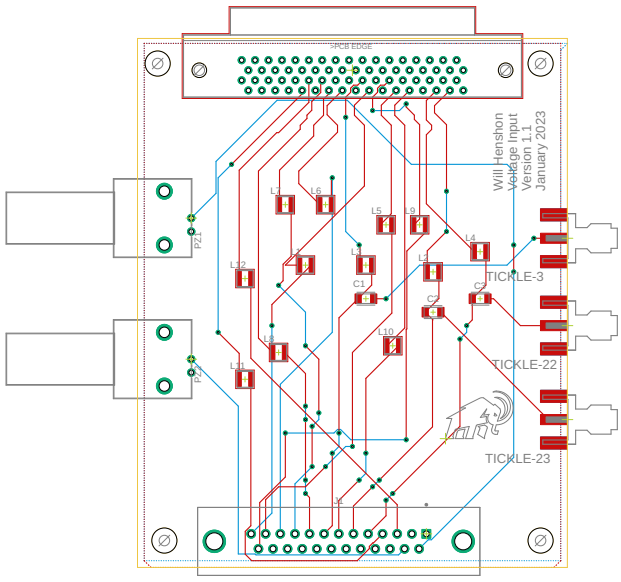

Figure 3.13: Design of the completed PCB for the auxiliary circuit. Dimensions of the board are 7.54 cm x 10.16 cm.

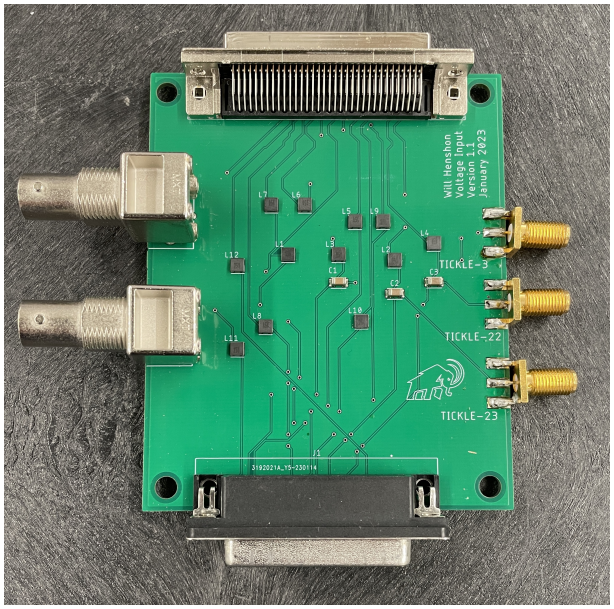

Figure 3.14: Fully built printed circuit board with components for the trap circuit. Dimensions of the board are 7.54 cm x 10.16 cm.

Chapter 4

Results

Once the circuit was constructed, it was important to gather detailed data on its behavior so that we had a strong understanding of how the ions would behave in the trap when using each of the circuit functions. Therefore, the circuit setup, with both the auxiliary board and the RF board, was tested extensively using the trap apparatus. Then, the circuit was used to trap ions, and the circuit features were tested based on how they affected the ion behavior. The results below detail the parameters of the circuit and how it can be used for experiments in the future as well as how it worked for the trapping of Be^+ ions.

4.1 Circuit Parameters

4.1.1 DC Offset

The most straightforward part of the circuit to test is the DC offset. To do this, wires were attached to the RF circuit board directly at the pins of each of the twelve trap contacts (5 each for the segmented electrodes, 1 each for the solid electrodes), allowing for the direct measurement of the DC offsets on the trap electrodes. Voltages were then applied to each electrode using a PXI-6723 voltage source. The potentials at each electrode were measured with both an Agilent Technologies DSO-X 2004A oscilloscope, and a Fluke 179 Multimeter;

while we confirmed that the potentials at each electrode were applied correctly, we did find that the scope was less accurate than the multimeter for measuring these voltages, which may be important to consider for future data collection. When accounting for the impedance of the multimeter compared to the impedance in the filters on the circuit board, we found that the measured voltages were exactly as expected.

4.1.2 Voltage Limits

In order to push the ions into the desired location in the trap, as described in section 3.4.1, we need to be able to produce a wide range of voltage differences across the diagonally opposite trap electrodes. In order to measure these voltage shifts, the circuit was tuned to resonance at a wide range of capacitance values, according to the resonance of the coil being measured. Examples of the circuit potentials at two of the capacitor values can be seen in fig. 4.1. For the solid electrodes, the resonances of the circuit ranged between 10.95 and 11.57 MHz, and for the segmented electrodes, these ranged between 10.99 and 12.10 MHz, both depending on the capacitance. These resonance variations can be seen in the bottom plot of figs. 4.2 and 4.3. Qualitatively, this is as expected: the resonance should go up when there is less capacitance in the circuit.

As can be seen in figs. 4.2 and 4.3, in both cases, we can vary the voltages on each electrode from ratios of about 0.3 to 1.0 (we can go above 1.0, but this is equivalent to the inverse ratio on the other electrode). From this ratio, we can predict the range at which we can push the ions in the trap, using the models described in section 3.3. It should be noted that with careful tuning of the variable capacitors, we can get potential ratios under 0.3; however, this leads to unstable trap conditions (see section 4.2.3).

4.1.3 Phase Issues

There are some difficulties associated with changing the voltages across the electrodes in this manner. Initially, when testing this method, we assumed that the changes in resonance

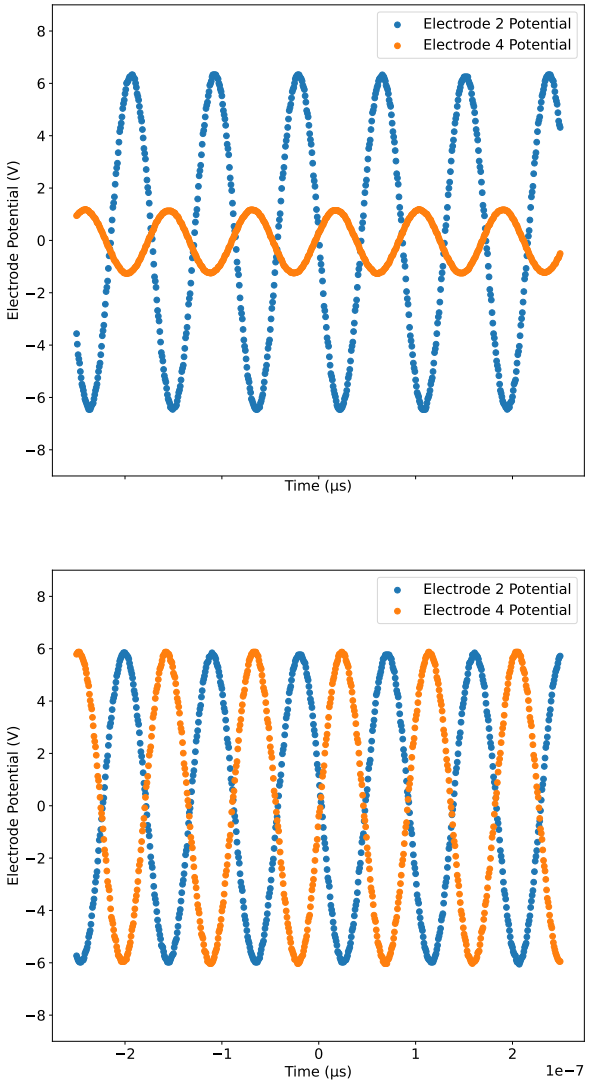

Figure 4.1: Potential vs. Time for solid electrode 4 and segmented electrode 2 with varying capacitance on electrode 4. All the other electrodes had 700 pF of capacitance in series, and the input signal had a peak-to-peak voltage of 2 V. (Top) 30 pF of capactiance were added in series to electrode 4, decreasing the voltage by a factor of about $\frac{3}{4}$. The circuit was driven at resonance with an RF frequency of 11.6 MHz. Note that the signals are not exactly 180° out of phase. (Bottom) 700 pF of capacitance were added to electrode 4, resulting in even amplitude signals. The circuit was driven at resonance with an RF frequency of 11.06 MHz.

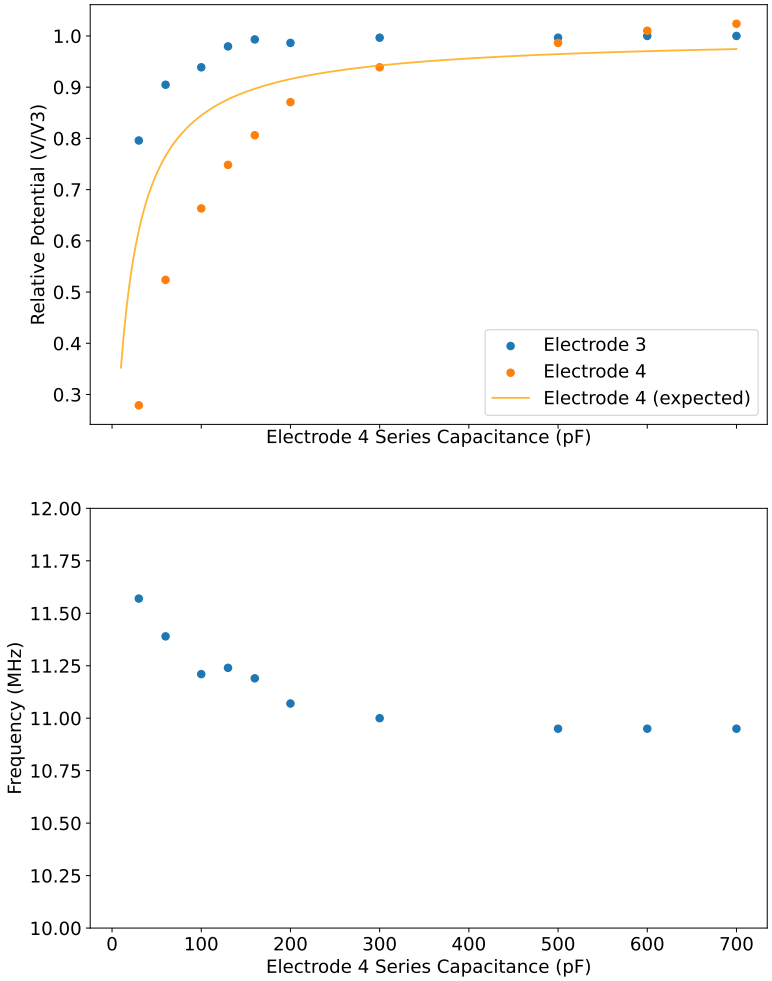

Figure 4.2: Electrode potential of the two solid electrodes with a varying capacitance in series on electrode 4. The three other electrodes each had 700 pF of capacitance in series, and the capacitance in series on electrode 4 was varied. (Top) Potential on electrode 3 and electrode 4 relative to the potential of electrode 3 at 700 pF as capacitance on electrode 4 is varied. (Bottom) Resonant frequency of the circuit as the capacitance is varied.

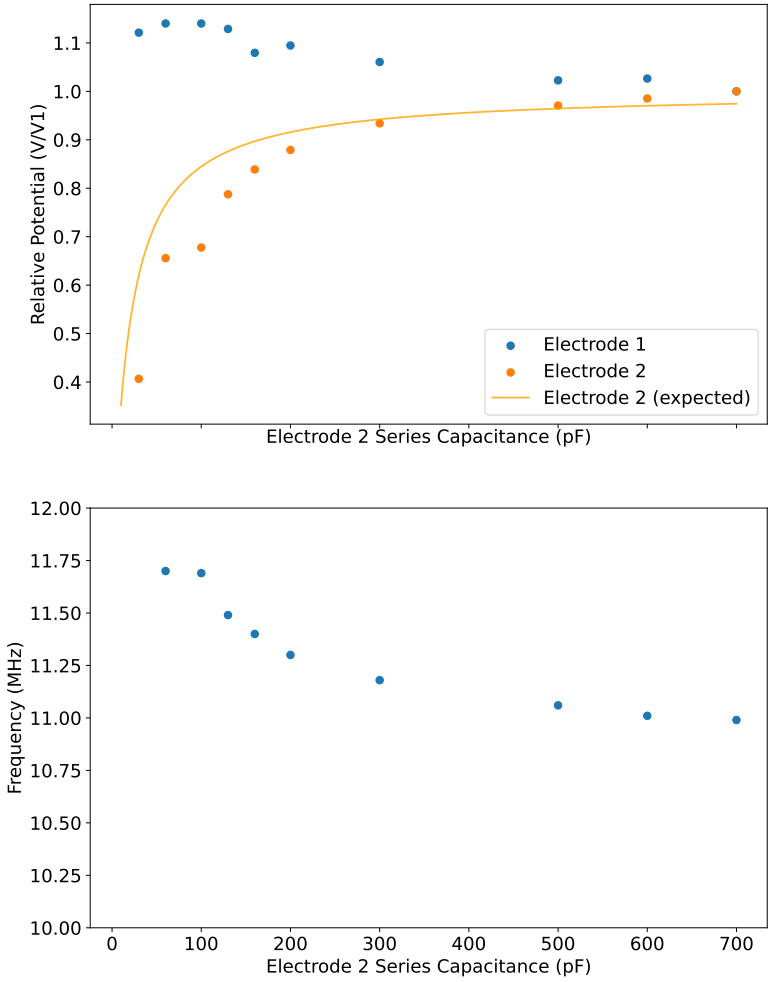

Figure 4.3: Electrode potential and resonant frequency with two segmented electrodes with a varying capacitance in series on electrode 2. The three other electrodes each had 700 pF of capacitance in series, and the capacitance in series on electrode 2 was varied. (Top) Potential on electrode 1 and electrode 2 relative to the potential of electrode 1 at 700 pF as capacitance on electrode 2 is varied. (Bottom) Resonant frequency of the circuit as the capacitance is varied.

due to the introduction of these variable capacitors would be advantageous: we want to keep the resonant frequency as high as possible (up to about 15 MHz), but various physical factors in the circuit add small amounts of capacitance, pushing down the resonance. The variable capacitors reduce the capacitance in the circuit, allowing for the resonance to be higher, as intended. However, by adding asymmetry to the two branches by adjusting these capacitors, the resonant frequencies of the two branches shift so that they are not exactly the same. There is also some asymmetry in the coil wrapping, causing differing resonance: although I was careful to evenly wind the coils on the two branches, without machining there is no feasible way to get perfect symmetry. This means that we need to be able to tune the capacitances of the branches of the circuit, likely by adding new variable capacitors in parallel.

