Progress towards FrAg molecules for nuclear CP violation

by

Andrew Lagno

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics

Waterloo, Ontario, Canada, 2025

© Andrew Lagno 2025

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Ultracold francium silver is a promising experiment that has the potential to set a new upper bound on nucleon electric dipole moments. In working towards making francium silver molecules, our short term goal is to develop the knowledge and ability to evaporatively cool francium and silver. This entails finding the scattering properties of francium and silver using photoassociation spectroscopy and developing the ability to sub-Doppler cool silver. In this thesis, I talk about my work towards this goal, including attempting photoassociation at TRIUMF during francium beam time, work at the University of Chicago towards photoassociation and gray molasses in silver. Even though these efforts weren't successful, the next steps are clear. Additionally, I talk about what I've accomplished at Waterloo when I'm not working on the francium silver project. This includes working towards better control over and stabilization of lasers and experimental optics and beginning optimization of the Cs Zeeman slower.

Acknowledgements

I would like to thank Forouzan Forouharmanesh, Omar Hussein, Paul Del Franco, Megan Byres, Addison Okell and Max Jones at Waterloo for their willingness to answer a question or lend a hand around the lab.

I would like to thank Mohit Verma, Shaozhen Yang, Louis Croquette, and Erik Frieling for making Chicago and Vancouver incredibly welcoming, both in and out of the lab.

Most of all, I would like to thank Prof. Alan Jamison for his guidance, understanding, and willingness to answer my many questions.

Table of Contents

Author's Declaration ii

Abstract iii

Acknowledgements iv

List of Figures vii

List of Tables x

1 Introduction 1

1.1 Baryogenesis 1

1.2 Francium Silver 2

1.3 CP Violation Precision Measurements 3

1.4 The Ultracold Atoms Experiment 5

1.5 Thesis outline 6

2 Laser Cooling 8

2.1 Set-up 8

2.2 Doppler Cooling 9

2.3 Sub-Doppler cooling 16

3 Cs and Li at the University of Waterloo **21**
3.1 Cs Zeeman Slower Setup . . . 23
3.1.1 Cs Oven Design . . . 23
3.1.2 Cs Doppler-Free Absorption Spectroscopy . . . 24
3.1.3 Cs Doppler-Sensitive Fluorescence Spectroscopy . . . 27
3.1.4 Cs Zeeman Slower . . . 30
3.2 Controlling & Stabilizing Lasers . . . 36
3.2.1 Cs Laser Locking . . . 36
3.2.2 AOM Boards . . . 40
3.2.3 Optics Shutters . . . 44

4 Fr at TRIUMF **47**
4.1 Experimental Apparatus . . . 48
4.2 Photoassociation Spectroscopy . . . 51

5 Ag at the University of Chicago **59**
5.1 Silver MOT . . . 60
5.2 Compensation Coil Design . . . 64
5.3 Current Controller Box . . . 68

6 Conclusion **72**

References **74**

List of Figures

2.1 A 1 dimensional representation of the atom's energy as a function of position in the MOT. The magnetic field $B(x) = B'x$ increases the $m_J = +1$ state's energy to the right of the trap and decreases the $m_J = -1$ state's energy, which brings the $m_J = -1$ onto resonance, making that transition more likely, and pushing the atom back to the left when it interacts with the photon. The opposite is true to the left of the trap.[11] 13

2.2 The Sisyphus cooling mechanism. In traveling up the potential, the atom turns some kinetic energy into potential energy. Upon traveling up the potential, the atom is pumped into the lower state. The kinetic energy is lost with the emitted photon. [7] 19

3.1 Caesium's D1 and D2 transitions. [34] The D1 and D2 transitions are the lowest and second-lowest electric dipole moment transitions. 22

3.2 The DFAS optics set-up. The red arrow indicates the path of the laser beam through the oven. 25

3.3 A plot of Cs flux after the DP tube. The error bars were set by the noise on the oscilloscope. 26

3.4 On the left, a zoomed out view of the DSFS set-up. The APD is the black box with the wires coming out of it on the right. The pick-up and focusing lenses are covered with the black box and paper tube to shut out background light. On the right, the Zeeman slower beam and probe beam paths. 28

3.5 The optics imaging the centre of the main chamber onto the APD. The real version has a mirror between the $f = 100$mm and $f = 50$mm lenses, that has been left out for clarity. 29

3.6 Zeeman slower detuning needed to move a peak to the same spot after applying a new bias field. 32

3.7 The top shows the saturated absorption signal, the atoms' absorption of the probe beam, as a function of frequency. The bottom shows the pump and probe beam's behaviour as a function of frequency. On the left and right, the probe beam and pump beam are addressing different velocity classes due to coming in from opposite directions. In the middle, they're both addressing the v=0 velocity class. At this frequency, the pump beam is pumping atoms out of the state that the probe beam can address, leading to a decrease in absorption. [11] 37

3.8 Our error signal, found through modulation transfer spectroscopy, which we use to lock our laser. The error signal is overlayed with the Doppler-sensitive spectrum to identify which peak is being locked to. 39

3.9 On the left, the two Cs spectroscopy boards in the lab. On the right, a simplified diagram of the spectroscopy set-up from Daniel McCarron's thesis. [22] PD1 here is used to observe MTS, while PD2 is used to observe FMS. The flipper mirror is not used in our design. 40

3.10 The typical AOM board. The laser comes onto the board at the top left. The red arrows represent the undiffracted light, the yellow arrows represent light that has been diffracted once, and the blue arrows represent the light that has been diffracted a second time. 41

3.11 The designs for the first and second optics condos. The black rectangle is the optics table, the red rectangles delineate the condos, and the orange squares show the optics boards. The orange arrows indicate which way the boards slide out. It is clear to see that in the initial design, 8 "plots" can be filled with optics boards instead of 6 in the new one. 43

3.12 The finished optics condo, already containing some AOM boards. 44

3.13 On the left, Prof. Isaac Chuang's design for the optics shutter, with the mounted circuit board. On the right, the final design with a BNC port for triggering. 45

4.1 ^{211}Francium's D2 line, at 417.431643THz.[31][15] 48

4.2 The francium ion-catching scheme. The foil, which used to be made of yttrium but is now made of zirconium, catches the francium ions and neutralizes them. The foil flips up to release the francium into the capture chamber, where it's trapped and cooled. While this happens, the Faraday cup detects the francium beam. This signal is used as a diagnostic tool for the francium beam. [35] 49

4.3 The Solidworks model of the foil and capture MOT. The pure glass cell is visible in blue within the exoskeleton cage, which holds the MOT optics in place. [35] 50

4.4 2-photon photoassociation laser scheme. The red arrow labelled PA represents the first PA photon and L2, in blue, represents the second photon in 2-photon PA. The dotted line represents due to from single-photon PA. [21] 54

4.5 The francium MOT near the beginning of the beam time. 56

5.1 The D1 ($S_{1/2} \rightarrow P_{1/2}$) and D2 ($S_{1/2} \rightarrow P_{3/2}$) lines lines of ^{107}Ag and ^{109}Ag. The values in the brackets represent the values for ^{107}Ag. [18][36] 60

5.2 On the left, the replacement lower Bitter coil for our MOT. On the right, the Bitter coil design from Prof. Cheng Chin's lab.[26] 63

5.3 The silver MOT we achieved. This MOT had a temperature of 1mK and a density of 8×10^{10}cm^{-3}. 64

5.4 On the left, the magnetic coil mount for the compensation coil I designed in Solidworks. On the right, the final design of the magnetic field compensation coil, mounted onto the the experimental apparatus. The metal gate valve on the left sets the spatial constraint on this magnetic coil. 66

5.5 On the left, the inside of the current supply box. The rear panel is further from the viewer. On the right, the mounted box with the 4 power supplies and a labelled front panel. 68

5.6 The circuit for the MOSFET relay (or MOSFET switch). The clip-on ammeter that measured the shut-off time is hooked up in series with the magnetic coil. 69

5.7 On the left, the measurement of the time delay between the shut-off signal, in blue, and the magnetic coil shut-off, in purple. The shut-off starts 330µs after the trigger. On the right, the time the magnetic coil shut-off takes, in blue. The current takes 90µs to shut off. 70

List of Tables

3.1 List of AOM boards. The -420MHz shift on the slowing laser AOMs comes from the shift from the spectroscopy board. The -352.5MHz shift on the repump laser AOMs comes from the fact that the laser is locked to a transition that's 352.5MHz lower than the F=3 to F' = 4 transition. 42

Chapter 1

Introduction

1.1 Baryogenesis

It's been observed that the universe contains more matter than antimatter. This is strange because all presently known processes that form matter form antimatter in the same quantity. Baryogenesis is the name given to the hypothesized process that would favour formation of matter over antimatter in the ratios that we see in the universe today. Andrei Sakharov proposed a set of conditions that needed to be true in the early universe for such a process to exist:

- Violation of baryon number conservation
- Interactions out of thermal equilibrium
- C (charge conjugation) symmetry violation and CP (charge conjugation and parity) symmetry violation [28]

Although some CP violation is already accounted for through K and B meson decay (assuming the CPT theorem is valid), not enough yet has been found for baryogenesis. Additionally, beyond the standard model physics, such as supersymmetry, requires even more CP violation. In matter that it couples to, CP violating interactions produces electric dipole moments. Since the CP violation can couple to protons, neutrons, and electrons, we can look for electric dipole moments in these common particles to look for the missing CP violation and to look for new physics.

If a nucleus had an electric dipole moment, the electrons that surround it would rearrange themselves to create an opposing electric dipole moment, shielding it fully from the outside world. But when relativity is brought into the picture, this is no longer the case. The Schiff moment is the leading order electric moment left over due to the incomplete shielding due to the relativistic motion of the electrons. [32] Experiments that look for nuclear electric dipole moments do so by looking for the nuclear Schiff moment.

The intrinsic electric dipole moment has to be locked to the spin of the particle in order to conserve Fermi-Dirac statistics. If they were not locked to each other, there would be a new set of quantum numbers which would affect how many fermions could occupy a given state. For example, if the electric dipole moment had two possible values (up or down), but wasn't tied to the spin's direction, there would be four states that an electron could occupy in an atom's S orbital instead of two. The fact that the intrinsic EDM and the spin are locked is the key to how most intrinsic EDM searches work. Applying an electric and magnetic field along the same direction and in opposite directions will lead to a slight difference in the energies of the two systems, proportional to $\boldsymbol{S} \cdot \boldsymbol{E}$, where $\boldsymbol{S}$ is the Schiff moment and $\boldsymbol{E}$ is the electric field.

1.2 Francium Silver

I'm working on a project that plans to measure nucleon electric dipole moments with ultracold francium silver molecules.

Ultracold molecules are useful because they have a much longer coherence time than warmer molecules, which allows for longer measurements. Additionally, working with such a small cloud of molecules makes it much easier to identify and eliminate systematic error. For example, stray electric and magnetic fields are easier to eliminate in a 1mm diameter at the centre of a vacuum chamber.

We specifically want to use ^{223}Fr for this experiment because it has an octupole-deformed nucleus. Octupole-deformed nuclei have much larger Schiff moments than non-octupole-deformed nuclei. As a result, an EDM measurement that uses an octupole-deformed nucleus is going to be more sensitive than one that doesn't.