From figs. 4.4 and 4.5, we can see that the resonances are very close, but not exactly equal. We can also see in the phase difference graphs that at lower frequencies, the phase difference is very close to 180° (as desired), and gets slightly worse through the resonance.[1] We can see an example of this at a specific frequency (11.6 MHz) in fig. 4.1 (top), where the signals have a slight phase shift. Because of this small difference, we only need a few tenths of a picofarad of tuning to be able to match the resonances. However, variable capacitors at these values are not produced, so we need to make our own by combining several capacitors in series; I did not leave space for this on the current board, since as stated above, we thought that this effect would be a benefit as opposed to a problem. In a future board, this will be important, and is discussed in section 5.1.1.

It is also possible to tune resonances with variable capacitors in series, as opposed to the ones in parallel. By tuning an 8-40 pF capacitor, we can slightly change the resonance of each half of the circuit. Since capacitances in series add reciprocally, when the variable capacitors are set to 500 pF, this only allows for only a few femtofarads of tuning, but when

[1] At higher frequencies, we would usually expect the difference to go back to 180°, but there is a second resonance at about 15-16 MHz changing this behavior. This resonance is not usable for our purposes because the signals are in phase instead of 180° out of phase.

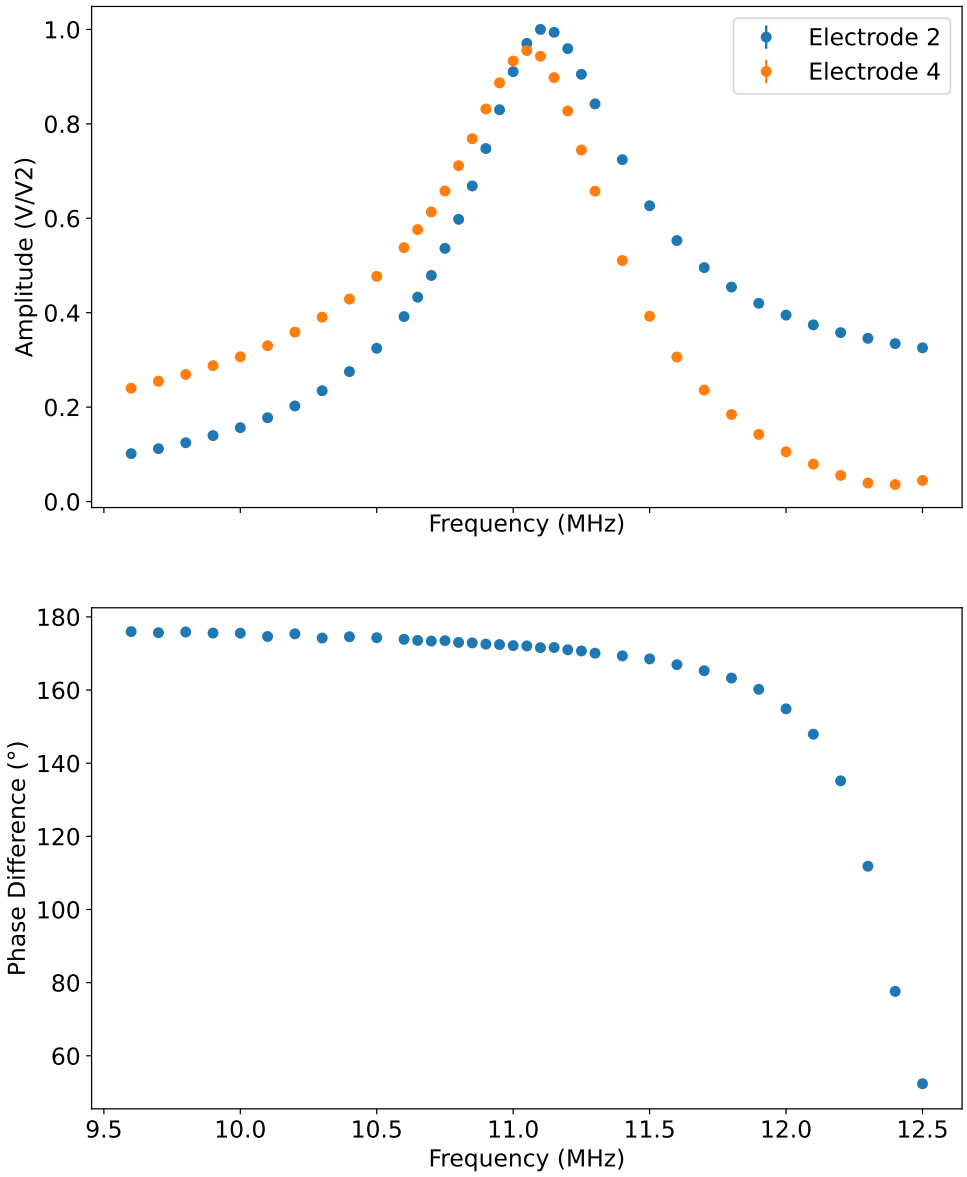

Figure 4.4: Amplitude vs. Frequency (top) and Phase Difference vs. Frequency (bottom) for the signal in electrodes 2 and 4. Electrode 2 is a segmented electrode, and electrode 4 is a solid electrode. The variable capacitors for all four electrodes were set to 700 pF. The voltages are expressed as ratios of the peak voltage at resonance for electrode 2. Note that error bars for these plots are included, but are small enough to be nearly invisible.

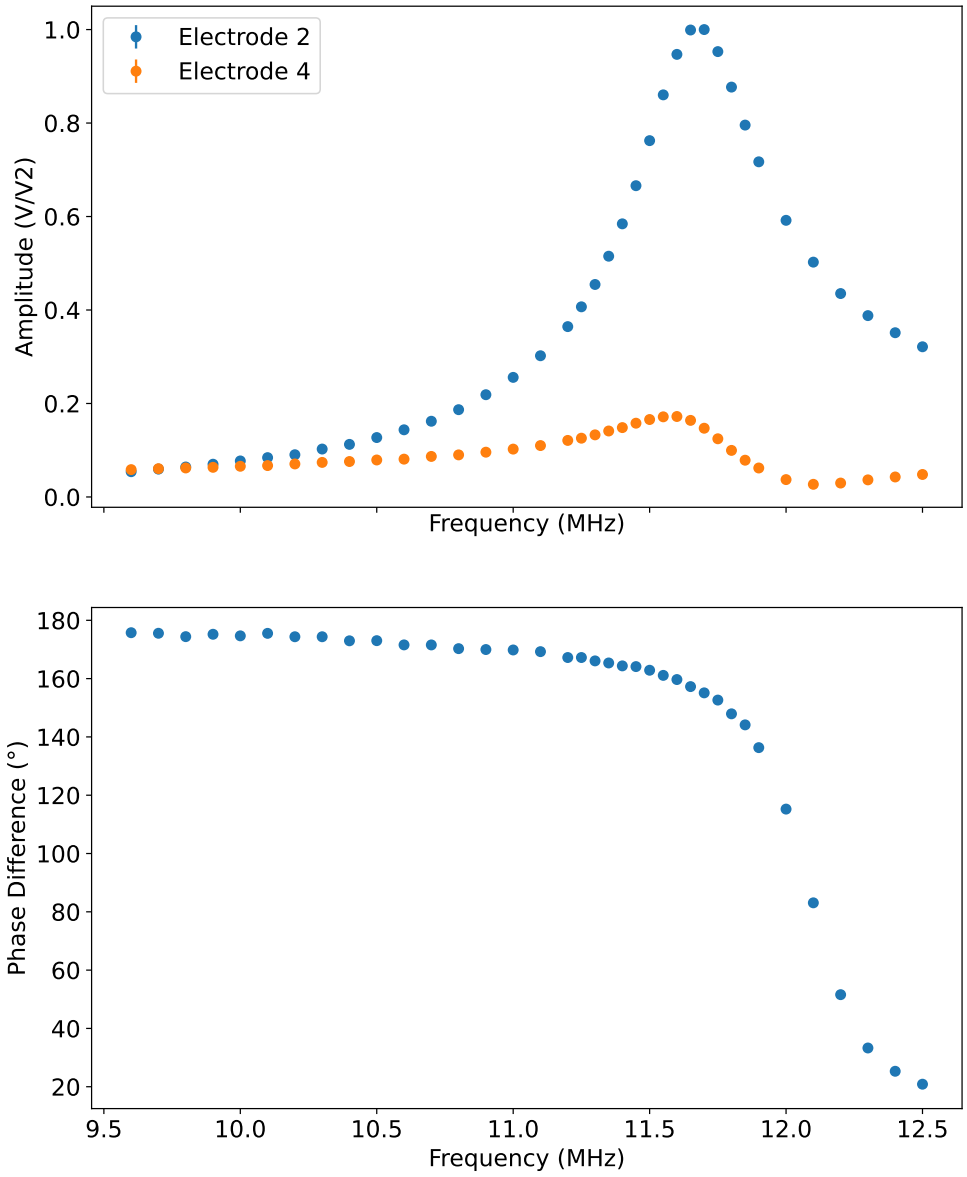

Figure 4.5: Amplitude vs. Frequency (top) and Phase Difference vs. Frequency (bottom) for the signal in electrodes 2 and 4. Electrode 2 is a segmented electrode, and electrode 4 is a solid electrode. The variable capacitors for electrodes 1, 2, and 3 were set to 700 pF, and electrode 4 to 40 pF. Voltages are expressed as ratios of the peak voltage at resonance for electrode 2. Note that error bars for these plots are included, but are small enough to be nearly invisible.

the variable capacitors are set to 100 pF, it allows for a few hundredths of pF. We found that this tuning was successful at 200 pF, meaning that about 0.01 pF is enough to bring the circuit branches to the same resonant frequency when the capacitance on all four of the electrode branches are the same. While this is useful for finding the difference in the resonances, this method is not a good long term solution because we want to be able to use the series variable capacitors for their intended purpose (moving the ions with RF pushing) instead of needing to tune them to maintain the desired phase shift.

This change in resonance is a problem because the non-180° phase difference in the signals causes micromotion in the ions in the trap; the position of the ions, accounting for the micromotion, can be written as

$$u_x(t) = [u_{0x} + u_{1x}\cos(\omega_{xt} + \phi_{Sx})][1 + \frac{1}{2}q_x\cos(\Omega t)] - \frac{1}{4}q_x R\alpha\phi_{\text{ac}}\sin(\Omega t), \tag{4.1}$$

where u_{0x} and u_{1x} are first- and second-order terms describing the secular motion of the ions, ω_{xt} and ϕ_{Sx} describe the secular frequency and phase of the ions, q_x and Ω are as described above, R is the radius of the trap, α is a geometric factor characteristic of the trap, and ϕ_{ac} is the difference in AC phases [34]. We focus on the micromotion term here, since secular behavior is discussed in section 3.1; the only factor we are missing is α. The relevant literature only defines α for a cylindrical Paul or Penning trap, where it has a value of 0.8; therefore, as a likely overestimate, we take α to be 1.0 [35]. By plugging in the relevant values, and taking the AC phase difference to be about 7° by using scope measurements, we expect to see a ion micromotion of about 10 microns.[2] In our images, though, we did not see this level of micromotion. The sensor on the camera is 1004 x 1002 pixels, and given our field of view of 736 microns, a 10 micron shift should be reflected in about 13 pixels, which would be noticeable to the eye, but we do not see this. This means that we are overestimating

[2]This value ranges from about 8-11 microns, depending on how we measure the phase. The scope is the easiest way to do this, but the measurements are not stable, so we do not have a very accurate value here. Figs. 4.4 and 4.5 use sinusoidal fits of scope data, but this method is too slow for sustained use. Since we eventually hope to reduce the phase significantly (see section 5.1.1), this should not cause long-term problems for the experiment.

either the phase difference or the α parameter. This result is promising, but because of the potential issues with the second-order Doppler shifts and the experimental uncertainty, we still want to do everything we can to equalize the branch resonances to ensure the smallest possible amount of micromotion.

4.1.4 AC Modulation

The AC modulation, which is used to measure the secular frequency of the ions in the trap, was also tested with the trap. Using equations 3.11, 3.15, and 3.16, we estimate the secular frequency to be about 1.0 MHz, and the axial frequency to be about 180 kHz; therefore, we expected to use AC frequencies in this range to measure the exact secular frequency of the ions and therefore measure the charge-to-mass ratio, confirming that we have oxygen ions. As described in section 3.4.3, we use a lowpass filter in order to filter out high-frequency noise on the AC modulation, and we also found that there was a second, unintended filter being formed. Although it is possible to quantitatively measure the properties of these filters, the exact values are not important:[3] we can test them easily by looking for resonances in the trapped ions because the ions only need a tiny signal to display visible effects. This process is described below in section 4.2.4.

4.2 Trapping Ions

Having measured the parameters of the trapping circuit setup, I will now describe the use of the circuit to trap ions.

[3]We tried to measure the attenuation at various frequencies using a spectrum analyzer, but the data was inconsistent with scope data at lower frequencies (where the scope was sensitive enough to collect data) and therefore unreliable, so it is not included.

4.2.1 Loading and Imaging

Although the molecular beam apparatus was not prepared for use at this point, Be^+ ions from the beryllium oven could still be trapped. The beryllium oven simply consists of a set of two beryllium coils. By running a current through these coils, they heat up, and beryllium atoms are ejected from the coils, forming a cloud inside the trap chamber. The electron gun then emits electrons with enough energy to remove a valence electron from beryllium atoms in the cloud, ionizing them so that they can be trapped [36]. The electron gun functions by heating up a metal with a relatively small work function, and by applying a voltage (70 V), which both lowers the work function further, and accelerates the electrons away from the gun.

In general, beryllium atoms have two pairs of electrons: one pair in the 1s shell, and one pair in the 2s shell. By ionizing the atoms, we remove an electron in the 2s shell, leaving an unpaired electron there. Generally, this 2s electron behaves a lot like a 2s electron in a hydrogen atom, but with one important difference: when the electron is close to the nucleus, it interacts with the 1s electrons. This means that if we excite the 2s electron to the 2p shell, we get an observable difference in its behavior relative to the nucleus. So, by exciting the electron to the 2p shell, and then allowing it to emit a photon to send itself back to the 2s state, we get an observable emission of photons. This is called a cycling transition, since the process happens repeatedly to the same ions. These photons allow us to use our imager to observe the position of the ions.

Our imager simply uses a lens to magnify the ions so that we can view them using an iXon3 camera. To calibrate the imager's field of view, we located an object in the imager field, moved the imager so it was exactly on the edge of the field, then moved the imager so that it was barely off the screen on the other edge. By measuring how far we moved the imager, we could determine the field of view. In our case, the field of view was 736 microns for most of the trapping, though we changed it to 655 microns for some of the data. The imager also contains a translational stage, allowing the lens tube to move in three directions:

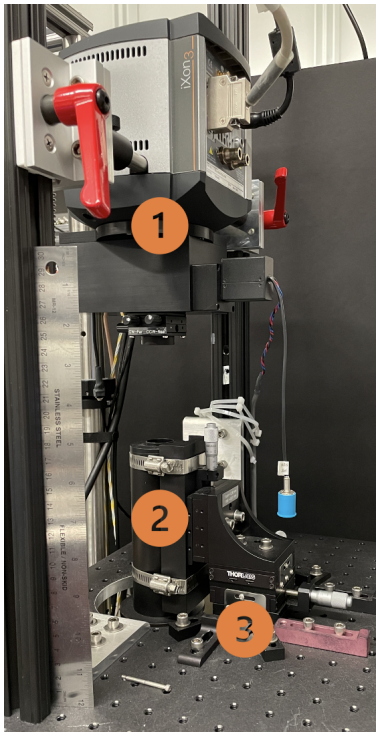

Figure 4.6: Photograph of the imager used to image trapped ions. (1) is the camera, which detects the photons and converts them to an image. (2) is the lens tube, which magnifies the ions, and (3) is the translational stage, which allow us to move the lens tube to find the location of the ions. The lens tube goes through the optical table to the top viewport of the trap chamber. Note the ruler, included for scale. This setup generates a field of view of 655 x 655 microns, though most of the data used a 736 x 736 micron field.

along the trap axis, across the gap between the electrodes, and vertically. We refer to these directions as along, across, and focus.

4.2.2 Trapping Parameters

Since the capacitance and therefore the resonance of the circuit were changed over the course of the data collection, the Be^+ ions were trapped using a range of RF frequencies from about 10.8 MHz to 12.2 MHz. One important aspect of the trapping is the DC voltages used to trap the ions. As described in section 3.4.3, the DC voltages have two main purposes: the first is to confine the ions along the trap axis (the RF does not do this), and the second is to account for asymmetries in the trap geometry that lead to micromotion and may prevent the ions from properly crystallizing. Since this was not an issue at the precisions we required,

we only considered the confinement properties of the DC offsets. In all of the measurements, we maintained a central electrode potential of 0 V, and set the outside electrodes to 10 V, and varied the second and fourth electrodes between 0.1 V and 5 V depending on our desired effects. These voltages are specified in the form 10/X/0/X/10 on each measurement, where each number represents the potential on an electrode.[4] As expected, with higher DC voltages, the ion crystals were compressed, although there was also some lateral movement, showing evidence of asymmetry in the trap. An example of trapped ions can be seen in fig. 4.7, although it should be noted that this trapping only used RF on the solid electrodes.

4.2.3 RF Pushing

The RF push, for which parameters are described in section 4.1.2, were tested for the extent to which they could move ions in the trap. To collect the data, the four electrode branches were set to equal capacitance (700 pF), and then tuned to resonance, which was about 10.90 MHz. All of the data was taken at DC potentials of 10/1/0/1/10. The imager was then used to bring the ions into the center of the field of view and into focus, and the laser was positioned to maximize the amount of light emitted by the ions and reaching the camera and to ensure the ions had crystallized. After recording the positions, the RF was stopped, the capacitance on one of the electrode branches was changed (electrode 2 for the segmented testing, and electrode 4 for the solid testing). The circuit was then tuned back to resonance, new ions were trapped, and then the process was repeated. The capacitances measured in each branch were 700 pF, 600 pF, 500 pF, 300 pF, 200 pF, 160 pF, 130 pF, 100 pF, 60 pF, and 30 pF.[5] Using this data, the distance between the ions was calculated in two different ways. The first is the difference in imager position to bring the ions into focus in the center of

[4]There is nothing special about the 10 V on the outside electrodes, we simply chose it because it was effective to trap the ions and gives us a wide range for the second and fourth electrode potentials.

[5]The variable capacitors were both tuned to about 30 pF by eye, so these values are estimates; the only way to accurately measure the capacitances is to remove the variable capacitors from the circuit altogether and measure using an impedance meter. As can be seen below, our understanding of the push depends on voltage through the branches, which can be directly measured, so it is not important to have accurate capacitance values.

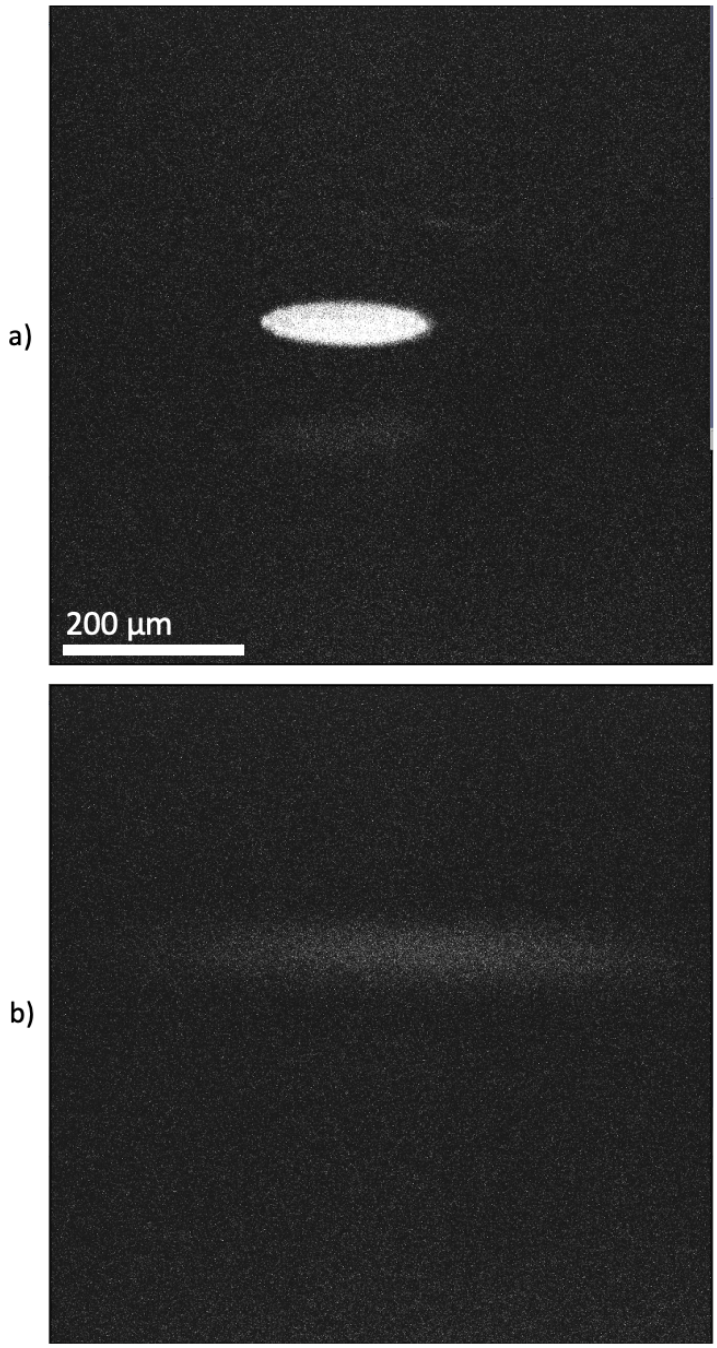

Figure 4.7: Photograph of trapped ions at 12.8 MHz, 108 Vpp on the solid electrodes, and no RF on the segmented electrodes. Each of the images has length and width of 736 microns. a) DC offsets set to 10/2/0/2/10. b) DC offsets set to 10/0.1/0/0.1/10. We see that varying the DC offsets on the segmented electrodes allows us to induce crystallization of the ion cloud. We also noticed that there was a lateral shift in ion position when the DC offsets were changed; this is likely due to geometrical asymmetries in the trap assembly.

the field of view, and the second is the difference in laser position to scatter the most light off of the ions. The results of these measurements can be seen in figs. 4.8 and 4.9. In addition to the measured data, the plots contain simulated data and two theory curves. The simulated data was generated using SIMION with the exact RF voltages as the measured points. The theory curves are based on the derivation in ref. [33], discussed in detail in section 3.3, which shows that by reducing the voltage on one electrode relative to another in-phase electrode, we can push the ions between those two electrodes. The expected distance curve is based on the full derivation, which, maintaining the convention that $U_B = (1 - \delta)U_A$ so that the potential on electrode A is greater than that on B, gives us the relation that

$$\delta x = \frac{\delta}{2 - \delta} \frac{r_0}{2}. \tag{4.2}$$

In our case, r_0 is 1.25 mm. Additionally, we can consider the case for which the potentials on electrodes A and B are similar, which gives $\delta x = \frac{\delta r_0}{4}$. Although this should only fit the data near a voltage ratio of 1.0, we see that it actually fits the data (especially the simulated data) better than the full theory curve. Although it is not obvious why this may occur, one possibility is that one of the assumptions in the derivation, that the two sets of opposing electrodes behave independently, is not safe to make, which would mean that we will need to investigate this behavior further.

It is important to note that the data was collected using imager translation. The dimensional shifts from the imager's translational stage are at 45° from the ion pushing axes because of the orientation of the trap electrodes. We only include the across and focus directions in the distance measurement, since there is no motivation for changes in position along the trap axis. We did observe some of this, but it was not correlated with any type of push. This could be due to the fact that the focus knob on the imager was not as precise as we would have liked. Not only is it relatively arbitrary to tell when the ions are fully in focus, but when the focus knob is turned, the whole imager shifts, including in the other two

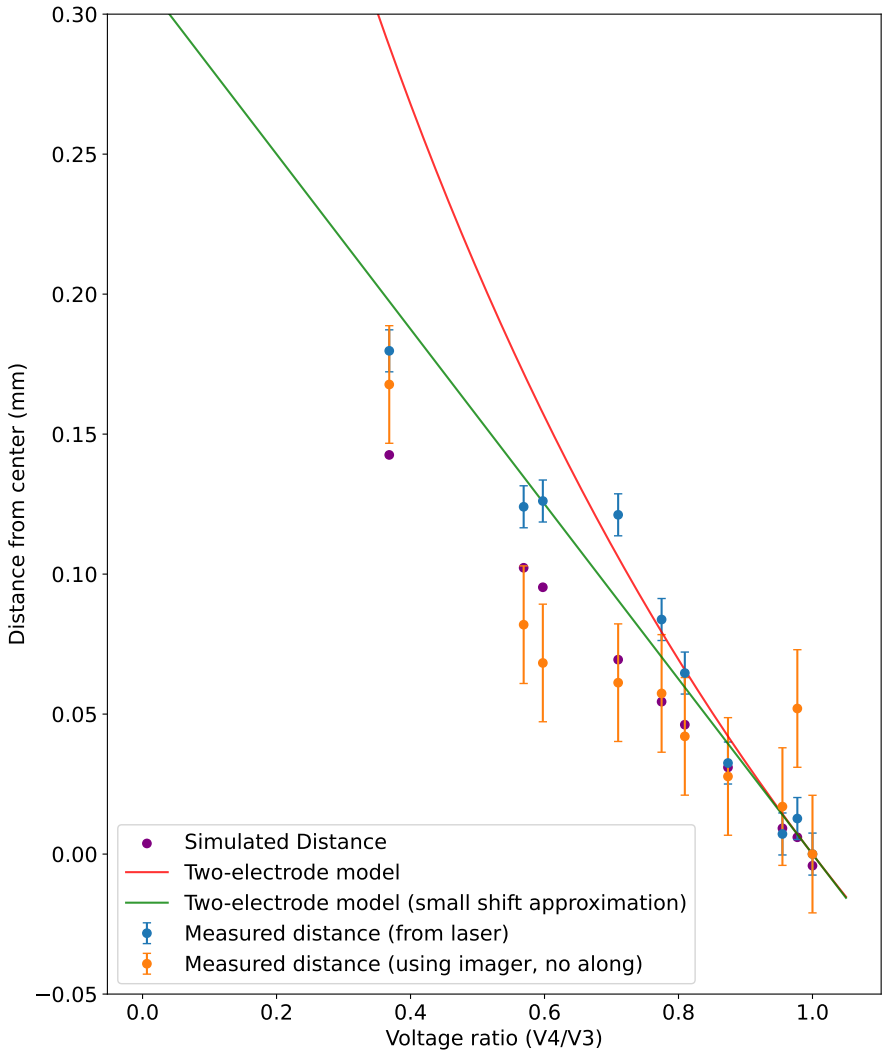

Figure 4.8: Ion position as a function of RF amplitude ratio on the segmented electrodes.