Silver is different from any alkali atom in that it has a much higher electron affinity. Alkali-silver molecules are much more polar molecules than alkali-alkali molecules. A FrAg molecule's internal electric field is orders of magnitude larger than the highest fields that can be applied externally. A good example of the largest fields that can be applied in such an experiment is found in the nucleon EDM measurement using mercury, which is discussed

later. The fields achieved in this experiment were 10kV/cm. [13] In radium silver, which has a dipole moment 5/9 the size of FrAg, the field experienced by the radium nucleus is calculated to reach 63.9GV/cm. To polarize the cloud of molecules, a field needs to be applied[10][16]

The first step of this experiment is making FrAg molecules. This will be done by laser cooling both kinds of atom, and making excited state molecules with Feshbach resonance (a technique that involves sweeping a magnetic field across an avoided crossing). From there, the molecules will be transferred to their ground state through stimulated Raman adiabiatic passage (STIRAP). Since current literature involving these atoms only goes as far as putting them in a magneto-optical trap (MOT), the current short to medium-term goal is working up to evaporatively cooling them. In particular, during my master's I have been working towards finding the scattering properties needed to evaporatively cool both francium and silver, and towards sub-Doppler cooling silver.

1.3 CP Violation Precision Measurements

Molecules have been shown to be very effective in precision EDM measurements. In 2018, an upper limit on an electron EDM was set by the ACME collaboration using molecules. This experiment used cold (1K-10K) ThO molecules produced in a cryogenic buffer gas beam. The molecule was chosen based on the large dipole moment and polarizability of one of its state manifolds.

The beam is rotationally cooled and enters a magnetically shielded chamber with uniform, parallel electric and magnetic fields ($\vec{E}$ and $\vec{B}$, setting the quantization $\hat{z}$) where state preparation and then measurement takes place. The atoms are initially pumped into a mixture of $|M = \pm 1, \tilde{N} = \pm 1\rangle$ states, where M is the spin, oriented along the magnetic field and $\tilde{N}$ is the orientation of the molecular dipole moment with the electric field. The atoms are pumped into a state with angular momentum perpendicular to the quantization axis $\hat{z}$ using light that is linearly polarized perpendicular to $\hat{z}$. This results in a superposition of M states in the $\hat{z}$ basis.

Upon reaching this state, the molecules start precessing, accumulating 2ϕ in relative phase difference between the two basis states. The phase winding is described by

$$\phi = -(\vec{B} g \mu_B + \tilde{N} \tilde{\varepsilon} d_e \varepsilon_{\text{eff}}) \tau$$

where on top of the previously stated quantities, g is the state's g-factor, μ_B is the Bohr magneton, $\tilde{\varepsilon}$ is the orientation of the electric field relative to the magnetic field, ε_{eff} is the

field experienced by the electron, τ is the interaction time, and d_e is the electron's electric dipole moment. The final state is measured through pumping along two perpendicular axes and measuring the fluorescence due to each. With careful selection of the two axes, the signal is linear in ϕ. The measurement is repeated with all combinations of flipped $\vec{B}$, $\tilde{N}$, or $\tilde{\varepsilon}$, which allows ϕ and therefore the electron EDM d_e to be backed out. Since the electron EDM wasn't measured, an upper bound of $1.1\times10^{-29}e$ cm (90% confidence) was set, based on the expected sensitivity of the experiment. [2][3]

A second important experiment using molecules set the current limit on an electron EDM in 2023. This experiment used HfF+ ions in an RF trap, which applies a quadrupole electric field rotating at radio frequencies to create a time-averaged local minimum in electric field to trap the ions. The RF trap supplied the electric field for this measurement, and a rotating quadrupole magnetic field was applied to create a time-averaged radially emanating magnetic field.

Like in ACME, the molecules were initially prepared in the four combinations of interatomic axis aligned and anti-aligned with electric field and valence electron spin aligned and anti-aligned with magnetic field. The atoms are put into a superposition of spin states orthogonal to the $\vec{B}$ field with a $\pi/2$ pulse, allowed to precess for time τ, and $\pi/2$ pulsed back. One of the spin states is pushed out of the trap and the number of molecules in each of the interatomic axis states is measured. This measurement is repeated with the opposite spin state. This allows each state's accumulated precession to be mapped onto a population difference between four states. Depending on the initial state from which the superposition was prepared, the molecules experience a different phase build-up due to the electron EDM, which is backed out through the population difference. This experiment set an electron EDM upper limit of $4.1\times10^{-30}e$ cm (90% confidence). [25]

As a measurement of a nuclear EDM and not the electron EDM, FrAg is competing against upper bounds set not by a molecule, but by atomic mercury. This experiment takes place in four hot vapor cells containing ^{199}Hg. Each cell has one of the four possible combinations of magnetic and electric field direction, and they're all placed in a magnetically shielded chamber. Optical pumping is used to spin-polarize the atoms into a spin state orthogonal to the applied magnetic field, as with the other two experiments.

The mercury atoms undergo Larmor precession in the magnetic field, which is measured through Faraday rotation. Linearly polarized light is sent in before and after the precession period to measure the rotation in polarization over this time. The magnetization of the atoms, which rotates as they precess, rotates the polarization of the light. The precession time is extracted from the rotation of the light, and the difference in precession between the cells is used to determine the contribution due to the Schiff moment. The upper bound is

first set on the Schiff moment of the nucleus, and then the maximum possible contribution to the Schiff moment from each kind of nucleon is calculated to set the upper limit on that nucleon. This experiment set upper limits of $2\times10^{-25}e$ cm for protons and $1.6\times10^{-26}e$ cm for neutrons. [13]

We expect that the FrAg experiment will beat this limit. ^{199}Hg does not have an octupole-deformed nucleus and FrAg has increased sensitivity gained by using a polar molecule.

1.4 The Ultracold Atoms Experiment

The goal of the archetypal ultracold atom experiment is to capture and cool a cloud of atoms in a magneto-optical trap (MOT) and to perform experiments with the trapped atoms, with further cooling or trapping steps if needed. Our experiments roughly follow this procedure, with some variations pointed out as encountered. Atoms in a MOT can be cooled to hundreds of microkelvins in temperature, and even colder with other techniques that can only be used after trapping and cooling in a MOT.

The source of the atom gas is typically a chunk of alkali metal being heated in an oven. Since the experiments often need a lower pressure than what is achievable in the oven where the metal is heated, the oven is typically inside a separate chamber from the trapping or experiment chamber, with only a narrow opening to maintain as large a pressure gradient between the two chambers as possible. The oven has a set of apertures that focus the emitted gas into an atomic beam, and often a differential pumping chamber that maintains a pressure gradient between the oven and the experiment chamber. The constraint on the temperature at which the metal is heated is also typically the pressure in the rest of the apparatus. Maintaining as low of a pressure as possible in the experiment chamber is important because the presence of any gases can lead to loss of trapped atoms, resulting in a less precise or even non-existent experiment.

In the laser cooling chapter, we calculated the MOT capture velocity for Cs and Li to be 12m/s to 25m/s and 61m/s to 136m/s respectively. In contrast, an atomic beam of Cs at 100° C and lithium at 350° C have an average velocity of 264m/s and 1604m/s respectively. Since most of the atoms in the atomic beam are moving faster than can be captured by a MOT, they need to be cooled before reaching it, which is often done by a Zeeman slower or a 2D MOT. These techniques are explained in the laser-cooling chapter. A 2D MOT acts as a funnel that cools the atoms transversely and channels them into the

experiment chamber. A Zeeman slower slows a portion of an atomic beam longitudinally, increasing the number of atoms that can be trapped by the MOT.

The Zeeman slower typically stays on until the MOT is fully loaded. The loading time and final size depends on the atom being trapped, the pressure, and the MOT's properties. A normal loading MOT time is on the scale of a few to tens of seconds. After the Zeeman slower shuts off, the MOT stays on for a few hundred milliseconds to a few seconds during which the atoms are thoroughly cooled before moving on to the next stage. Often, the experiment requires the atoms to be colder or confined in a tighter space than what is possible in a MOT. In these cases, the MOT is compressed and shut off, the atoms are sub-Doppler cooled, and/or the atoms are retrapped into an optical dipole trap (ODT). Sub-Doppler cooling are a category of cooling techniques different from Doppler cooling, which can get the atoms to a temperature lower than what's achievable with Doppler cooling. However, the commonly used sub-Doppler cooling techniques are not able to trap the atoms. As a result, this step is typically done for a short period of time ($\sim$ tens of milliseconds) before transfer into another trap, such as an ODT or magnetic trap. These techniques include polarization-gradient cooling and gray molasses.

Optical dipole traps use the dipole force to trap atoms.[11] The experiment phase is often done in an ODT. Evaporative cooling also takes place in an ODT. This is a technique that involves slowly decreasing the depth of the trap by decreasing the intensity of the laser to allow atoms with higher energy to escape. This lowers the average temperature of the trapped atoms over time and in doing so cools them down. This technique is used to create Bose-Einstein Condensates.

1.5 Thesis outline

The thesis body is broken up into four parts: an overview of the laser cooling theory relevant to my work, the work I did at the University of Waterloo, the work I did at the University of Chicago, and the work I did at TRIUMF.

At Waterloo, I worked on getting the Zeeman slower working and characterizing it and I worked on various projects towards the goal of more effectively controlling and stabilizing lasers and experiment optics. The third chapter contains a discussion of the lab and experiments at Waterloo. Next, I talk about the Zeeman slower, and the steps in going from characterizing the flux in the oven to characterizing the effectiveness of the slower. Finally, the third chapter discusses the optics setups and mechanical devices and structures I built and worked on in furthering control over the experiment's optics. This

includes shutters for optics boards, optics set-ups that allow the laser's frequency to be locked to an atomic transition, optics set-ups that allow for the frequency shifting of laser beams, and an optics "condo" that can store built optics set-ups without vibration.

In the fourth chapter I talk about my time at TRIUMF. I describe the unique francium ultracold atom apparatus and talk about doing photoassociation spectroscopy on francium in order to find its scattering properties.

The fifth chapter is about my work at the University of Chicago. I talk about the issues we had with the silver MOT, and the work I did there that was less directly related to the MOT, such as the magnetic field compensation coil I designed and the magnetic coil control box I built.

Finally, in the last chapter I talk about the next steps in each of the experiments.

Chapter 2

Laser Cooling

2.1 Set-up

To begin looking at laser cooling, let's consider the atom as a 2 level system with transition energy $\hbar\omega_0$. When the atom absorbs a photon from a laser at frequency $\omega_L = \omega_0$, it's excited to its higher energy state. Additionally, the atom experiences a momentum kick equal to the photon's momentum $\hbar k$. The two mechanisms for decay from the excited state are through spontaneous emission and stimulated emission. [11]

In the absence of stimulated emission, the excited state has a lifetime τ, giving the rate of spontaneous emission $\Gamma = \tau^{-1}$. Because of the uncertainty relation between the energy and the lifetime of the state, $\hbar\Gamma$ characterizes the energy uncertainty, or linewidth, of the excited state. For this reason, Γ is called the natural linewidth. As a result of this uncertainty, a photon can be detuned by some $\delta = \omega_L - \omega_0$ from the atomic transition and still drive the atomic transition, although with a smaller probability. In undergoing spontaneous emission, the atom decays into its ground state, emitting a photon in a random direction and experiencing a momentum kick in the opposite direction.