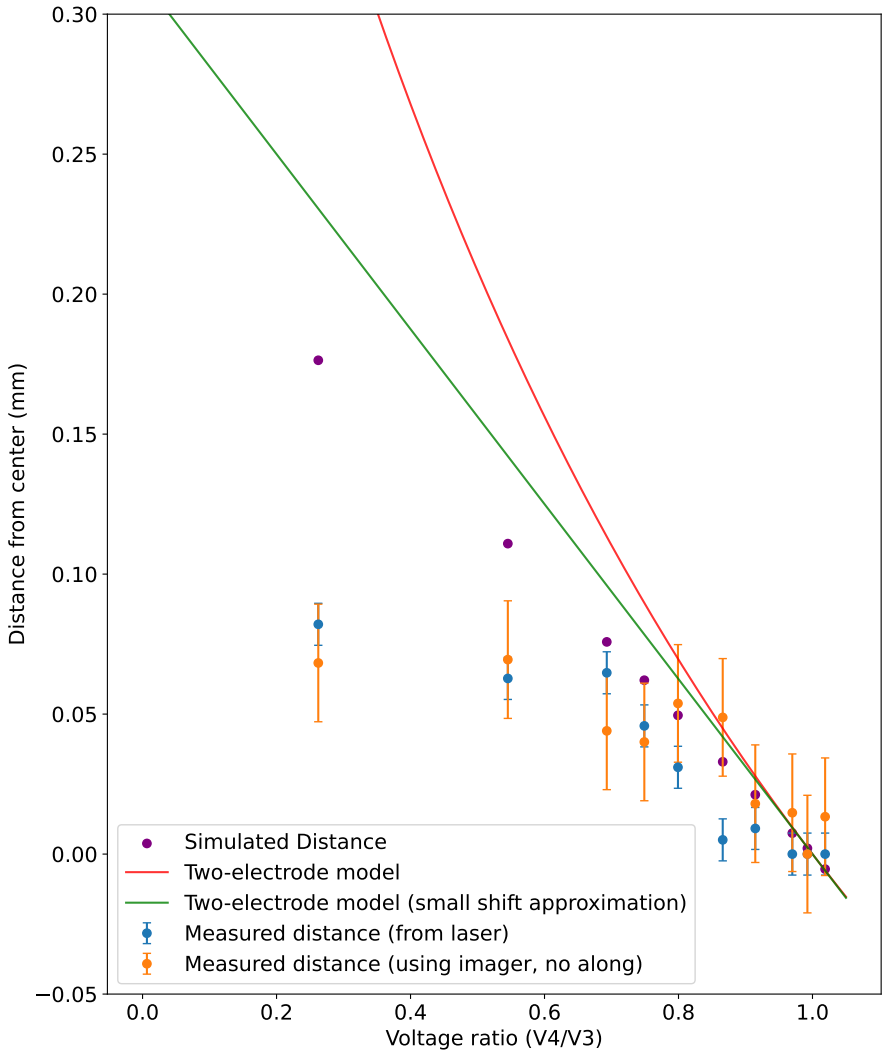

Figure 4.9: Ion position as a function of ratio of RF amplitude on the solid electrodes.

directions. Although it usually shifts back when the knob is released, this is not always true, which means that the other two directions may be affected. From my use of the imager, I estimate the error in the focus to be about 20 microns, and the error on the along measurement to be about 5 microns, giving a net error of about ± 20.1 microns. Since the laser was sent through the trap along the trap axis, there is no along measurement for the laser data, so that concern does not apply; however, like the focus, it is difficult to tell exactly when the most light is scattered off of the ions. The laser is a lot more accurate, though; I estimate the error in the vertical direction to be about 2.5 microns, and in the horizontal direction to be about 10 microns, giving a net error of ± 10.3 microns. These error bars are included on the plots in figs. 4.8 and 4.9.[6]

From the data, we see that we can move the ions by about 180 microns towards either of the segmented electrodes, and about 80 microns towards the solid electrodes. It is possible that with careful tuning of the variable capacitors to under 30 pF, a bigger shift could be generated; however, at the limits tested here, the trap started to become unstable, with ions slowly leaving the trap over the course of the measurement. Therefore, these limits are probably good ones for our ability to push the ions using RF. Since we were able to see clear crystallization of the ions even at the furthest bounds, micromotion was likely not a significant issue despite the problems arising in 4.1.3. It is not clear why the segmented electrodes were able to push so much more than the solid electrodes; it is likely that asymmetries in the geometry of the trap or small inefficiencies in the circuit are responsible, although these may be hard to identify and/or fix, if they exist. Still, the amount of movement we observed allows for sufficient control over the ions in the trap to be able to move the crystal into the required location for intersection with the laser and molecular beam paths needed for the experiment.

[6]It would have been preferred to gather enough data to get statistical uncertainties for the RF push, but due to time constraints, this was not possible.

4.2.4 AC Modulation

Finally, in addition to the RF push, the AC modulation functionality of the circuit was tested, allowing us to measure the geometry of the trap potential and make an accurate measurement of the RF amplitude.[7] To make this measurement, we loaded ions at 11.1 MHz RF with all four electrodes at 91 V_{pp} and the DC offsets at 10/2/0/2/10. The modulation blurs the ions, so they do not appear to be crystallized when these measurements are made.

Fig. 4.10 shows the results of adding an AC modulation to the second segment on electrode 2, with an amplitude of 10 V_{pp}. Since this segment is off-center, finding the resonance here acts as a measurement of the axial frequency of the trap ions. We found that the measured axial frequency depended on the amplitude of the modulation signal: using an 8 V_{pp} amplitude, we found that $\omega_z = 2\pi * 243.5$ kHz; with a 9 V_{pp} amplitude, we found that $\omega_z = 2\pi * 236$ kHz; and with a 10 V_{pp} amplitude, we found that $\omega_z = 2\pi * 223$ kHz. Referring back to section 3.1, recall that the equation for axial frequency is

$$\omega_z = \sqrt{\frac{2\kappa Z|e|V}{mz_0^2}}. \tag{4.3}$$

Since we now have a measured value for ω_z, and the rest of the values in this equation are constants for the beryllium ions and the trap conditions, we can calculate κ. However, nothing in the equation suggests that we should expect different measurements of axial frequency at different modulation amplitudes. The most likely explanation for this effect is that our approximation of the axial oscillations as harmonic is not fully accurate. This would mean that the smallest amplitude modulation is the most accurate, since it would drive the oscillations the least into the non-harmonic regime. Still, it is worth investigating this behavior further. Using these measured values, we get $\kappa = 0.124 \pm 0.001$ for the 8 V signal, $\kappa = 0.116 \pm 0.001$ for the 9 V signal, and $\kappa = 0.103 \pm 0.001$ for the 10 V signal.[8]

[7] In general, the equation that we use to relate the AC modulation to the RF amplitude can be used to measure the charge-to-mass ratio of the ions given an RF amplitude, but in this case, we already have the charge-to-mass ratio since we are trapping beryllium ions.

[8] I estimate that ω_z is accurate to about 0.5 kHz and V is accurate to about 1%.

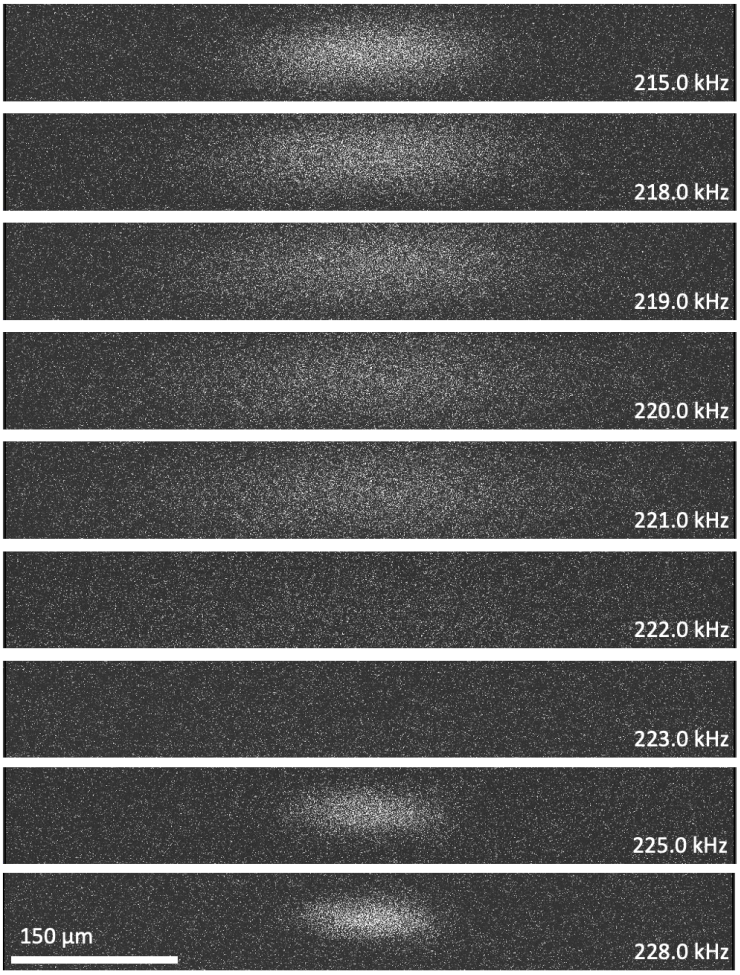

Figure 4.10: Measurement of the axial frequency of the ion motion using AC modulation on electrode 2, segment 2. The width of each of the images is 655 microns. The AC modulation was applied as a 10 V_{PP} signal. Resonance appeared at 223 kHz.

These values are all slightly below our simulated value of 0.17, but are reasonable for our trap conditions.

Fig. 4.11 shows the results of an AC modulation on electrode 3, a solid electrode. This segment has its center at the midpoint of the trap, so finding the resonance here acts as a measurement of the radial frequency. This time, we used a modulation voltage of 3 V_{pp}, and swept the signal frequencies, finding a resonance at 966 kHz. Referring again to section 3.1, we can use equations. 3.11, 3.13, and 3.16 to estimate a value for the RF amplitude; we

get a value of 88 ± 18 V_{pp},[9] which is in line with the scope-measured value of 91 V_{pp}.[10] Alternatively, we can assume the voltage is correct, and calculate the trap radius; doing this, we get a radius of 1.27 ± 0.06 mm.[11] From these measurements, we can conclude that our assumed trap dimensions are correct, and that we understand the geometry of the trap reasonably well. This measurement was repeated with the third segment on electrode 2, and gave the same result. This is as expected, and means that we can measure the motional frequencies of the ions in all three spatial dimensions.

In sum, the circuit was tested to characterize its behavior when applying DC offsets, varying RF amplitudes, and AC modulation. The DC offsets worked fully as expected; the RF pushing allowed us to move the ions up to about 180 microns towards each of the segmented electrodes and about 90 microns towards each of the solid electrodes. The AC modulation allowed us to measure the axial and radial frequencies of the beryllium ions in the trap. However, the design of the circuit, which uses separately wound coils for the segmented and solid electrodes and differing variable capacitors on each branch, means that the resonant frequency of the two sections of the circuit is not exactly even. This means that the circuit will need to be modified in the future to allow for fine tuning of the resonance of each branch to ensure that the resonances match.

[9] I assume an uncertainty of 1 kHz in the AC modulation measurement, 0.1 kHz in the RF input signal, and 10% (0.125 mm) in the trap radius.

[10] Recall from section 4.1.1 that the scope voltage amplitude measurements are relatively inaccurate.

[11] I assume a 10% uncertainty in the scope voltage measurement.

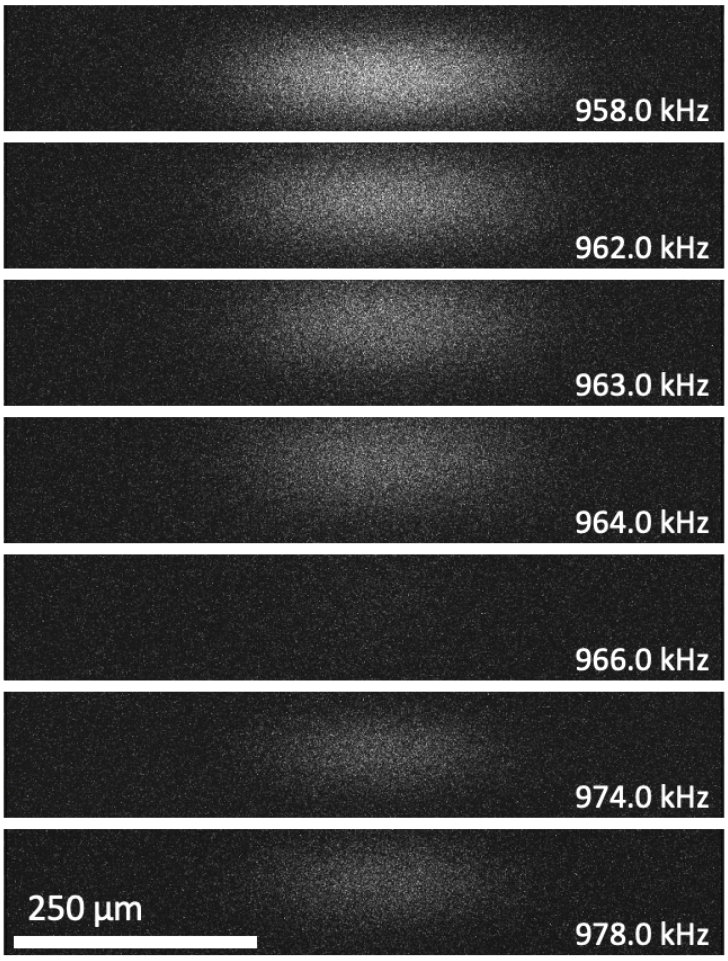

Figure 4.11: Measurement of the radial frequency of the ion motion using AC modulation on electrode 3. The width of each of the images is about 736 microns. The AC modulation was applied as a 3 V_{PP} signal. Resonance appeared at 966 kHz.

Chapter 5

Conclusion

5.1 Circuit Improvements and Next Steps

Based on these results, we can see that the current circuit setup works effectively to meet the constraints required for the ion trap to work. The RF trapping, DC offsets, AC modulation, and RF positional tuning all are in working order and provide a strong basis for the circuit. However, there are still improvements to be made before the circuit is ready for the purposes of the measurement of the proton-to-electron mass ratio.

5.1.1 Solving the Phase Issues

As stated in section 4.1.3, the current setup has the segmented and solid electrodes operating at phases that are not exactly 180° apart. This is an issue because the micromotion that this phase difference induces will lead to an increased time dilation effect on the laser frequency, which is important for the frequency measurement of the vibrational transition (see section 3.1). In order to fix this, we can simply add a (very small) variable capacitor in parallel directly after the series variable capacitor in each of the four branches of the circuit. This will allow for tuning of the resonance to the precision we need to prevent micromotion. There are some small logistical issues with this: currently, it is very difficult to access the

circuit when it is plugged into the trap (see fig. 3.10), since there are only 8.8 cm of clearance underneath it. Therefore, in order to make sure we can actually tune the variable capacitors while the circuit is plugged into the trap, it may make sense to introduce several daughter boards that plug into the main board perpendicularly. Regardless of the logistics of the solution, though, this is a small but important issue that must be addressed.

5.1.2 Spatial Constraints

During the design of the circuit, I made several assumptions about how various components would fit in the larger vacuum chamber setup that were later found to be inaccurate. Currently, the 25-pin DSUB connector between the auxiliary board is on the front of the main circuit board, which I thought would be the best place to access it. However, shielded 25-pin DSUB cables are much bigger than I expected, which means that there is not enough room between the board and the table for the cable to come straight downwards; this means we needed to use a 90° connector for access. This is a problem because it causes the cable to block access to two of the circuit branches, meaning that it is very difficult to change the variable capacitors on those branches. On a future version of the board, it makes more sense to put this connector on the edge so that the cable comes straight outward and doesn't interfere with other parts of the circuit.

Additionally, I also thought that the circuit would be rotated 180° from its current orientation; however, one of the adapters flipped the pin locations from where I thought they would be. This meant that we were spatially limited on the side of the board containing the SMA connectors used for the built-in scope probes on the circuit, so we had to remove two of them to be able to use the circuit. In the future, these will need to be relocated. It is also worth noting that we did not use the circuit's built-in scope probes for data collection because their calibration relies on a fixed capacitance in the circuit. Much of our data collection relies on changing the capacitance, so it was easier to use independent scope probes. In later iterations of the circuit, once the position of the ions is determined, we plan to monitor

the circuit with the built-in scope probes.

We also plan to put the final version of the circuit inside of metal box to minimize noise in the RF signals. Construction on this box has already begun; however, due to these unexpected constraints and limitations described above, we decided to pause until we have a fixed plan of what the final version of the circuit will look like. Once we do have a more concrete vision of the final product, it will be relatively easy to create an updated box.

5.1.3 Miscellaneous Issues and Possibilities

In addition to physical fixes, there is some circuit analysis that we did not have time to complete. Most significantly, we should collect more comprehensive pushing data while varying all of the electrode voltages. As argued in section 4.2.3, it seems that our derivation, which assumed independent behavior between the two sets of opposing electrodes, was likely an oversimplification. Therefore, it is important that we collect more data in order to understand and quantify the effects of shifting the RF potentials on all four of the electrodes on the ion position. Our current model is a good approximation that works as an order-of-magnitude prediction for ion location, but given the need for precision when positioning the ions at the intersection of so many beam lines, we need to have a very quantitative understanding of how the positioning works. Based on a conceptual picture of the saddle potentials in the trap, we would expect the presence of the second set of RF electrodes to bias the ions towards the center of the trap compared to our model, which is what we see, but it would be valuable to explain this more precisely.

It is also worth spending more time investigating the axial frequency measurements. In our observed data, we found that the measured frequency depended on the amplitude of the modulation, which does not agree with our theoretical understanding of the motion. In the future, the behavior should be investigated further at a wider range of amplitudes to see if the measurements stabilize or continue their behavior, and which measurements produce an accurate value for κ.

There are also a few minor issues with the circuit, which are less important to address but still should be considered. Notably, the built-in scope probe for electrode 1 has an issue with capacitive coupling with one of the secondary outputs from the transformer coil. These two signals each have traces on the circuit board that run parallel to each other on the opposite side of the board for 8.5 cm. This means that the secondary output, which has much higher voltage, leaks over to the scope probe output, making the data unreliable. We can use independent scope probes at a test point to collect this data for now, but this is worth fixing in the future.

One peculiarity that is not necessarily a concern is the presence of a DC voltage on the coil side of each of the blocking capacitors on the circuit branches. This potential is consistently measured at about 2.75 V, but has no obvious source. This has not had any effects on the DC voltages elsewhere in the circuit, so there is no effect on the trapping, but it is still worth investigating at some point, since it is possible that other changes to the circuit could cause this to become an issue.

In sum, the circuit as currently constructed works very well, but there are a series of minor improvements, some more important than others, that should be made before it is fully usable for the experiment.

5.2 Closing Thoughts

As I complete my year in working in the Hanneke lab, there remains much to do to build towards the point where we can make a meaningful measurement of the electron-to-proton mass ratio. Some of the next steps, such as using the electronics system to trap oxygen ions in addition to beryllium, follow directly from the work I completed. Others, such as the design and implementation of the spectroscopy laser, will require significant thought, testing, and design. In particular, we had hoped as part of my work this year to implement the optical cavity for the spectroscopy inside of the vacuum chamber; however, some delays

on parts meant that this had to be pushed back into the summer, which means that the vacuum will need to be vented and pumped down yet again before work can be continued. We are also waiting on factors outside of our control for the timeline of the experiment. As mentioned in section 1.5, a measurement using an atomic clock requires a reference clock: in our case, over the course of working up to the final measurement, we will use a few different references. The first will be a neon lamp used in a High Finesse WS-7 wavemeter; then, a cesium microwave clock will be used. Finally, for the experiment itself, an atomic optical clock will be used. It is anticipated that a clock of this form will be available commercially for purchase in the next few years, which is the same timescale it will take to prepare for the measurement.

While the continuations to my work and to the experiment in general will certainly be challenging, I am very excited to see where the experiment goes in the future. We are currently at the very beginning of an ambitious and significant project, and though we have made considerable progress, there is much still to do. I am interested to see how Professor Hanneke and future students modify and improve on my electronics, and I hope that in the next few years, the experiment will be able to produce a meaningful measurement for the variation of μ.

More generally, Standard Model physics is currently at a fascinating crossroads. Over the past 50 years, particle accelerators have been the primary mechanism used to discover new particles; however, now that much of the natural WIMP parameter space has been ruled out, it seems that AMO experiments are positioned to find evidence of non-Standard Model particles, including dark matter. As we rule out more and more of the parameter space for dark matter, it seems likely that we will eventually find evidence of a particle that explains the longstanding cosmological observations of the phenomenon. Precision measurement experiments, such as this one, are the primary mechanism we can use to find this evidence, and hopefully soon in the future, perhaps even in this experiment, we will see some kind of dark matter.

Appendix A

Circuit Diagrams

I include circuit diagrams and PCB designs below for each of the final circuits I used for the RF ion trap.

A.1 Radiofrequency Board

This is the circuit diagram for the main board for ion trapping, containing all of the RF electronics. The details of the board are described in section 3.4.

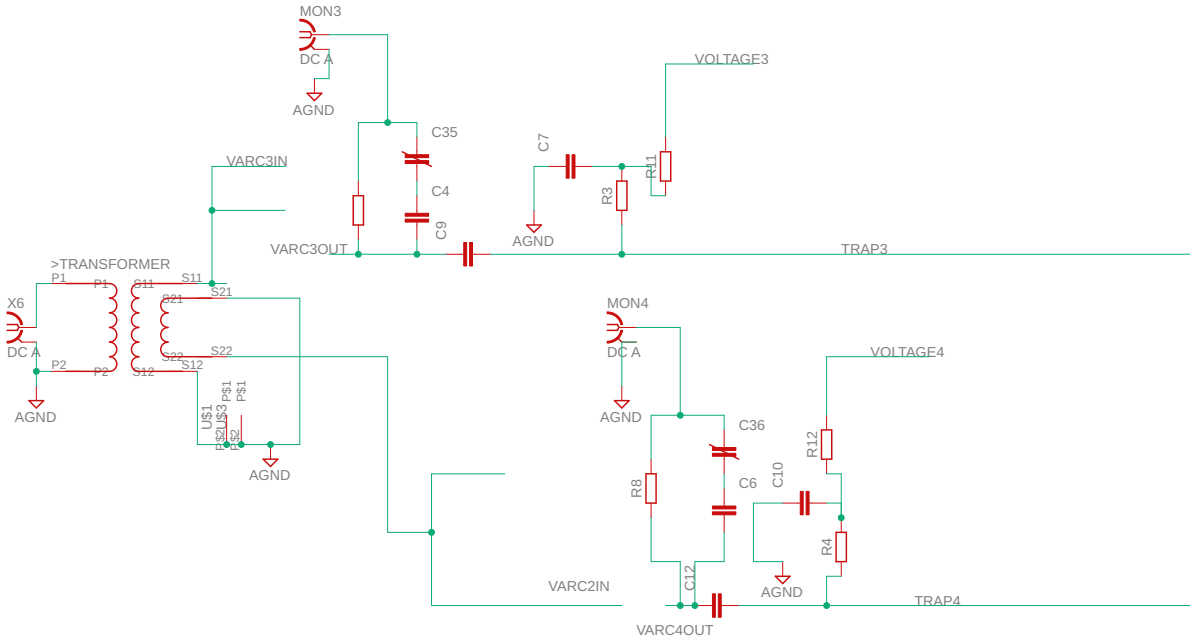

Figure A.1: Circuit diagram containing the transformer coil and the electronics on the solid electrodes (excluding the variable capacitors for RF pushing).

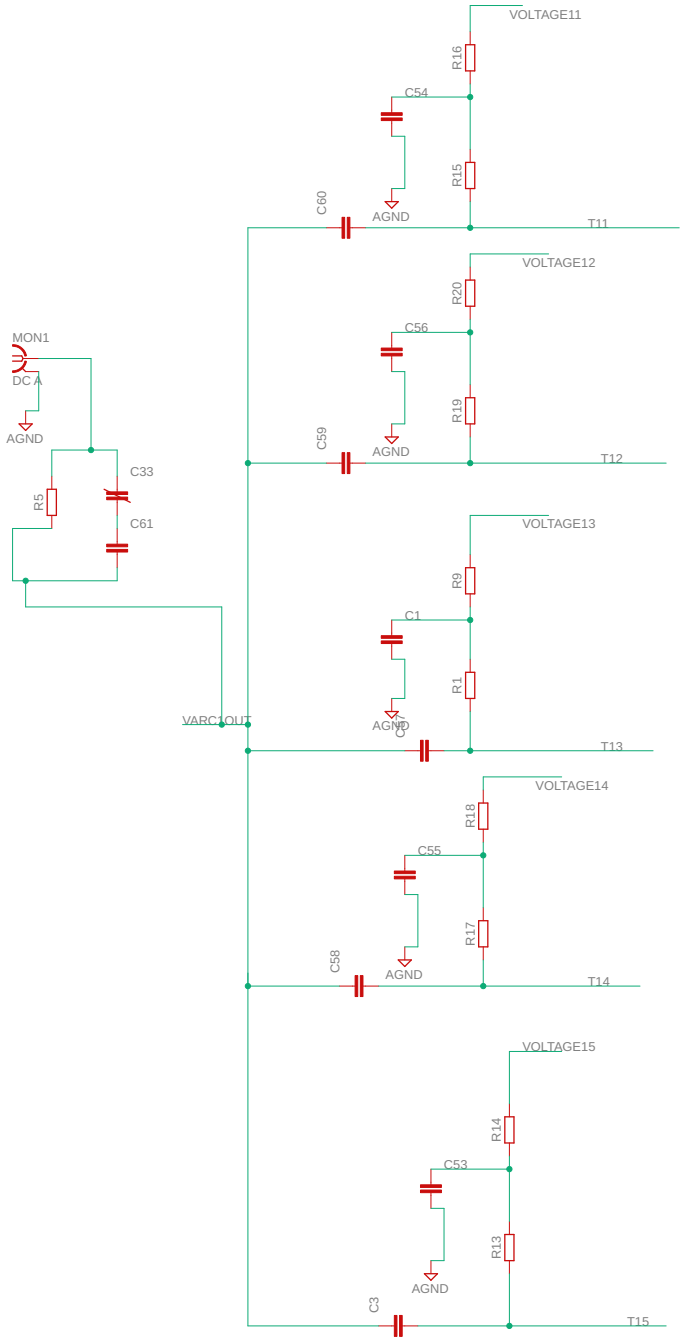

Figure A.2: Circuit diagram containing the electronics on segmented electrode 1 (excluding the variable capacitors for RF pushing).

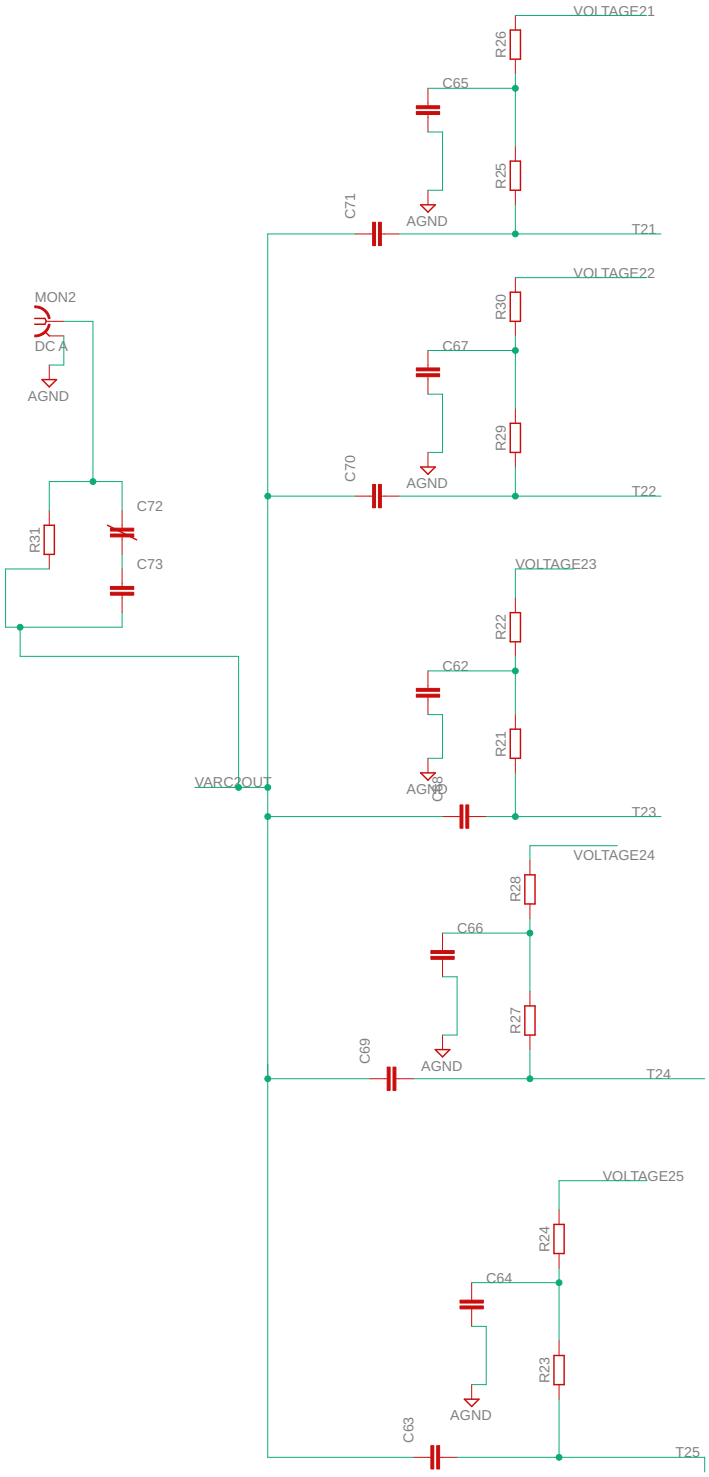

Figure A.3: Circuit diagram containing the electronics on segmented electrode 1 (excluding the variable capacitors for RF pushing).