Stimulated emission describes the process of an atom in the excited state being de-excited by an incoming photon. In de-exciting, the atom releases a second photon identical to the first and experiences a resulting momentum kick. An important result of stimulated emission is the phenomenon of saturation. The characteristic parameter for this phenomenon is the saturation intensity I_{Sat}. As the intensity I of resonant light becomes much larger than I_{Sat}, the rate of excited atoms undergoing stimulated emission becomes much greater than the rate of spontaneous emission Γ and the population of atoms in the

ground state becomes equal to the population of atoms in the excited state. Additionally, the linewidth broadens as off-resonant light becomes intense enough to drive the transition. This effect of saturation is called power broadening.

Atomic scattering of light means an atom or cloud of atoms absorbing and re-emitting incoming photons, and includes both spontaneous and stimulated emission. The scattering rate due to spontaneous emission is given by

$$R_{Scatt} = \frac{\Gamma}{2} \frac{I/I_{Sat}}{1 + I/I_{Sat} + (\frac{2\delta}{\Gamma})^2} \tag{2.1}$$

where Γ is the natural linewidth, I is the intensity of the light, I_{Sat} is the saturation intensity, δ is the laser detuning $\omega_L - \omega_0$. The fraction I/I_{Sat} is also often replaced with a single variable s, known as the saturation parameter. Finally, the net force experienced by the atoms is equal to the momentum change per photon scattered multiplied by the rate of scattering:

$$\begin{aligned} F_{Scatt} &= \hbar k R_{Scatt} \\ &= \frac{\hbar k \Gamma}{2} \frac{s}{1 + s + (\frac{2\delta}{\Gamma})^2} \end{aligned} \tag{2.2}$$

Now, let's make the atom move with some velocity $\vec{v}$. The frequency of light seen by the atom is Doppler shifted, changing the detuning to $\delta = \omega_L - \omega_0 - \vec{k_L} \cdot \vec{v}$, where $\vec{k_L}$ is the photon's wave vector. This change in detuning leads to a change in the atom's scattering rate, which now depends on the atom's velocity.

2.2 Doppler Cooling

Doppler cooling, the first understood laser cooling mechanism, can be understood using the tools set up in the last section. This mechanism was first applied in the optical molasses technique, which involves shining 3 orthogonal pairs of counterpropagating laser beams that are red detuned from the the atomic transition. (Red detuning means the laser's frequency is lower in frequency than the atomic transition.) The 1 dimensional force due to a single counter-propagating pair of laser beams is given by:

$$
\begin{aligned}
F_{Scatt} &= \frac{\hbar k \Gamma s}{2}\left(\frac{1}{1+2s+(\frac{\omega_L-\omega_0+kv}{\Gamma})^2}-\frac{1}{1+2s+(\frac{\omega_L-\omega_0-kv}{\Gamma})^2}\right) \\
&= \frac{\hbar k \Gamma s}{2}\left(\frac{1}{1+2s+(\frac{\delta_L+kv}{\Gamma})^2}-\frac{1}{1+2s+(\frac{\delta_L-kv}{\Gamma})^2}\right)
\end{aligned}
\tag{2.3}
$$

where $\delta_L = \omega_L - \omega_0$ is the laser's rest frame frequency detuning from the atom's rest frame transition frequency. The sign of each term was chosen by setting the direction of the atom's movement as positive. The saturation parameter in the denominators of each term has a new factor of 2 to account for the fact that each term represents the scattering caused by one of two beams, each at saturation paramater s. The Doppler shift can be treated as small compared to the laser detuning, meaning we can Taylor expand the force:

$$
\begin{aligned}
F_{Scatt} &= \frac{\hbar k \Gamma s}{2}\left(1-2\cdot 4\left(\frac{1}{\Gamma}\right)_L^{\delta}\left(1+2s+4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2}kv\right. \\
&\quad \left.-\left(1+2\cdot 4\left(\frac{1}{\Gamma}\right)^2\delta_L\left(1+2s+4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2}kv\right) \\
&= \frac{\hbar k \Gamma s}{2}\left(-2\cdot 2\cdot 4\left(\frac{1}{\Gamma}\right)^2\delta_L\left(1+2s+4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2}kv\right) \\
&= -8\hbar k^2 s\frac{\delta_L}{\Gamma}\left(1+2s+4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2}v \\
&= -\alpha v
\end{aligned}
\tag{2.4}
$$

This results in a force proportional to the atom's velocity, with a negative coefficient: a frictional force. Optical molasses is so named because of this effect. The atoms' momentum is lost through scattering of light, causing them to slow down.

A technique for trapping and cooling atoms using Doppler cooling that has become ubiquitous in atomic physics is the magneto-optical trap, or MOT. However, Doppler cooling alone cannot trap atoms, only cool them. Trapping happens through the introduction

of the Zeeman effect, which describes the energy shift of an atoms' states in the presence of a magnetic field:

$$\begin{aligned}\Delta E &= -\vec{\mu}\cdot\vec{B} \\ &= -g\mu_B\vec{J}\cdot\vec{B} \\ &= -g\mu_B B m_j\end{aligned} \tag{2.5}$$

The state m_j is defined here along the direction of the magnetic field. The light's detuning from an atomic transition becomes $\delta = \delta_L - \vec{k_L}\cdot\vec{v} + \frac{\mu_B B}{\hbar}(g_e m_{j,e} - g_g m_{j,g})$. The laser's detuning becomes dependent on which of the previously degenerate states is being excited. For simplicity, we will work with a $J = 0$ ground state, which only has projection $m_j = 0$ and a $J = 1$ excited state, which has $m_j = -1, 0, +1$. This means that the $m_{j,g}$ term in the detuning drops out. Finally, we will create a magnetic field with a linear gradient: $\vec{B} = \hat{x}B'x$. This gives us a detuning $\delta = \delta_L - \vec{k_L}\cdot\vec{v} + g\mu_B B' x m_j = \delta_L - \vec{k_L}\cdot\vec{v} + \beta m_j x$.

The 1D force acting on the atoms in these conditions is dependent on which m_f states are excited:

$$\begin{aligned}F = \frac{\hbar k\Gamma s}{2}\Bigg(&\frac{P(\hat{k}\cdot\hat{v}=-1, m_j=+1)}{1+2s+(\frac{\delta_L+kv+\beta x}{\Gamma})^2} - \frac{P(\hat{k}\cdot\hat{v}=+1, m_j=+1)}{1+2s+(\frac{\delta_L-kv+\beta x}{\Gamma})^2} \\ &+ \frac{P(\hat{k}\cdot\hat{v}=-1, m_j=0)}{1+2s+(\frac{\delta_L+kv}{\Gamma})^2} - \frac{P(\hat{k}\cdot\hat{v}=+1, m_j=0)}{1+2s+(\frac{\delta_L-kv}{\Gamma})^2} \\ &+ \frac{P(\hat{k}\cdot\hat{v}=-1, m_j=-1)}{1+2s+(\frac{\delta_L+kv-\beta x}{\Gamma})^2} - \frac{P(\hat{k}\cdot\hat{v}=+1, m_j=-1)}{1+2s+(\frac{\delta_L-kv-\beta x}{\Gamma})^2}\Bigg)\end{aligned} \tag{2.6}$$

where $P(\hat{k}\cdot\hat{v}, m_j)$ denotes the proportion of each beam that excites the atoms to the given m_j state. The $m_j = +1$ state is excited by σ^+ polarized light and the $m_j = -1$ state is excited by σ^- polarized light. The $m_j = 0$ state is excited only by light traveling transversely to the quantization axis (and thus the direction of travel). Therefore, these scatters don't contribute to slowing (or heating) and can be ignored

Linearly polarized light is made up of a superposition of σ^+ and σ^- in equal parts, meaning that $P(\hat{k}\cdot\hat{v}, m_j) = 1/2$ for each non-$m_j = 0$ state. Putting these $P(\hat{k}\cdot\hat{v}, m_j)$

values in and Taylor expanding in small kv and βx results in force

$$F = -8\hbar k^2 s \frac{\delta_L}{\Gamma} \left(1 + 2s + 4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2} v = -\alpha v \tag{2.7}$$

This result is identical to that of molasses. This is to be expected because both sets of beams target both states, leading to an x dependence that averages out to zero.

When both beams consist entirely of σ^+ light, $P(\hat{k} \cdot \hat{v} = \pm 1, m_j = +1) = 1$. The resulting force is asymmetric:

$$F = -8\hbar k^2 s \frac{\delta_L + 2\beta x}{\Gamma} \left(1 + 2s + 4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2} v = -\alpha(1 + 2\beta x/\delta_L)v \tag{2.8}$$

While such a configuration introduces a dependence on position, it preferentially cools in one direction than in the other. Using σ^- light leads to a similar result.

When the laser beam pointing towards the right is σ^- and the beam pointing to the left is σ^+, $P(\hat{k} \cdot \hat{v} = -1, m_j = +1) = P(\hat{k} \cdot \hat{v} = +1, m_j = -1) = 1$ and the rest are zero. Under these conditions, the force is

$$F = -8\hbar k s \frac{\delta_L}{\Gamma}(kv + \beta x) \left(1 + 2s + 4\left(\frac{\delta_L}{\Gamma}\right)^2\right)^{-2} = -\alpha v - \frac{\alpha\beta}{k} x \tag{2.9}$$

This configuration leads to a force that looks like a damped harmonic oscillator. This atom will oscillate around the $x = 0$ position, slowing down the whole time.

This force depends on the saturation parameter and the ratio of the laser detuning to the natural linewidth. When the detuning is small and/or the saturation parameter is large, there is no net force due to both beams scattering in equal amounts. However, a very large detuning and/or small saturation parameter also leads to no force, in the former due to the fact that the beam is so far off-resonance that no scattering occurs, and in the latter simply because the intensity of the light driving the transition is too small. The typical detuning of a MOT is 2 to 3 Γ and the typical intensity used in each MOT beam is roughly the saturation intensity, giving a saturation parameter $s = 1$. A saturation parameter of 1 is chosen to maximize scattering without power broadening the linewidth, which happens increasingly as the saturation parameter increases past 1. Additionally, past a saturation parameter of 1, the transition starts saturating, leading to diminishing returns in scattering rate with increasing power. The detuning which gives the coldest MOT is $-\Gamma/2$. However, a MOT with this detuning can capture only the coldest of atoms, so a higher detuning is

chosen to capture more atoms. A common practice is to compress the MOT by ramping the detuning and intensity down before moving on to the next cooling or trapping step. The magnetic gradient is often also ramped up to trap the atoms more densely.

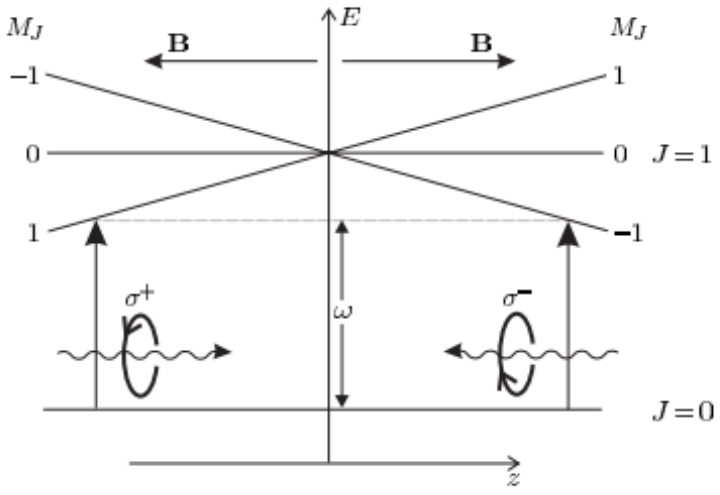


Figure 2.1: A 1 dimensional representation of the atom's energy as a function of position in the MOT. The magnetic field $B(x) = B'x$ increases the $m_J = +1$ state's energy to the right of the trap and decreases the $m_J = -1$ state's energy, which brings the $m_J = -1$ onto resonance, making that transition more likely, and pushing the atom back to the left when it interacts with the photon. The opposite is true to the left of the trap.[11]

Finally, an important quality of the MOT is its capture velocity. This is the maximum velocity of atoms which the MOT can capture. Using scattering force $\hbar k R_{Scatt}$ and setting the stopping distance equal to the diameter of the MOT beams d, we find $v_c = \sqrt{\frac{\hbar k}{m} \cdot R_{Scatt} \cdot d}$. (In the following calculation we choose $d = 2$cm, which is the expected size of our MOT beams.) Assuming a saturation parameter of 1, a detuning of -2Γ, and a permanently negligible Doppler shift, a rough estimate for the minimum capture velocity can be made. Using scattering rate $R_{Scatt} = \frac{\Gamma}{2}\frac{s}{1+s+2(\delta/\Gamma)^2} = \frac{\Gamma}{2}\frac{1}{10}$, we find $v_c = \sqrt{\frac{\hbar k}{m} \cdot \frac{\Gamma}{20} \cdot d}$. For ^{6}Li, this results in a capture velocity of 61m/s, and for Cs, 12m/s. An upper estimate can be made by assuming the laser is always on resonance ($\delta = 0$) with a saturation parameter of 1, giving a constant scattering rate of $\frac{\Gamma}{4}$: $v_c = \sqrt{\frac{\hbar k}{m} \cdot \frac{\Gamma}{4} \cdot d}$. This results in a capture velocity of 136m/s for lithium and 25 m/s for Cs. The real capture

velocity is between these two values, as the typical incoming atom initially sees the light Doppler shifted onto or near resonance, and upon being slowed to low velocities experiences a constant detuning of -2Γ. Finally, the chosen detuning and intensity of the cooling light also affects the capture velocity through the scattering rate.

As mentioned earlier, the $J = 0$ ground state and $J = 1$ excited state is a simplified case. In reality, all alkali atoms have multiple hyperfine ground states and excited states denoted by F. The F angular momentum quantum number comes from the addition of the electron's J angular momentum quantum number to the nuclear angular momentum I. These F states often have frequency spacings on the order of several natural linewidths. As a result, light that is resonant with a transition from one ground state is often far off-resonance from a transition from another ground state. MOTs and optical molasses run on what is known as cycling or closed transitions, in which the highest m_f states of the highest F states of the $S_{1/2}$ ground and $P_{3/2}$ excited states are used. (Atomic physicists refer to the $S_{1/2} \rightarrow P_{3/2}$ transition as the D2 line.) Since transitions have a maximum Δm_f of 1, the excited state can only ever decay back into the state it started in. The only time the atoms end up in a lower ground hyperfine state is when a transition to a lower excited hyperfine state is driven off-resonantly. As a result, a second frequency of light is used to pump the atoms in their lower ground hyperfine state back into the cycling transition. This repump light is usually either shone in as a single beam or coupled into some or all of the MOT arms.

In atoms with relatively large excited state hyperfine splittings, such as Cs and Fr, not very much repump light is needed. ^{211}Fr, which has one of the largest excited state hyperfine splittings, has an excited hyperfine splitting over 70 times larger than Γ, and, on average, takes 10ms to be pumped into its lower state. [15] For reference, the MOT stage generally lasts on the scale of a few hundred milliseconds to tens of seconds, depending on the experiment. To match this scattering rate with near-resonant repump light, a saturation parameter of less than 10^{-3} is needed. However, some atoms, such as both stable Li isotopes and ^{41}K, have excited state hyperfine splittings on the order of or even smaller than a Γ, meaning that non-cycling transitions are driven resonantly or near-resonantly. With such atoms much more repump is needed. Additionally, larger repump intensities are typically used even in the atoms with larger excited state hyperfine splittings to minimize the time spent in the dark state.

The Zeeman Slower also uses Doppler cooling and a magnetic field, but its purpose is to slow an atomic beam rather than to trap a cloud of atoms. The slower is comprised of a long pipe with a non-linear magnetic gradient and a slowing laser beam. The magnetic gradient is designed to keep the slowing light on resonance with the transition as the atoms are slowed. The detuning is linear in velocity due to the Doppler shift being linear

in velocity. Since the velocity lost per scatter is constant ($\Delta v = \hbar k/m$), this leads to a quadratic rate of change in position. Therefore, the desired Zeeman shift and therefore the slower's magnetic field is quadratic in position. The chosen intensity is I_{Sat} and either the detuning, in the case of a decreasing field slower, or the magnetic field's highest value at the end of the slower, in the case of an increasing field slower, is chosen based on the desired final velocity, usually the MOT's capture velocity. The last parameter follows from the largest addressed velocity, which is the velocity addressed at the beginning of the slower. The largest velocity the slower can address depends on the length of the slower, as the field's gradient has to match the atoms' rate of change in velocity as a function of position, set by the scattering rate and the change in velocity per scatter.

Unfortunately, Doppler cooling cannot get a cloud of atoms to 0K. The spontaneous emissions, whose net force averages to zero, heat the atoms. The random emission gives the atom a kick in a random direction each time it scatters light. Additionally, atoms near zero velocity are equally likely to scatter light from either beam. These two effects each result in a random walk. While these two forces average to zero, they create a non-zero mean square velocity $\overline{v^2}$. In other words, they cause velocity fluctuations from the average that represent heating. At steady state, the power absorbed by this heating and dissipated by Doppler cooling are equal. We can find the mean square velocity along the z axis using 3 sets of beams:

$$\begin{aligned} P &= \frac{1}{2}M\frac{d\overline{v_z^2}}{dt} \\ &= 2E_r(2R_{Scatt}) - \alpha\overline{v_z^2} \\ &= 0 \\ \overline{v_z^2} &= 2E_R\frac{2R_{Scatt}}{\alpha} \\ &= 4\frac{\hbar^2k^2}{2m}\frac{R_{Scatt}}{\alpha} \end{aligned} \tag{2.10}$$

where $2E_r(2R_{Scatt})$ is the heating rate and $\alpha\overline{v_z^2}$ is the cooling rate. The first factor of two in the heating comes from the two mechanisms causing this heating, the random walk and the scattering. The second factor of two comes from two laser beams contributing to this effect.

From here, we use the equipartition theorem to find the equilibrium temperature:

$$\frac{1}{2}m\overline{v_z^2} = \frac{1}{2}k_B T$$
$$T = \frac{\hbar\Gamma}{4k_B}\frac{1+2s+4(\frac{\delta_L}{\Gamma})^2}{-2(\frac{\delta_L}{\Gamma})} \tag{2.11}$$

With saturation parameter s approaching 0 and detuning $\delta_L = \omega_L - \omega_0$ at $-\Gamma/2$, the atoms reach the Doppler temperature: $T_D = \frac{\hbar\Gamma}{2k_B}$. [11] For most alkali atoms, the natural linewidth is roughly $2\pi\times$6MHz, giving a Doppler temperature of roughly 140µK.

2.3 Sub-Doppler cooling

When optical molasses was first developed and implemented, one group was amazed to find that the temperature observed in their atomic vapor was unmistakeably smaller than the Doppler temperature. [23] MOTs can also cool below the Doppler temperatures. It turns out that there's a second cooling mechanism present in an optical molasses set-up that can cool atoms below the Doppler temperature. (There are actually two mechanisms that can be present in optical molasses. The mechanism that is present depends on the polarization configuration of the molasses.) These are members of a class of cooling mechanisms called sub-Doppler cooling. The cooling mechanisms seen in optical molasses are called polarization gradient cooling (PG cooling).

Sisyphus cooling is one of the two PG cooling mechanisms and involves optical molasses with 2 orthogonal linear polarizations. [7][11] This creates a polarization gradient that cycles between linear, σ^-, orthogonally linear, and σ^+ on a length scale of $\lambda/2$. For simplicity, we will work with a $J = 1/2$ ground state, with $m_j = \pm 1/2$. We also need to briefly look at the AC Stark shift and the Rabi frequency. The Rabi frequency is the frequency at which an atom's wavefunction oscillates between two states when those two states are coupled by resonant light. In calculating it, we assume that the wavelength is much bigger than the size of the atom, which allows us to treat the light simply as a uniform electric field oscillating in time. The Rabi frequency frequency is then given by

$$\Omega = \frac{\langle 1|\vec{E}\cdot\vec{d}|2\rangle}{\hbar}$$
$$= \frac{e|E_0|\langle 1|\vec{x}|2\rangle}{\hbar} \tag{2.12}$$

where $\vec{E}$ is the electric field of the light, $\vec{d} = e\vec{x}$ is the dipole moment operator, e is the electron charge, and $|1\rangle$ and $|2\rangle$ are the two coupled states. [11] The position overlap of the two states $\langle 1|\hat{x}|2\rangle$ is the product of the radial and angular overlaps of the two states' wavefunctions. The radial overlap doesn't change between the transitions we are looking at, whereas the angular overlap is proportional to the Clebsch-Gordan (CG) coefficient between the two states c_{12}. The Clebsch-Gordan (CG) coefficient describes the relative couplings between different states. The square of the CG coefficient $|\langle 1|\hat{x}|2\rangle|^2$ is proportional to the transition probability.

When a laser beam resonantly couples two states, they experience a frequency shift away from each other. When the light is red-detuned, the upper energy state gets shifted to a higher energy and the lower state gets shifted to a lower energy. This frequency shift is known as the AC Stark shift. The energies of the two states coupled by off-resonant light is given by

$$\omega_{\pm} = \pm\frac{\sqrt{\Omega^2 + \delta^2}}{2} \tag{2.13}$$

where Ω is the Rabi frequency and δ is the laser's detuning from the $|1\rangle \rightarrow |2\rangle$ transition, including the Doppler and Zeeman shifts. (Although we expect both to be negligibly small or nonexistent in PG cooling.) When the Rabi frequency is much smaller than the detuning, we find

$$\omega_{\pm} = \pm\left(\frac{\delta}{2} + \frac{\Omega^2}{4\delta}\right) \tag{2.14}$$

The splitting of each state due to the AC stark shift is given by

$$\Delta\omega_{\pm} = \pm\frac{\Omega^2}{4\delta} \tag{2.15}$$

Using the Rabi frequency we found earlier, we get

$$\begin{aligned} \Delta\omega &= \frac{e^2|E_0|^2|\langle 1|\hat{x}|2\rangle|^2}{4\hbar^2\delta} \\ \Delta\omega &\propto (I/I_{Sat})(\Gamma/\delta)c_{12}^2 \end{aligned} \tag{2.16}$$

The AC Stark shift is proportional to the intensity and the square of the CG coefficient c_{12}^2, and inversely proportional to the detuning.

The σ^+ light AC Stark shifts the $m_j = 1/2 \rightarrow m_j' = 3/2$ transition more than the $m_j = -1/2 \rightarrow m_j' = 1/2$ because its CG coefficient is larger between the two. On the other hand, the σ^- light AC Stark shifts the $m_j = -1/2 \rightarrow m_j' = -3/2$ transition more than the $m_j = 1/2 \rightarrow m_j' = -1/2$. Since the polarization cycles between σ^-, linear, and σ^+ over space, the two ground states oscillate in energy perfectly out of phase with the polarization gradient.