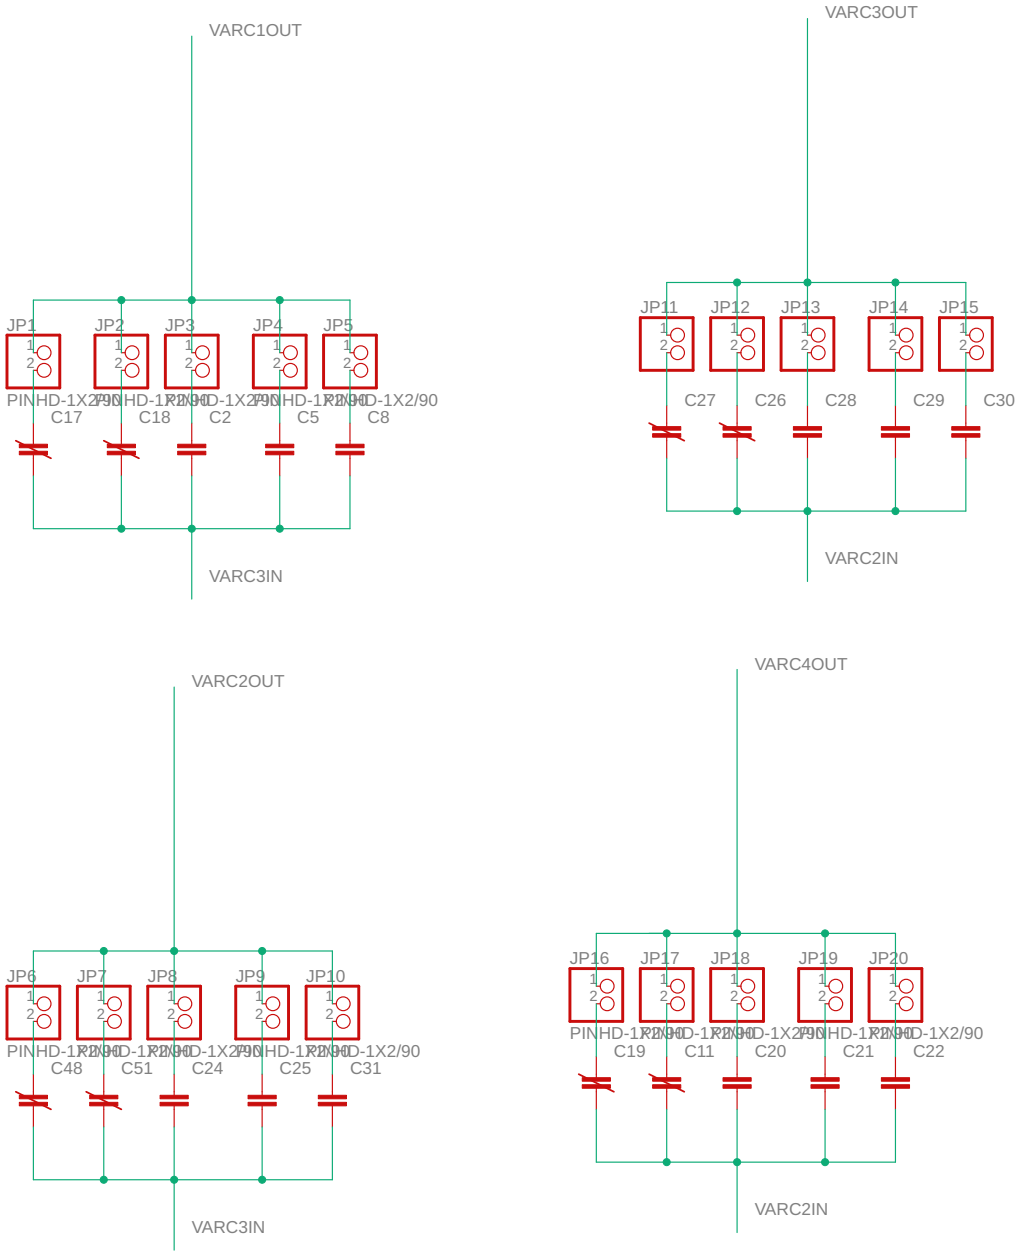

Figure A.4: Circuit diagram containing the schematic of the variable capacitors on each of the four electrode branches for RF pushing.

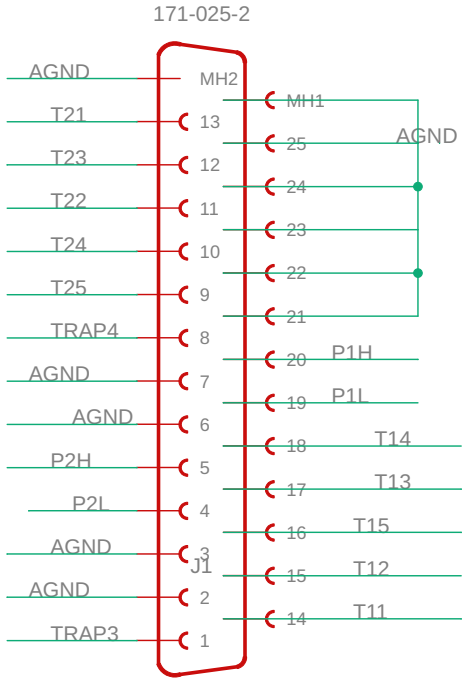

Figure A.5: Schematic diagram for the 25-pin DSUB adapter connecting the RF trap circuit board to the trap chamber.

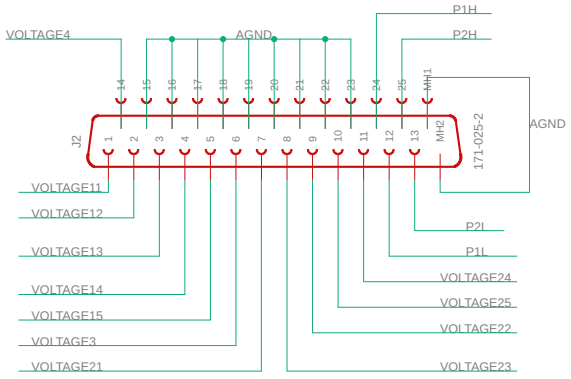

Figure A.6: Schematic diagram for the 25-pin DSUB adapter connecting the auxiliary board to the main RF board.

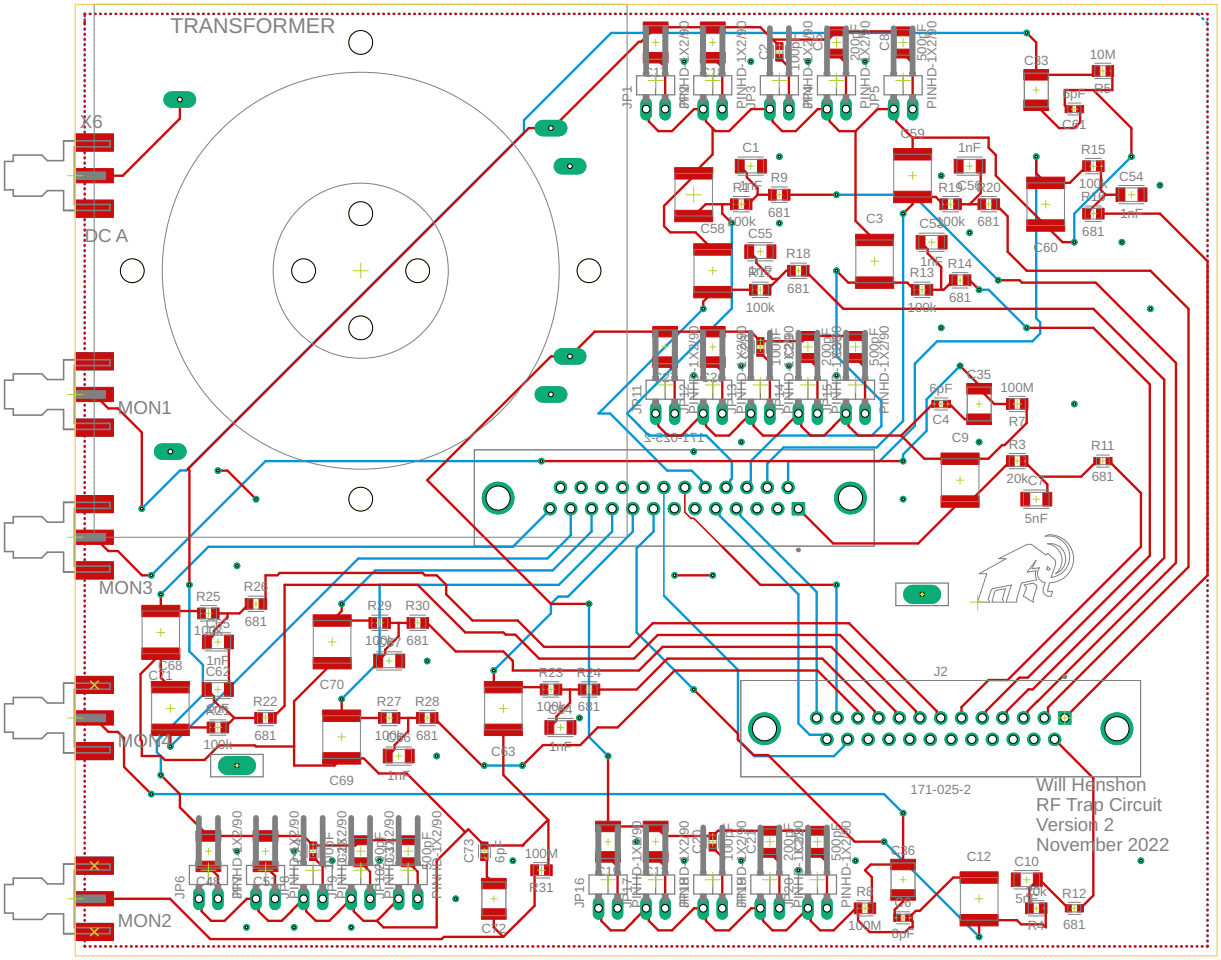

Figure A.7: EAGLE PCB design for the main RF board. Board dimensions are 15.24 cm x 12.7 cm.

A.2 Auxiliary Board

This is the circuit design for the auxiliary board, which supplies the main board with the necessary DC offsets at each electrode segment and the AC modulation.

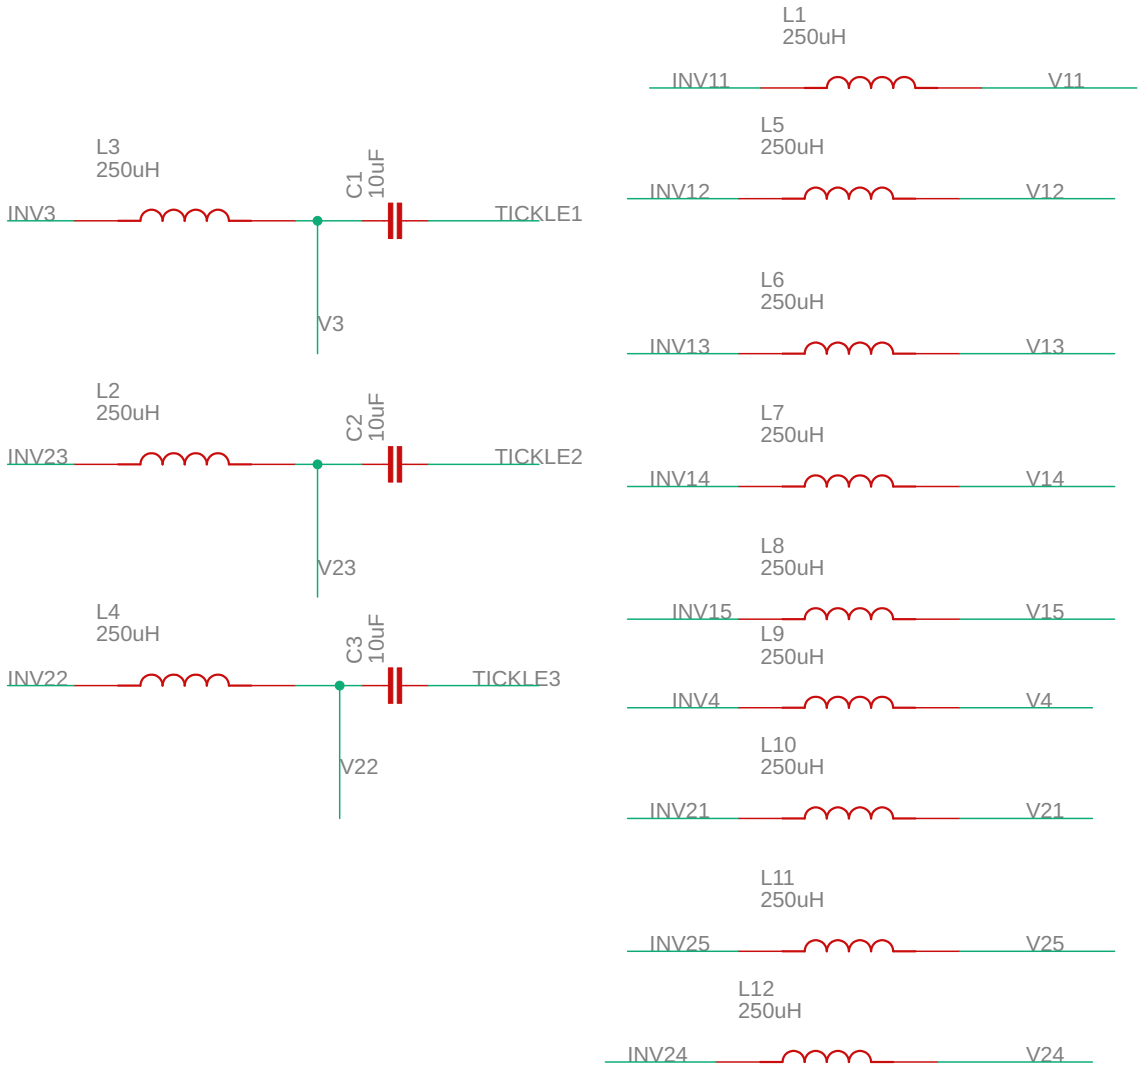

Figure A.8: Circuit schematic for the AC and DC filters on the auxiliary circuit board.