An atom at the bottom of the potential in one of the states travels to the top of the potential energy hill, slowing down, and scatters light. The higher energy ground state whenever the polarization is circular is the ground state in the $m_j = -1/2 \rightarrow m_j' = 1/2$ or $m_j = 1/2 \rightarrow m_j' = -1/2$ transition because that's the transition with the smaller CG coefficient, leading to a smaller AC Stark shift. Upon scattering light, the atoms are more likely to end up in the lower energy state because the excited state has a higher CG coefficient with that state. (For example, the $m_j' = +1/2$ state has a higher coupling to the $m_j = +1/2$ state than $m_j = -1/2$, meaning that when the light is scattered, the atom scatters to $m_j = +1/2$ rather than its original state.) This results in the atom at the bottom of the potential again, but without the kinetic energy it lost climbing the potential hill. This mechanism is called Sisyphus cooling because of its similarity to King Sisyphus' mythical punishment: rolling a boulder up a hill in perpetuity, only for the ball to roll back down once he reaches the top.

As seen earlier, the AC Stark shift $\Delta\omega$ is proportional to the the saturation parameter $s = I/I_{Sat}$, and inversely proportional to the detuning δ/Γ. Since the energy dissipated per scatter is proportional to the AC Stark shift, we might guess that as the atoms become slower, we can ramp $I/I_{Sat} \rightarrow 0$ or $\delta/\Gamma \rightarrow \infty$, bringing the temperature down to zero. However, even when an atom is stuck at the bottom of a well at v=0, it still scatters light, getting a momentum kick from the photons. The net power dissipated is zero when the energy dissipated per scatter is equal to the recoil energy. Using equipartition again, we find the recoil temperature, where the average energy is equal to the recoil energy:

$$\begin{aligned} \frac{1}{2}k_B T &= \frac{1}{2}mv_r^2 \\ T &= \frac{mv_r^2}{k_B} \\ &= \frac{\hbar^2 k^2}{mk_B} \end{aligned} \tag{2.17}$$

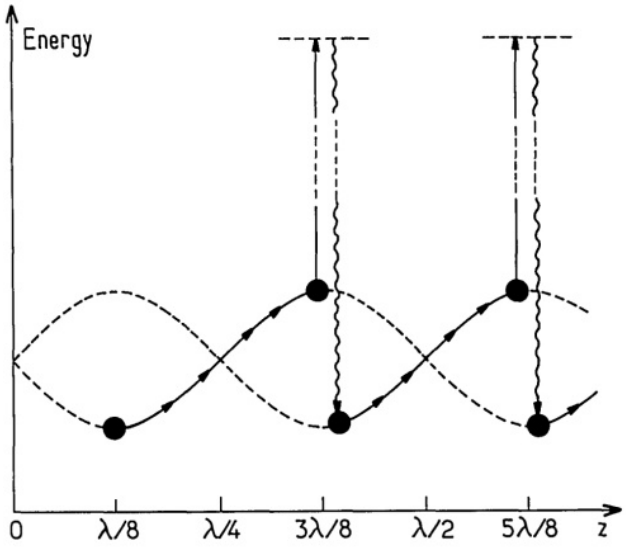


Figure 2.2: The Sisyphus cooling mechanism. In traveling up the potential, the atom turns some kinetic energy into potential energy. Upon traveling up the potential, the atom is pumped into the lower state. The kinetic energy is lost with the emitted photon. [7]

This temperature is very small for heavier elements, such as Cs, Rb, and Fr. For Cs and Fr, the recoil temperature is roughly 0.2µK. For Li, the lightest alkali, it's 7µK. However, the recoil temperature is rarely achieved with PG cooling because of Doppler heating.

A second PG cooling mechanism exists that similarly cools optical molasses to sub-Doppler temperatures, but with counter-propagating σ^+ and σ^- beams. [7] This polarization configuration creates a linear polarization that rotates in space. The rotating polarization results in moving atoms being pumped into the state more likely to scatter light from the counter-propagating beam. The rate at which they're pumped into this state is dependent on their velocity and the AC Stark shift due to the light, leading to a force dependent on the atoms' velocity that can also be tuned through intensity or detuning. Like the other PG cooling mechanism, this mechanism can cool to the recoil limit. Its scaling in intensity and detuning is the same as well.

Both PG cooling mechanisms only work for clouds with average velocities on the scale of $\lambda\Gamma$, where λ is the light's wavelength and Γ is the natural linewidth. As a result, PG cooling mechanisms only take over after Doppler cooling has brought the atoms to such temperatures. Additionally, when multiple excited hyperfine states are present, they need to be well-resolved in order for PG cooling to work. This is because of their CG coefficients,

the magnitudes of which are different for different angular momentum excited states. When many excited F' states are present, the light shifts from the different m'_f states add up to a less effective potential for cooling.

A final note about PG cooling is that these mechanisms can work on a $F \rightarrow F' = F$ transition. (This is true when F is bigger than 1/2.) However, this transition has inverted CG coefficients, meaning that all else being equal, PG cooling becomes a PG heating mechanism. Using a blue detuning turns this back into a cooling mechanism. This technique, called blue molasses, is useful when the hyperfine splitting on the D2 transition (The $S_{1/2} \rightarrow P_{3/2}$ transition) are too small for PG cooling to work properly, such as with lithium, because it allows you to sub-Doppler cool on the D1 transition ($S_{1/2} \rightarrow P_{1/2}$), which otherwise would not be possible due to the lack of cycling transition.[8]

Finally, gray molasses is a sub-Doppler cooling technique that can cool more effectively through the creation of states which don't scatter light, called dark states. The source of these dark states is the coupling of two m_f states through the presence of linear light tuned to a $F \rightarrow F' = F$ or $F \rightarrow F' = F - 1$ transition. This beam is blue-detuned for reasons mentioned in the previous paragraph. Instead of a PG cooling mechanism, there is an AC stark shift arising from an intensity gradient, caused by a standing light wave. This intensity gradient results in a Sisyphus cooling mechanism, with the atoms being pumped into the dark state. Moving at a large enough velocity $v \sim \frac{k}{s\delta}$, the atomic states do not adiabatically follow the intensity gradient, leading to the dark states remixing into bright states and allowing the atom to be further cooled. (The name gray molasses comes from using dark states which aren't fully dark - gray states) At small velocities, the dark state follows adiabatically, leading to the atom staying in the dark state and not being heated by scattering. States at v=0 accumulate in the dark state, meaning that this cooling mechanism isn't subject to the recoil limit. [38]

A newer innovation with gray molasses uses a second frequency, coupling the second ground hyperfine state to the excited state. This technique is called Λ-enhanced gray molasses. When the difference in detuning between the two lasers is small compared to the Raman linewidth Δ, the two ground states become coupled. The coupling of two F states leads to more effective cooling through the introduction of a larger energy scale for energy dissipation. Because it isn't dependent on Raman coupling between two hyperfine states, this mechanisms also seems to allow the use of PG cooling mechanisms instead of a standing intensity wave. The Raman linewidth depends on many different factors. Literature gives linewidths on the scale of a few 0.01Γ for potassium [29][33] and rubidium [24] and on the scale of 0.1Γ for lithium. [12][33]

Chapter 3

Cs and Li at the University of Waterloo

While I did not do any work towards the FrAg experiment at Waterloo, I learned most of the experimental techniques and physics background that I applied at TRIUMF and Chicago here.

At Waterloo, we use fermionic ^{6}Li and bosonic ^{133}Cs. Unlike silver or francium, these are relatively common and simple alkali atoms. Cs has large (compared to its linewidth) hyperfine splittings between its ground and excited states, which means it can be sub-Doppler cooled easily. ^{6}Li has small excited state hyperfine splittings, which means that it's difficult to cool and requires a lot of repump light. However, it's easier to tune interactions in lithium than in other alkali atoms, which makes it desirable for our research purposes.

Using two different species of atoms allows for a wide variety of available experiments. In our current experimental apparatus, we will be doing quantum simulation experiments with bosonic Cs, fermionic Li, and boson-fermion mixtures. In our second experimental apparatus, which will be built after the first one is finished, we will be doing quantum chemistry experiments using LiCs molecules.

The Waterloo experimental apparatus follows the structure outlined in the first chapter, with the exception that we are using two species of atoms, lithium and caesium. As a consequence, there are two ovens, two DP chambers, and two Zeeman slowers oriented orthogonally, that lead into a single main chamber. Since the different atoms have different atomic transition wavelengths, there need to be two sets of MOT and sub-Doppler cooling optics. Finally, holding Cs and Li together in overlapping MOTs leads to atom loss, so

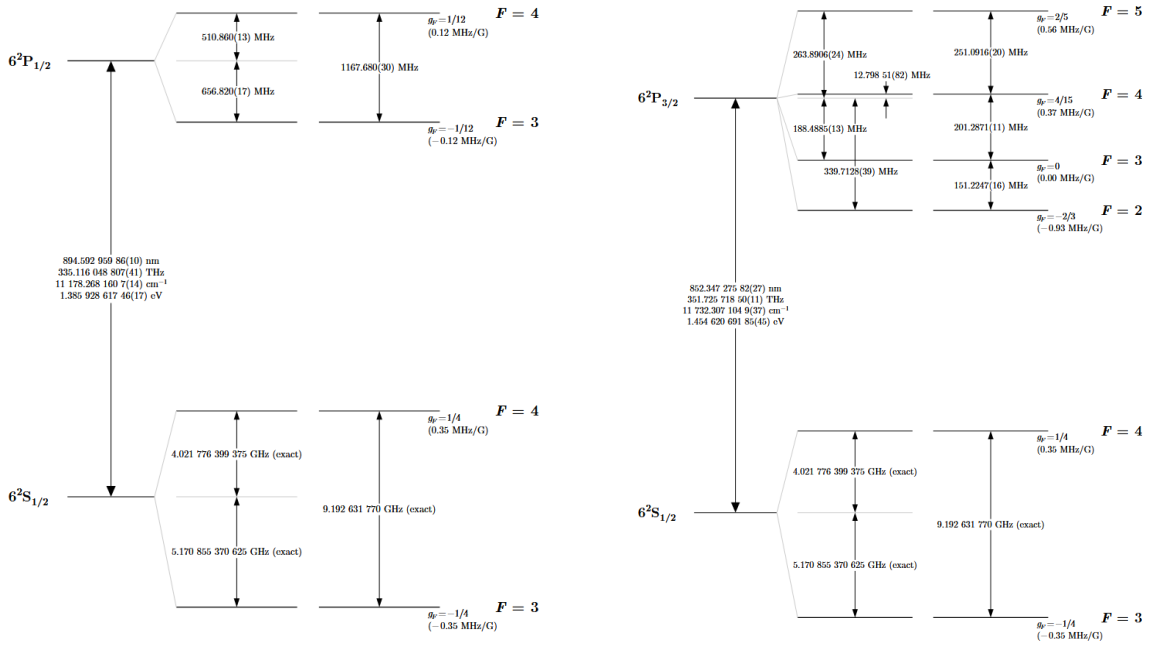


Figure 3.1: Caesium's D1 and D2 transitions. [34] The D1 and D2 transitions are the lowest and second-lowest electric dipole moment transitions.

there needs to be optics that accommodate two different ODT beams. In experiments that will use both Li and Cs, after one species is captured, cooled, and loaded into an ODT, the ODT will move the atoms out of the way of the second species, and then the process will be repeated with the other kind of atom. The first ODT will be brought back to the centre of the chamber to mix the two trapped species before the experiment begins.

At Waterloo, I've worked on two notable projects: optimizing our Cs Zeeman slower and controlling and stabilizing our lasers and experiment optics. In preparation to set up our Zeeman slower, I characterized the Cs atom flux out of our oven. After this, I assembled an optics setup that allowed us to see the velocity-space distribution of atoms that are reaching our main chamber. Finally, I worked on getting the slower working and finding the settings at which slowing is the most effective. The next step of this process is to build a MOT and to optimize the MOT loading rate by varying the Zeeman Slower settings. Towards the goal of controlling and stabilizing our lasers and experiment optics, I worked on the Cs laser locks, modularized optics set-ups that allow for shifting the frequencies of the Cs laser beams, stable storage for these and other optics set-ups, and home-built optics shutters. The next steps in this project are setting up locks for the Li lasers and

optimizing the frequency-shifting optics set-ups for the Li lasers.

3.1 Cs Zeeman Slower Setup

3.1.1 Cs Oven Design

Knowledge of the geometry of the oven and Cs-side vacuum chamber is important to understand the steps taken in setting up the slower.

The Cs metal is kept in a cup, connected to an elbow. A nozzle is on the other end of the elbow. The nozzle is made up of an array of hypodermic needles which act as a set of apertures which collimate the gas that travels through it into an atomic beam. After coming out of the oven, the atoms travel through a chamber made up of a six-way cross. Two opposite ports on the cross are attached to the nozzle and the differential pumping (DP) tube, through which the atoms travel to reach the next chamber. The bottom port is connected to a ribbed bellows that allows an attached stiff metal flag to be moved in and out of the way of the atomic beam to block or unblock it, and the three remaining ports each have viewports.

After passing through the differential pumping (DP) tube, the atoms travel through a differential pumping (DP) chamber and enter the Zeeman slower. The DP chamber is made up of a six-way cross with two opposite ports attached to the DP tube and Zeeman slower. The remaining ports sport three viewports and a 75L/s ion pump. The purpose of the DP chamber is to maintain a larger pressure gradient between the oven and the main chamber. There are three parts to this: creating a point where only a small amount of the background gas from the oven can enter, pumping to decrease the equilibrium pressure in that chamber, and creating an opening through which only a small amount of background gas can escape into the main chamber. This technique is called differential pumping.

The DP tube is a tube with a large length-to-width ratio that is used as the opening to the DP chamber. This makes it more selectively allow through atoms that are a part of the atomic beam, which is co-axial with the differential pumping tube, than atoms that are just a part of the background gas. This is based on the fact that atoms that hit the side of the tube are as likely to bounce forwards as they are backwards. Atoms that have a small transverse velocity are a part of the atomic beam and thus pass through the tube without bouncing, whereas atoms that have a large transverse bounce multiple times and have a high chance to bounce back into the oven. The Zeeman slower tube connects to the

other end of the DP chamber. It's much longer than it is wide, so it acts as a DP tube on the other side of the chamber. Finally, after the Zeeman slower is the main chamber.

3.1.2 Cs Doppler-Free Absorption Spectroscopy

Doppler-free absorption spectroscopy (DFAS) is a way of measuring the caesium flux coming out of the oven. It's the first step of optimizing the Zeeman slower and is important because it makes sure the oven is working as expected before attempting to slow the atoms. The data taken here can be used in the future to diagnose potential issues with the oven. Additionally, it will be helpful for calibration when setting up the slower.

This process involves sending a laser beam transversely through the atomic beam coming out of the nozzle and measuring the light absorbed by the atoms. The probe laser is swept in frequency around the transition from ^{1}33Cs's ground hyperfine $F = 4$ state to the $P_{3/2}$ manifold (the D2 transition) in order to see the three transitions (with excited hyperfine states $F' = 3,\ 4$, and 5). When the probe's frequency is resonant with an atomic transition, the atoms scatter the light, leading to a sharp drop in transmitted intensity at that point in time. Using the size of the drop in transmission, the atomic flux can be calculated with Beer's law.

The atoms' maximum transverse velocity in the first post-oven chamber is set by the hypodermic needles, which have a width to length ratio of 0.017":0.25". This gives a transverse velocity of $(0.017/0.25) \cdot 260\text{m/s} = 17.7\text{m/s}$ (with 260m/s being the atoms' most likely transverse velocity at 90°C). This velocity has a 20.8MHz Doppler shift inside the oven. The maximum transverse velocity in the differential pumping chamber is set by the aperture formed between the nozzle and the end of the differential pumping tube. The nozzle has a radius of 2.5mm and the DP tube a radius of 5.1mm, and they are about 15cm apart. An atom flying from one edge of the nozzle to the opposite edge of the DP tube has a transverse velocity of $((5.1 + 2.5)/150) \cdot 260\text{m/s} = 13.2\text{m/s}$. This velocity has a 15.5MHz Doppler shift inside the oven. Since the atomic linewidth is narrower than the transverse Doppler spread of the atoms, we are slightly underestimating the atomic flux coming out of the oven by using Beer's law. However, this underestimate is fine, since the atoms that we see at the height of the peak are the ones most likely to make it to the centre of the experiment chamber.

The probe beam comes in through a fibre coupler mounted via cage mount to the viewport on one side of the differential pumping chamber. A photodiode is mounted to the viewport opposite to it. The cage mount aligns the centres of the probe beam source and the photodiode to the centres of the viewports and chamber. The atomic beam travels

through the centre of the chamber perpendicular to the probe beam. The laser beam used in the spectroscopy setup had a diameter of roughly 4mm. In contrast, the atomic beam has a diameter of 5mm, set by the size of the nozzle. The probe beam being smaller than the atomic beam guarantees that all of the laser is passing through the atomic spot if well aligned, which minimizes the signal's background. A small power of 36µW was chosen to avoid saturating the transition and to minimize the signal's background. The laser's intensity was 340µW/cm^2, which is 12.5% of caesium's saturation intensity, 2.71mW/cm^2. [34]

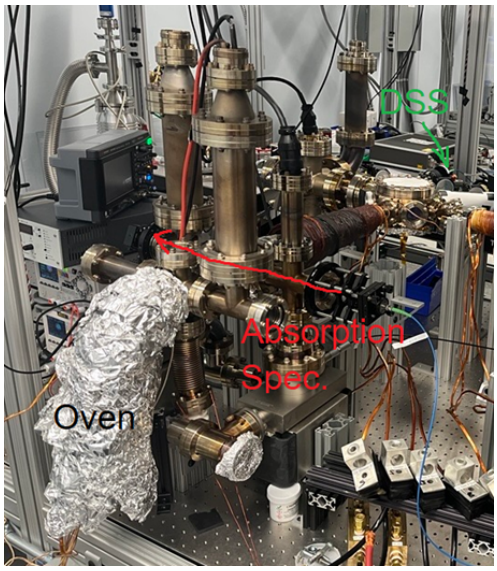


Figure 3.2: The DFAS optics set-up. The red arrow indicates the path of the laser beam through the oven.

The atomic flux is calculated by multiplying the atomic density in the beam by the most probable velocity in the beam. The most probable velocity is described by the formula $v = \sqrt{3k_BT/m}$, where k_B is Boltzmann's constant, T is the oven's temperature, and m is the atomic mass. [11]

The density is found using Beer's law: $I = I_0 \exp(-l\sigma n)$, where I_0 is the intensity going into the cloud, I is the intensity measured after the light passes through the atoms, l is the distance the light travels through the medium, n is the atom density, and σ is the optical scattering cross-section, described by $3\lambda^2/2\pi$. [11] One can rearrange to find the density as a function of absorption. Since a photodiode's signal is proportional to intensity, we use the signal from the photodiode for these quantities. The DC voltage of the photodiode

is used as a stand-in for the intensity going into the cloud, and the size of the dip is the amount of scattered light. The difference between these two is the transmitted intensity. Since the spacing of Cs's two ground states is 9.2GHz, [34] which is a larger frequency than the laser can sweep, we can only look at one ground hyperfine state at a time. So the density needs an extra factor to account for the fact that we're only looking at the atoms in one of the two ground states, $F = 4$. The ratio of these two states' Boltzmann factors is $\exp(-h \cdot 9.2\text{GHz}/k_B T)$ is very near 1 at room temperature. Therefore, the factor only needs to account for the multiplicities of the two states. $F = 4$ has 9 m_F states, whereas the other ground state, $F = 3$, has 7. This gives us a correction factor of $(9+7)/9 = 16/9$ to find the true density.

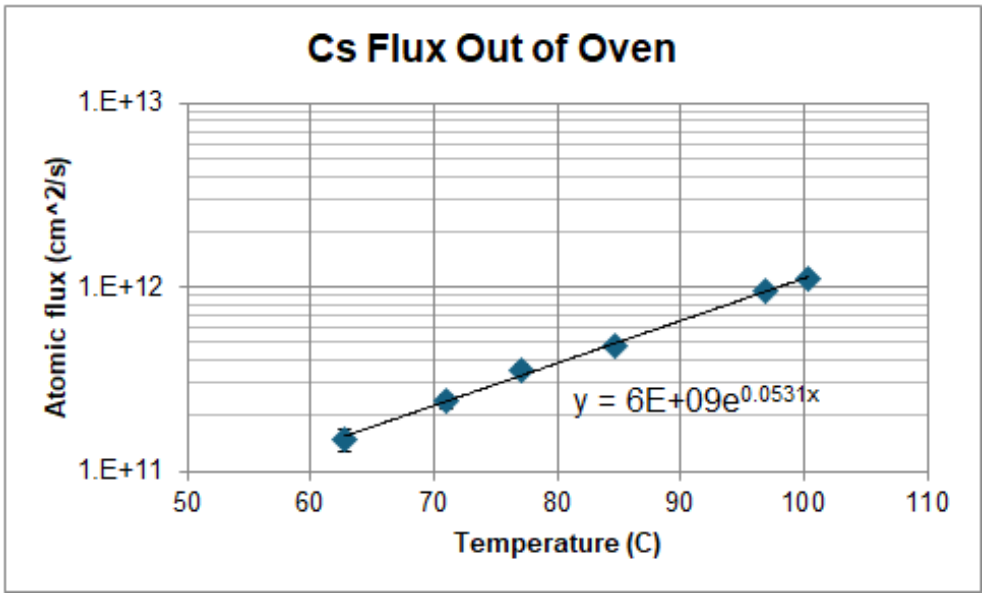


Figure 3.3: A plot of Cs flux after the DP tube. The error bars were set by the noise on the oscilloscope.

As expected, increasing oven temperature led to larger dips, signaling increasing flux. The flux was measured between 60°C and 100°C. These temperatures were chosen because the Cs vapor pressure at these temperatures is comparable to the pressures typically used for alkali atom ovens, including our Li oven. The flux we saw ranged from $1.48\times10^{11}\text{cm}^2/\text{s}$ at 62.8°C to $1.11\times10^{12}\text{cm}^2/\text{s}$ at 100.3°C. This data can be seen in Figure 3.3. Using this data, we can estimate the MOT loading rate for each temperature, which tells us how hot we need to run our oven. Calculations done by another student show that about 3.3% of the atoms in the atomic beam at the location of the MOT are slowed to a velocity

that the MOT can capture. The distance from the nozzle to the MOT is 80cm and the MOT's capture cross-section is roughly a square with 1.5cm sides. Assuming the atoms' flux is inversely proportional to r^2, and the flux was measured 15cm from the nozzle, we estimate that the flux at the MOT capture volume will be $5.20\times10^{9}\text{cm}^{-2}\text{s}^{-1}$ at 62.8°C to $3.90\times10^{10}\text{cm}^{-2}\text{s}^{-1}$ at 100.3°C. Assuming we can capture 3.3% of the atoms, and using our MOT capture cross-section, we get a MOT loading rate of $2.74\times10^{8}\text{s}^{-1}$ at 62.8°C to $2.05\times10^{9}\text{s}^{-1}$ at 100.3°C. A MOT loading rate greater than 10^{8}s^{-1} is considered good, so we are able to run the oven even at lower temperatures and get a good MOT loading rate.