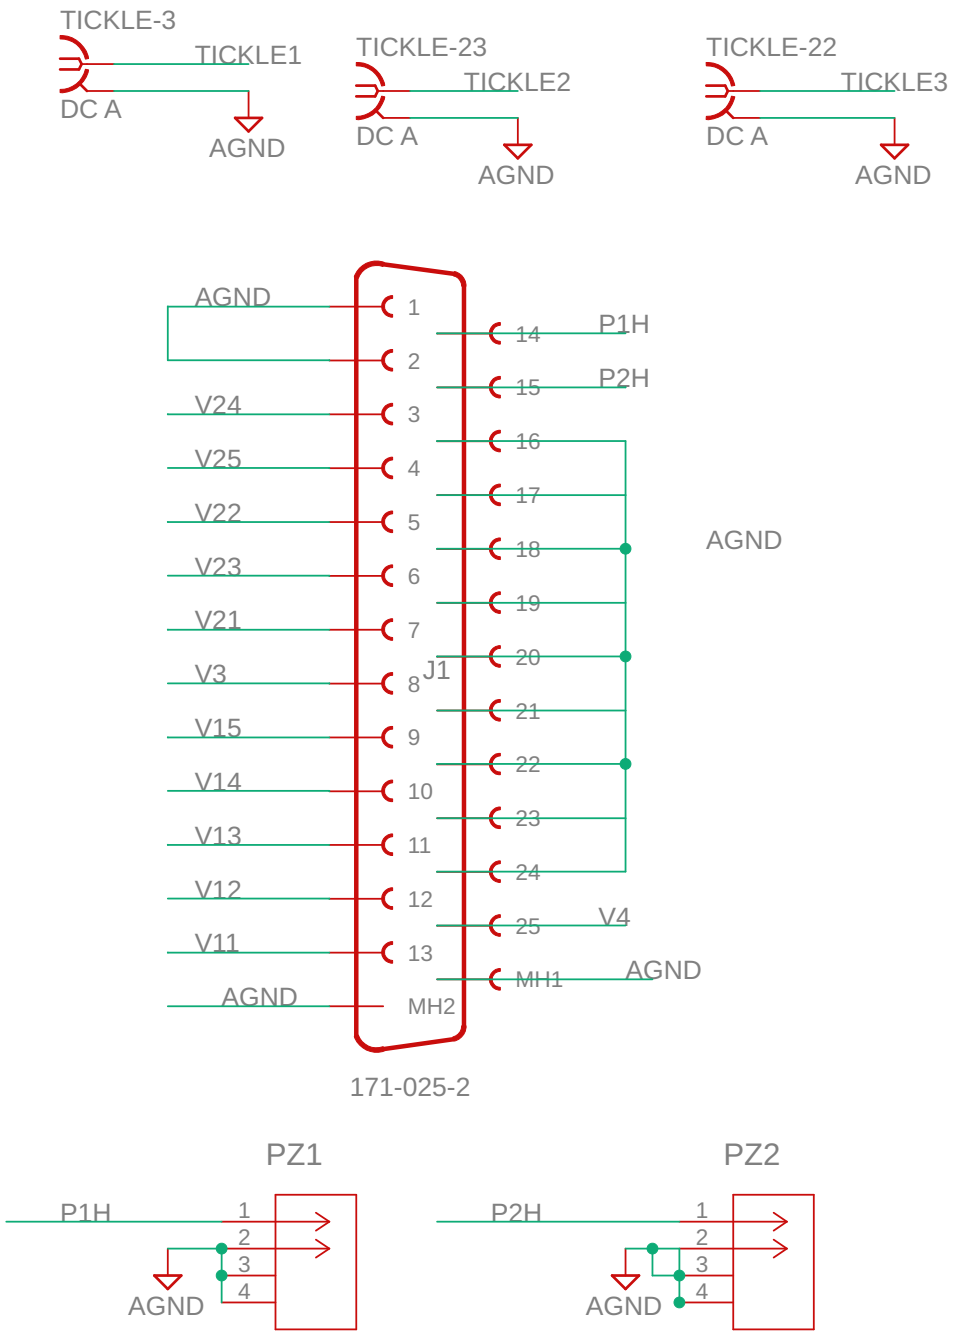

Figure A.9: Circuit schematic for the output adapter and connectors on the auxiliary circuit board.

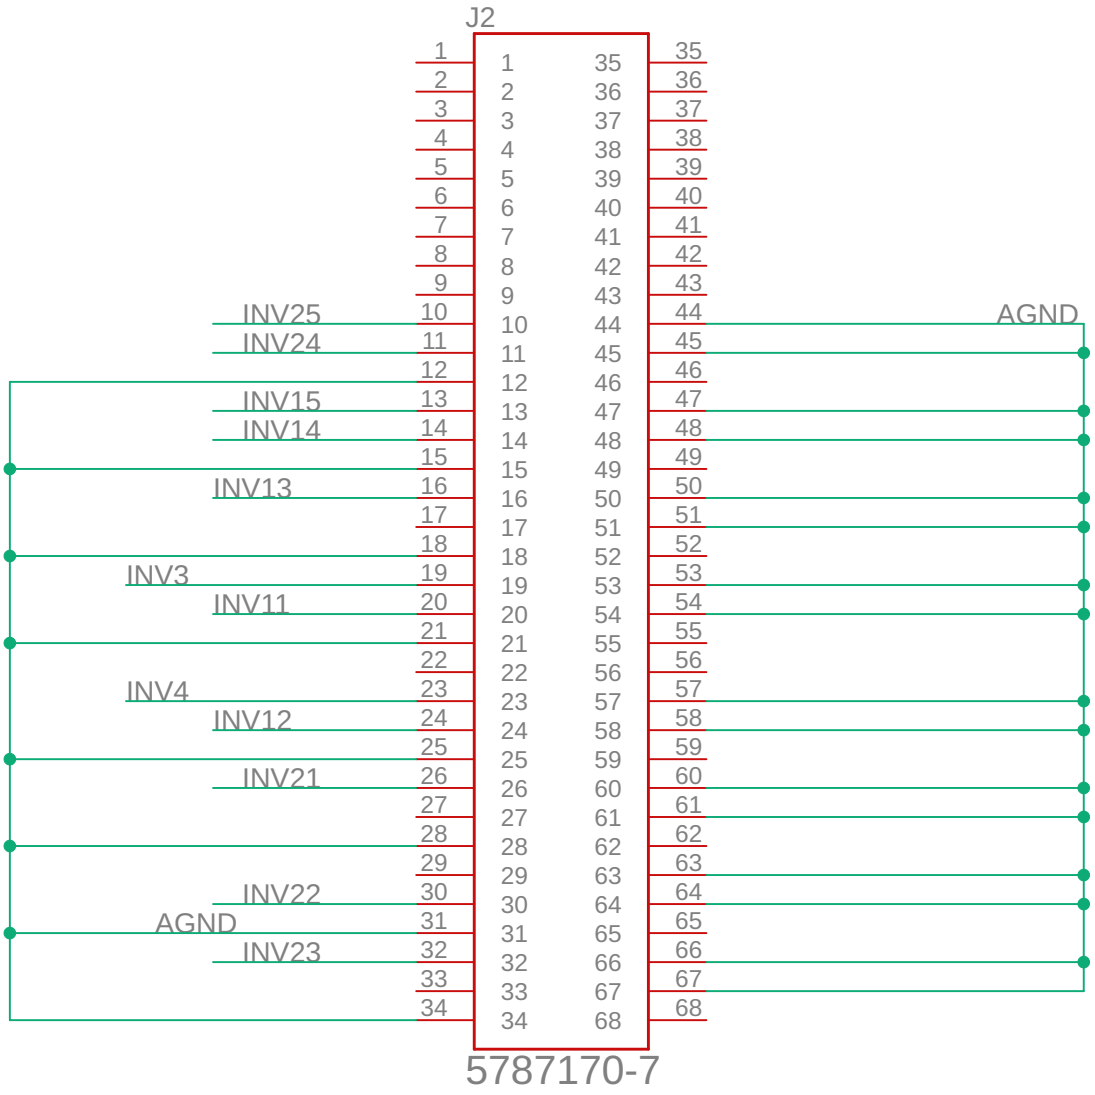

Figure A.10: Circuit schematic for the input 68-pin connector to the auxiliary board from the PXI-6723 DC source.

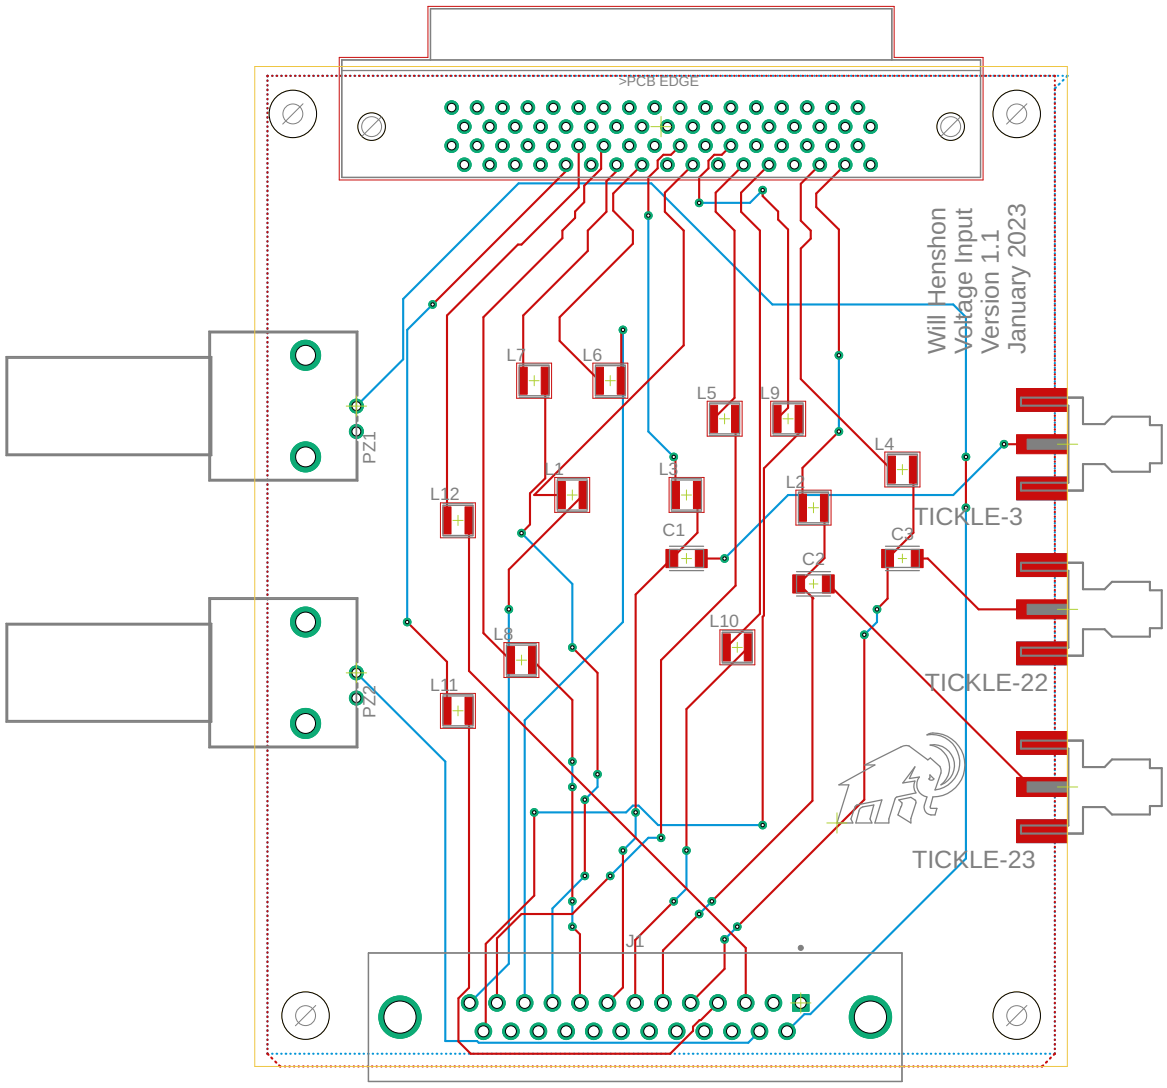

Figure A.11: EAGLE PCB design for the auxiliary circuit. Dimensions of the board are 7.54 cm x 10.16 cm.

A.3 Dummy Trap

This is the dummy trap I used to simulate the trap resonance in the circuit before the adapter on the trap chamber itself was ready for use. Eventually, we hope to phase out

the use of this, since it is an approximation and we should be able to use the full trap. However, because of the spatial concerns outlined in section 5.1.2, it will be difficult to tune the variable capacitors and built-in scope probes until upgrades to the board can be made. Therefore, the dummy trap will remain useful for approximating the tuning at this stage in the experiment.