3.1.3 Cs Doppler-Sensitive Fluorescence Spectroscopy

The next step of optimizing the Zeeman slower is Doppler-sensitive fluorescence spectroscopy. This is a way of measuring the velocity distribution of the atoms as they reach the location at which they'll be trapped. This technique is the main imaging technique used for optimizing the Zeeman slower, as it allows one to see the number of atoms at any given velocity.

Doppler-sensitive fluorescence spectroscopy (DSFS) has two important qualities: its sensitivity to Doppler shifts and the fact that it looks at atomic fluorescence rather than absorption. The first quality is what allows this technique to see the atoms' velocity distribution. The second quality allows us to see a smaller number of atoms, which is important due to the small number of atoms in the imaging volume. A beam from the same probe laser used in DFAS excites the atoms passing through that location, and the atoms fluoresce as they de-excite. (This beam will be referred to as the probe beam in this section.) The light scattered by the atoms is collected and measured. The probe beam is set up at a 135 degree angle to the atomic beam. This causes the probe beam's frequency to be Doppler shifted in relation to the atoms. Since the intensity of scattered light is proportional to the number of atoms for which that light is resonant, the result is a signal that shows the velocity distribution of the atoms. This is very useful in a preliminary optimization of the Zeeman slower, as seeing the number of atoms at a certain velocity gives you a good measure of whether any given settings are improving slowing.

The probe beam optics are set up on an optics board, with all of the optics level with the viewport. The beam was set up with a diameter of roughly 2.4cm. In this measurement, a larger spot size isn't as detrimental, because it doesn't directly increase the background of the measurement. Additionally, the signal is normally very small so it's helpful to make as many atoms fluoresce as possible. However, a larger spot means more light scattering in the chamber and contributing to a background, so after the beam was aligned an iris

was put in to shrink the beam until the the feature started to shrink. The final spot size was about a centimetre in diameter. The total power coming out of the fibre is 3.1mW, which gives an intensity of 685µW/cm^2, about a quarter of the saturation intensity. The intensity used in DSFS is higher than in DFAS because the signal is so much smaller.

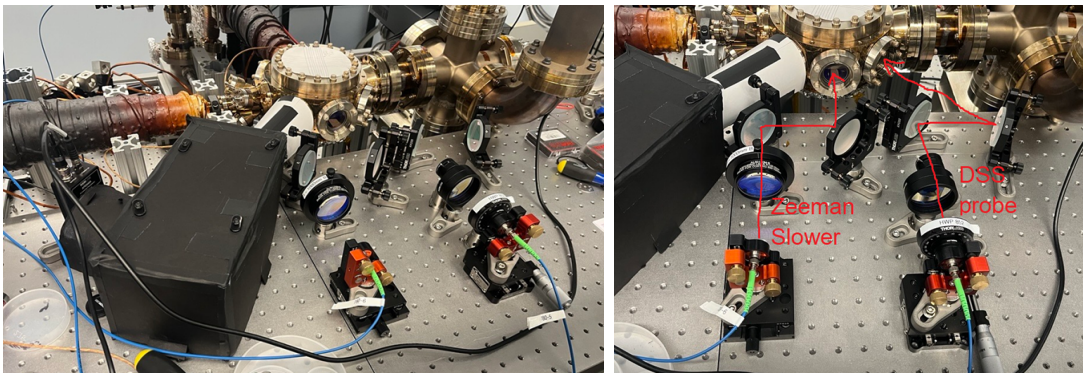


Figure 3.4: On the left, a zoomed out view of the DSFS set-up. The APD is the black box with the wires coming out of it on the right. The pick-up and focusing lenses are covered with the black box and paper tube to shut out background light. On the right, the Zeeman slower beam and probe beam paths.

The fluorescence is collected by a lens with focal length $f = 150$mm placed 150mm away from the centre of the chamber. This collimates the light, causing it to stop diverging. A $f = 100$mm and a $f = 50$mm lens spaced apart by 150mm are placed after it to increase the image's brightness and to shrink it down. Finally, a $f = 25$mm lens is placed 25mm away from an avalanche photodiode (APD) focusing the spot onto the sensor to measure all of the light collected by the first lens. An APD is used to read the signal because it's more sensitive than a regular photodiode. Additionally, an 852nm band pass filter (a filter that only lets through light within 5nm of 852nm) was used to cut down on the background and increase the APD's sensitivity to the fluorescence signal. Finally, an SRS Model SR560 low-noise preamplifier was used to amplify the signal and to filter out high-frequency noise. Noise above 1kHz was attenuated by 6dB/octave.

The signal seen was a bell curve-like distribution detuned red from the zero-velocity transition frequency. This makes sense, as the zero-velocity atoms are on resonance with the signal exactly at the transition frequency, whereas the higher velocity atoms which are more likely to make it to the chamber are resonant when the light is lower frequency, to make up for the Doppler shift. Since the atoms are traveling into the chamber, it makes sense that the atoms that make it into the chamber are moving at a positive velocity.

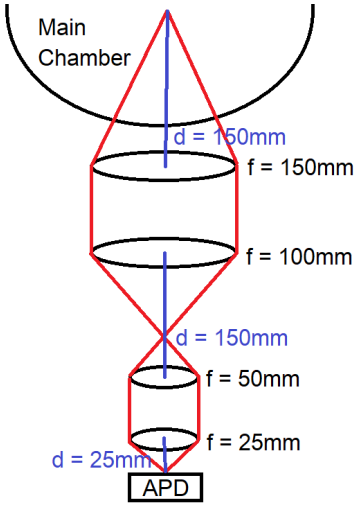


Figure 3.5: The optics imaging the centre of the main chamber onto the APD. The real version has a mirror between the $f = 100$mm and $f = 50$mm lenses, that has been left out for clarity.

The expected signal was the velocity distribution of an atomic beam which is described by $f(v) = v^3 \exp(-v^2/u^2)$ where $u = \sqrt{2k_BT/M}$.[11] The signal observed visually matches this distribution, meaning that an atomic beam was making it to the centre of the chamber. Since at any frequency, the laser is addressing three different velocity classes corresponding to three different atomic transitions, the signal we see is 3 overlapped velocity distributions spaced by the hyperfine splittings of the $P_{3/2}$ state. These velocity distributions are also compressed by $\sqrt{1/2}$ due to the probe beam being at 135 degrees to the centre axis of the atomic beam.

One difficulty encountered during DSFS was the appearance of subfeatures in the signal. The velocity distribution had a set of dips appear at higher probe intensities, that look like the dips seen in saturated-absorption spectroscopy used to lock the lasers. This caused some difficulties because a higher probe intensity leads to more light scattered by the atoms, and as a result a higher signal. This problem was most likely caused by a reflection of the probe beam from the viewport causing pumping when the laser was resonant with a transition, as described in the saturated-absorption spectroscopy section of this chapter. This was resolved by using an iris to decrease the beam size as much as possible without decreasing the signal size (which had the added benefit of decreasing the background) and by sending the beam into the chamber near the edge of the viewport. The beam was

aligned to the centre of the chamber by making sure the beam came out of the opposite edge of the opposite side viewport. This minimized the overlap of the reflection with the incoming beam. Additionally, the power was decreased to 3.1mW at this point in time, which was used throughout the rest of the measurement.

The next step after getting the Doppler-sensitive fluorescence spectroscopy optics set up and seeing a signal was to start slowing atoms with the Zeeman slower.

3.1.4 Cs Zeeman Slower

The Zeeman slower uses Doppler cooling and the Zeeman shift to slow atoms longitudinally. Since the MOT can only capture atoms below a certain velocity that is typically pretty low compared to the average velocity in the atomic beam, the slower can significantly increase the number of atoms being loaded into the MOT. The slower generally consists of a long, narrow pipe with a laser beam shining down it and a magnetic coil wrapped around it. The atoms travel through the inside of the pipe and are slowed by the light, with the magnetic field shifting their energy levels to account for their change in velocity as a function of distance travelled through the slower. This is explained more thoroughly in the laser cooling chapter. We use decreasing field slowers to maximize room for optics around the main chamber.

The Zeeman slower laser beam was sent in through a viewport on the main chamber. The main requirement for the spot was that the beam has to be 5mm wide at the nozzle, to hit every atom coming out of it. Additionally, we want the beam to converge on the nozzle to slow atoms that are coming out of the oven not perfectly collimated. This created a second constraint, that the beam had to fit through the viewport without clipping or distorting, which can happen at the edges of the viewport window. The viewport is 1.25 inches in diameter, so we decided on a $\sim$1 inch spot. The asphere used has a focal length of 100mm, and the fibre output was mounted on a translation stage in order to precisely space it from the asphere. Thorlabs 780nm PANDA core fibres have a numerical aperture of 0.12 so the spot size at the lens is given by $w = 2 \cdot 0.12 \cdot 100\text{mm} = 24\text{mm}$. The spacing chosen was such that the beam focused down to a diametre of 5mm at the position of the nozzle. Two mirrors were used to align the beam to conserve the alignment of the fibre output and asphere. The beam was aligned first by centring the beam on the viewport and overlapping the spot reflected off of the viewport with the beam source at the fibre output. It was aligned more finely by looking at the light scattered off the nozzle through a viewport with an IR viewer. When well-aligned, the DFAS signal shrinks in size, so the final alignments can be done by minimizing the DFAS signal. A quarter wave plate was also put

in and rotated to give the light left-handed circular polarization. The polarization needed to be left-handed because we wanted the atoms to be gaining positive angular momentum in every transition, because we had a magnetic field pointing towards the chamber. (We want σ^+ light, which is left-handed light when the quantization axis points opposite to the light's propagation axis. The magnetic field sets our quantization axis.) Finally, slowing was checked by setting the frequency to address the most populated velocity and looking for a dip at the addressed velocity class and a peak right beside it in the DSFS. The dip represents the loss in atoms at the velocity the laser is addressing due to interactions, and the peak is them showing up at an immediately slower velocity class due to being slowed by the laser.

The next step was to turn on the Zeeman slower's magnetic field. The Zeeman slower coil's current was set to 4A, at which it was designed to run. The expected effect is the dip growing in width to include higher velocities, and the slowed peak growing in height. With this being the case, it was time to characterize the slower. The slower laser frequency was set closer to resonance to slow a smaller-velocity group of atoms. The detuning chosen was roughly 80 MHz, which would address atoms at 68m/s at zero field. This frequency was chosen because there was a local maximum in peak height at that frequency. In characterizing the slower, the maximum slowed peak's dependence on the slowing beam intensity and the slower's magnetic field were studied. The general procedure was to set the coil current and laser power, and to move the laser's frequency to find the maximum peak height. The coil currents measured were 3A, 4A, 4.5A, and 5A. The powers looked at were 2.6mW, 3.1mW, 3.6mW, 4.1mW, and 5.1mW. The maximum peaks were all in the range of 64MHz to 80MHz in detuning.