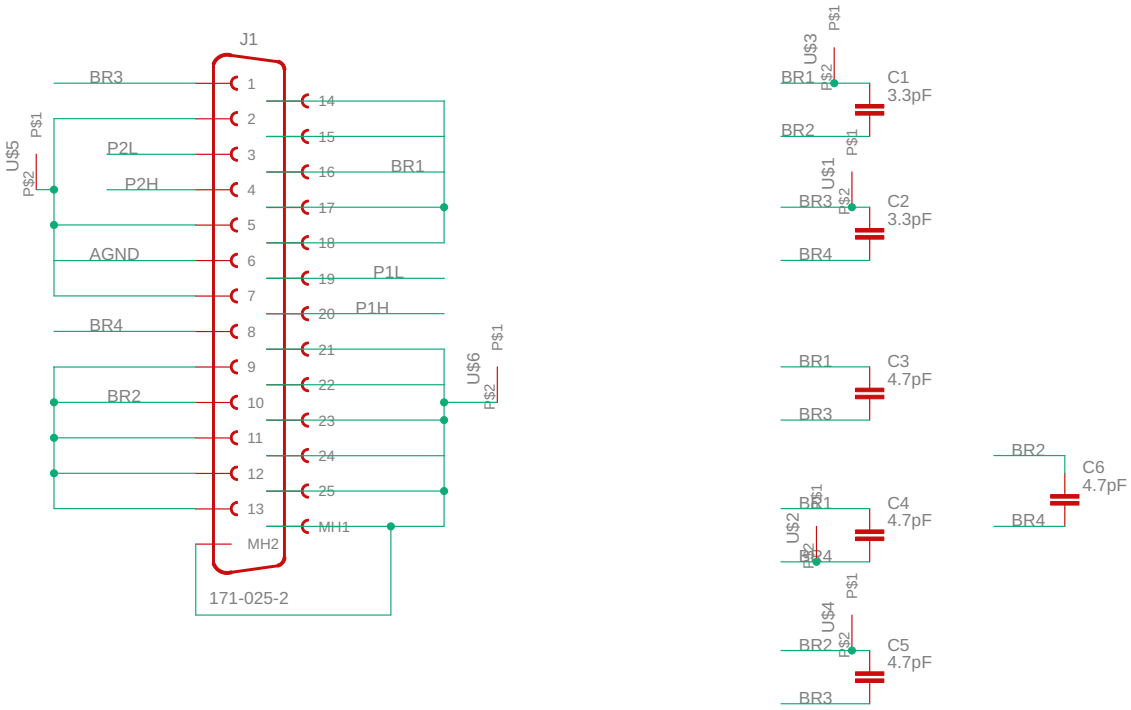

Figure A.12: Schematic diagram of the dummy trap circuit.

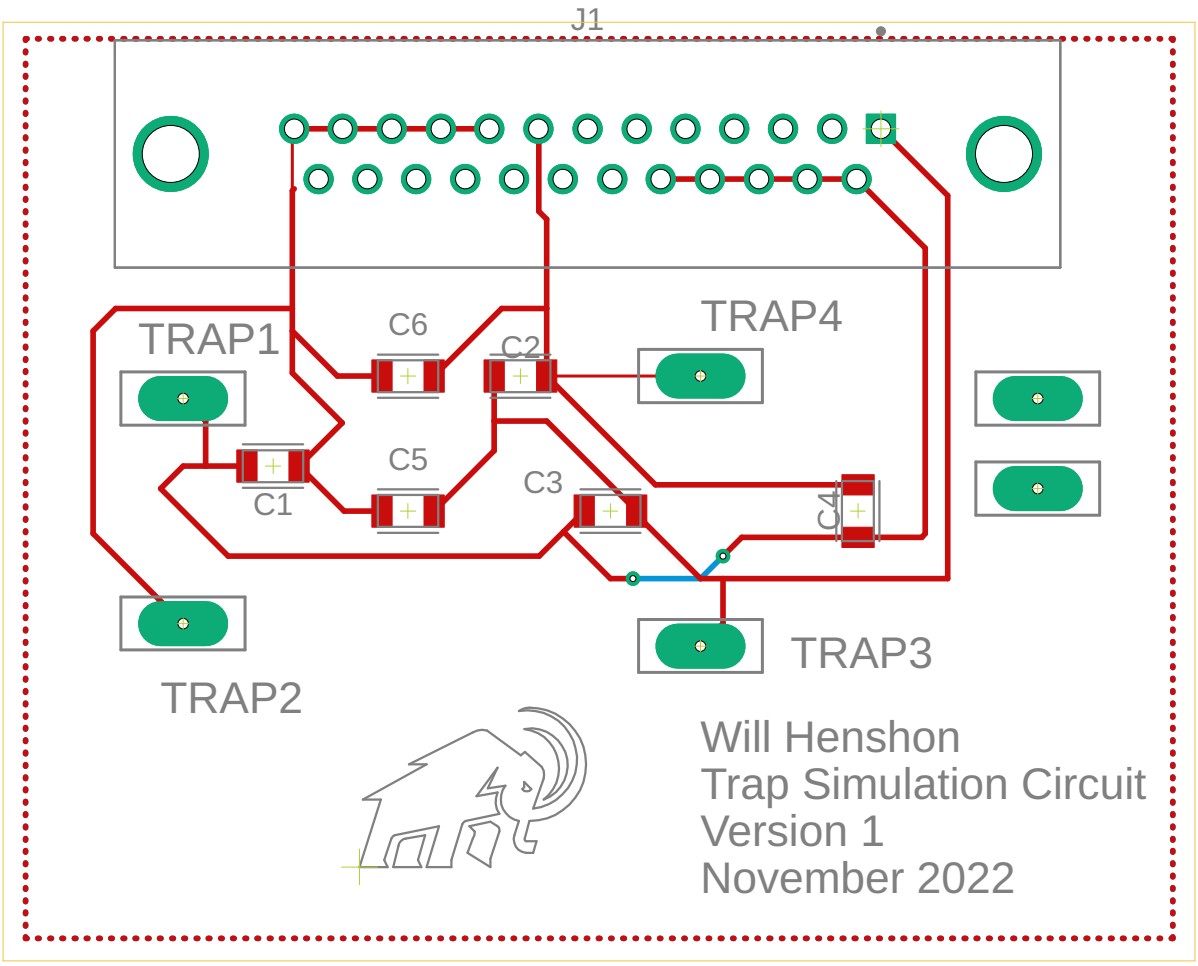

Figure A.13: Eagle PCB design for the dummy trap circuit. The dimensions of the board are 5.3 cm by 6.7 cm.

Bibliography

[1] G. Aad, T. Abajyan, B. Abbott, J. Abdallah, S. Abdel Khalek, A. Abdelalim, O. Abdinov, R. Aben, B. Abi, M. Abolins, O. AbouZeid, H. Abramowicz, *et al.*, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," Physics Letters B **716**, 1 (2012), ISSN 0370-2693, URL https://www.sciencedirect.com/science/article/pii/S037026931200857X.

[2] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, and B. Odom, "New Determination of the Fine Structure Constant from the Electron g Value and QED," Phys. Rev. Lett. **97**, 030802 (2006), URL https://link.aps.org/doi/10.1103/PhysRevLett.97.030802.

[3] NASA Science, "Dark Energy, Dark Matter," Webpage, URL https://science.nasa.gov/astrophysics/focus-areas/what-is-dark-energy.

[4] G. Bertone and T. M. P. Tait, "A new era in the search for dark matter," Nature **562**, 51 (2018), URL https://doi.org/10.1038/s41586-018-0542-z.

[5] J. F. Donoghue, "The Multiverse and Particle Physics," Annual Review of Nuclear and Particle Science **66**, 1 (2016), https://doi.org/10.1146/annurev-nucl-102115-044644, URL https://doi.org/10.1146/annurev-nucl-102115-044644.

[6] G. F. Giudice, "The Dawn of the Post-Naturalness Era," (2017), URL https://arxiv.org/abs/1710.07663.

[7] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, "Search for new physics with atoms and molecules," Rev. Mod. Phys. **90**, 025008 (2018), URL https://link.aps.org/doi/10.1103/RevModPhys.90.025008.

[8] E. W. Kolb, D. J. H. Chung, and A. Riotto, "WIMPZILLAS!" AIP Conference Proceedings **484**, 91 (1999), https://aip.scitation.org/doi/pdf/10.1063/1.59655, URL https://aip.scitation.org/doi/abs/10.1063/1.59655.

[9] L. Berezhiani and J. Khoury, "Theory of dark matter superfluidity," Phys. Rev. D **92**, 103510 (2015), URL https://link.aps.org/doi/10.1103/PhysRevD.92.103510.

[10] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess, "Did LIGO Detect Dark Matter?" Phys. Rev. Lett. **116**, 201301 (2016), URL https://link.aps.org/doi/10.1103/PhysRevLett.116.201301.

[11] J.-P. Uzan, "Varying Constants, Gravitation and Cosmology," Living Reviews in Relativity **14**, 2 (2011), URL https://doi.org/10.12942/lrr-2011-2.

[12] A. Hees, O. Minazzoli, E. Savalle, Y. V. Stadnik, and P. Wolf, "Violation of the equivalence principle from light scalar dark matter," Phys. Rev. D **98**, 064051 (2018), URL https://link.aps.org/doi/10.1103/PhysRevD.98.064051.

[13] D. Hanneke, B. Kuzhan, and A. Lunstad, "Optical clocks based on molecular vibrations as probes of variation of the proton-to-electron mass ratio," Quantum Science and Technology **6**, 014005 (2020), URL https://dx.doi.org/10.1088/2058-9565/abc863.

[14] A. Arvanitaki, J. Huang, and K. Van Tilburg, "Searching for dilaton dark matter with atomic clocks," Phys. Rev. D **91**, 015015 (2015), URL https://link.aps.org/doi/10.1103/PhysRevD.91.015015.

[15] R. Oswald, A. Nevsky, V. Vogt, S. Schiller, N. L. Figueroa, K. Zhang, O. Tretiak, D. Antypas, D. Budker, A. Banerjee, and G. Perez, "Search for Dark-Matter-Induced Oscillations of Fundamental Constants Using Molecular Spectroscopy," Phys. Rev. Lett. **129**, 031302 (2022), URL https://link.aps.org/doi/10.1103/PhysRevLett.129.031302.

[16] A. Shelkovnikov, R. J. Butcher, C. Chardonnet, and A. Amy-Klein, "Stability of the Proton-to-Electron Mass Ratio," Phys. Rev. Lett. **100**, 150801 (2008), URL https://link.aps.org/doi/10.1103/PhysRevLett.100.150801.

[17] A. Hartman, "Two-Photon Vibrational Transitions in O_2^+," Undergraduate Thesis, Amherst College (2022).

[18] B. Kuzhan, "A Molecular Beam Apparatus to Search for Time Variation of Fundamental Constants," Undergraduate Thesis, Amherst College (2021).

[19] D. A. Lane, "Developing a Quantum Toolbox: Experiments with a Single-Atom Harmonic Oscillator and Prospects for Probing Molecular Ions," Amherst College Undergraduate Thesis (2017).

[20] G. Weston, *Ultrahigh vacuum practice* (Butterworths, 1985).

[21] S. Qiao, "Constructing a Linear Paul Trap System for Measuring Time-variation of the Electron-Proton Mass Ratio," Undergraduate Thesis, Amherst College (2013).

[22] Morgan Advanced Materials, "MACOR Properties," Data Sheet, URL https://www.morgantechnicalceramics.com/en-gb/products/macor-machinable-glass-ceramic/macor-properties/.

[23] Leybold, "TURBOVAC 90 i Technical Data," URL https://www.leyboldproducts.us/media/pdf/69/d3/d0/TD_810031V3300_EN.pdf.

[24] Agilent, "Agilent TwisTorr 304 FS," Data Sheet (2020), URL https://www.agilent.com/cs/library/datasheets/public/Agilent_TwisTorr_304FS_turbo_pump_data_sheet.pdf.

[25] SAES Group, "NEXTorr Products Data," Data Sheet, URL https://www.saesgetters.com/sites/default/files/NEXTORR%20Products%20Data%20sheets_3.pdf.

[26] Kurt J. Lesker Company, "Ionization Vacuum Gauge with Integrated Controller and Display," User Manual (2011), URL https://www.thinksrs.com/downloads/pdfs/catalog/BAc.pdf.

[27] Stanford Research Systems, "Bayard-Alpert Gauge Specifications," Data Sheet.

[28] Stanford Research Systems, "Models RGA100, RGA200, and RGA300 Residual Gas Analyzer Manual," (2009), URL https://www.thinksrs.com/downloads/pdfs/manuals/RGAm.pdf.

[29] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions," Rev. Mod. Phys. **75**, 281 (2003), URL https://link.aps.org/doi/10.1103/RevModPhys.75.281.

[30] 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), URL https://link.aps.org/doi/10.1103/PhysRevA.45.6493.

[31] J. K. Pfatteicher, "Developing Rapid Quenching Electronics for Coupling an Ion Trap to a Mass Spectrometer," Undergraduate Thesis, Amherst College (2019).

[32] M. P. C. Solutions, "T184-0," Data Sheet (2019), URL https://datasheets.micrometals.com/T184-0-DataSheet.pdf.

[33] P. F. Herskind, A. Dantan, M. Albert, J. P. Marler, and M. Drewsen, "Positioning of the rf potential minimum line of a linear Paul trap with micrometer precision," Journal of Physics B: Atomic, Molecular and Optical Physics **42**, 154008 (2009), URL https://dx.doi.org/10.1088/0953-4075/42/15/154008.

[34] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Minimization of ion micromotion in a Paul trap," Journal of Applied Physics **83**, 5025 (1998), https://doi.org/10.1063/1.367318, URL https://doi.org/10.1063/1.367318.

[35] D. J. Wineland, W. M. Itano, J. C. Bergquist, and R. G. Hulet, "Laser-cooling limits and single-ion spectroscopy," Phys. Rev. A **36**, 2220 (1987), URL https://link.aps.org/doi/10.1103/PhysRevA.36.2220.

[36] C. M. Pluchar, "An Ultraviolet Laser for Beryllium Photoionization," Undergraduate Thesis, Amherst College (2018).