Additionally, I characterized the slowed peak's response to the Zeeman slower's bias coil. This coil adds a uniform magnetic field in the slower leaving the gradient the same while increasing the magnitude of the field in the slower. This shifts the transition energy of the atoms, which can be compensated for by decreasing the laser's detuning. This can be useful in the future if we decide to use an increasing field slower or spin flip slower, in either of which case we would hook up the bias coils in reverse. Bias coil fields of 0A, 5A, 10A, 15A, and 20A were measured. The results showed that when increasing the field, the slowing frequency needed to be increased for the peak to stay at the same spot. This suggests that the bias coils were working as intended, as the Zeeman splitting is proportional to the magnetic field.

There were several issues I faced when working on this. The first was that there was no repump light. This may or may not have made a difference, because it's unclear whether the repump light is useful in a decreasing-field slower. The relationship between the two ground hyperfine states changes with a magnetic field as well as the excited state hyperfine

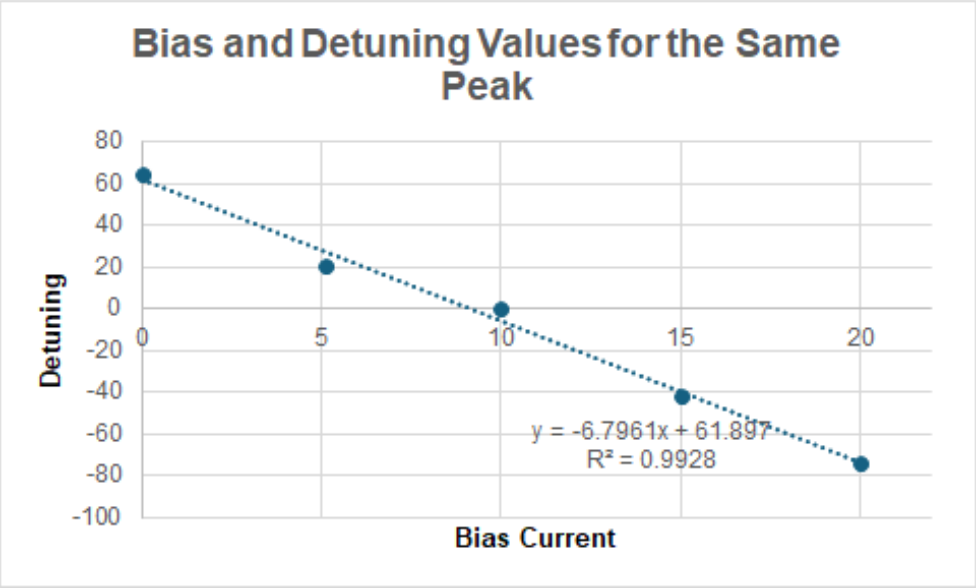


Figure 3.6: Zeeman slower detuning needed to move a peak to the same spot after applying a new bias field.

states. This means that the repump light is off-resonance inside the slowing pipe where the magnetic field is present. Additionally, Cs has a small chance of exciting into the $F' = 4$ state because the transition is many linewidths off. The second issue was that increasing intensity in the slowing beam led to the size of the signal becoming smaller. This made it hard to characterize the effectiveness of slowing, since one could not directly compare the sizes of the peaks. The solution for this was to use the ratio between the size of the slowed peak and the height of the Doppler peak. However, at higher intensities the signal size decreased to be comparable to the noise, and decreased to a size which the oscilloscope couldn't reliably measure (on the scale of mV). Third, the oscilloscope may not have been accurately displaying the measured signal. In some instances, zooming on the oscilloscope didn't change the size of the signal being displayed (eg. one could zoom out from 5mV per division to 10mV per division and the signal wouldn't change how many divisions tall it was). Another example of this was that when moving the signal up or down in height, the signal would float back to (roughly) the position on the screen where it was before moving the signal. Finally and most importantly, the slowed peak height changes over time, even while keeping the same settings. It's very unclear what's causing this, as the obvious culprits such as polarization or intensity fluctuations in the slower beam and the probe beam aren't large enough to be causing such large fluctuations in the signal.

Because of the issues mentioned in the previous paragraph, it's hard to make statements on which settings lead to the most effective slowing. Additionally, the settings on the low-noise preamplifier changed between days because for the first few days, I was getting familiar with it. The data from April 28 and 29 was not used in the conclusion, however it did influence which settings I looked at in future runs.

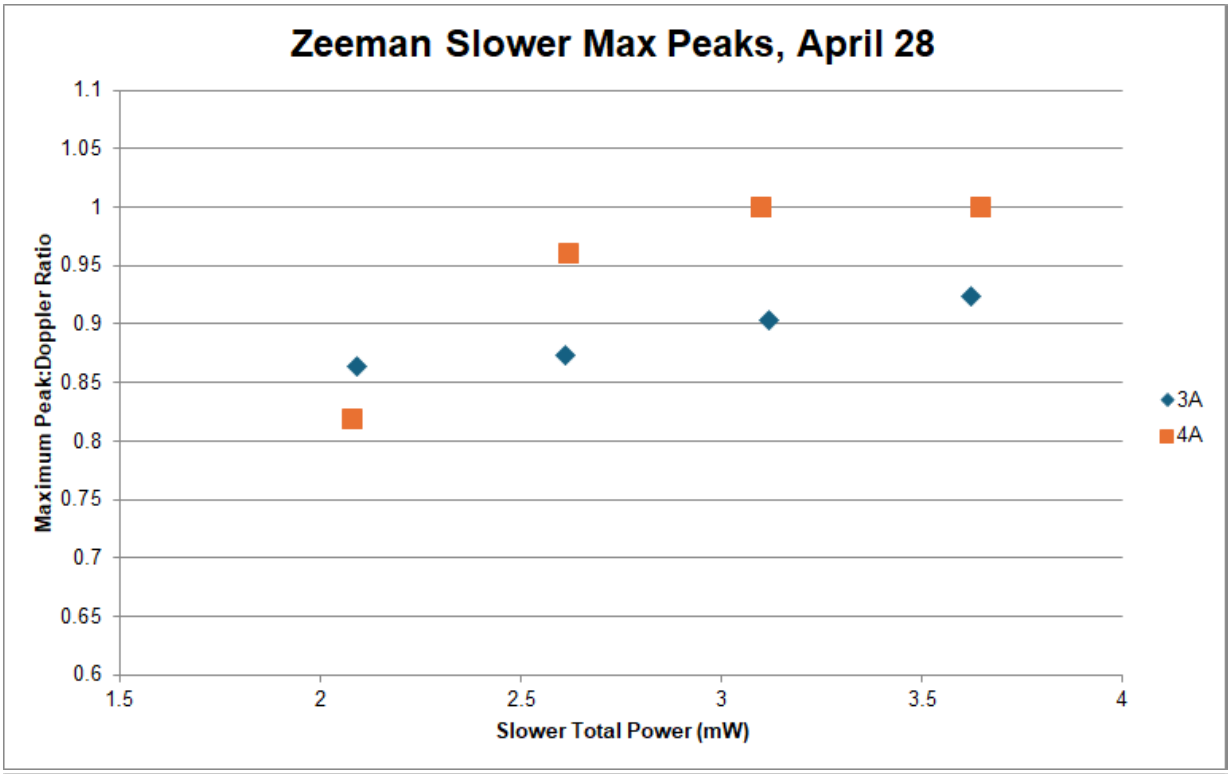
April 28 was the first day of characterizing the slower. Since the signal was getting smaller with higher intensity, I didn't take data at higher intensities. However, this graph suggests that slowing improves with increasing intensity, so I looked at higher intensities in the future. The data from April 29 supported this. Additionally, I saw that slowing was better at 4A than at 3A. The data from April 29 suggested that slowing at 4A was also better than slowing at 5A, however the difference between them is much closer than between 3A and 4A. This led me to look at 4.5A on April 30 and May 1. The gain on the low-noise pre-amp was increased at 4.1mW and 4.6mW on the 5A run on April 29.

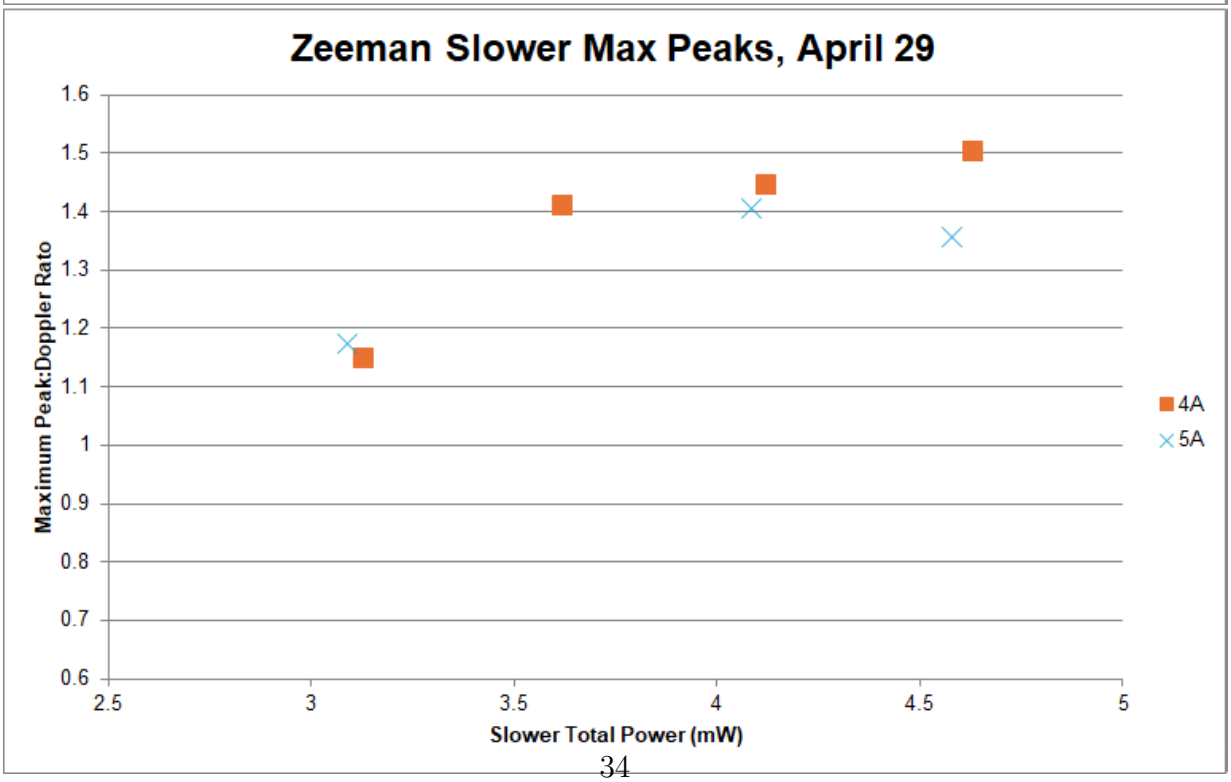
On April 30, no low noise pre-amp was used, so the signal was small (on the scale of 0.5-1mV in height). As a result, these results might be unreliable. It's unclear how inaccurate the oscilloscope is at such small signals. This data suggests that 4A is the best current to run through the Zeeman Slower, since it had the highest peaks at every power except 3.6mW. The trend with intensity is still that higher intensity leads to be better slowing, although this trend is much less clear than in the past two measurements.

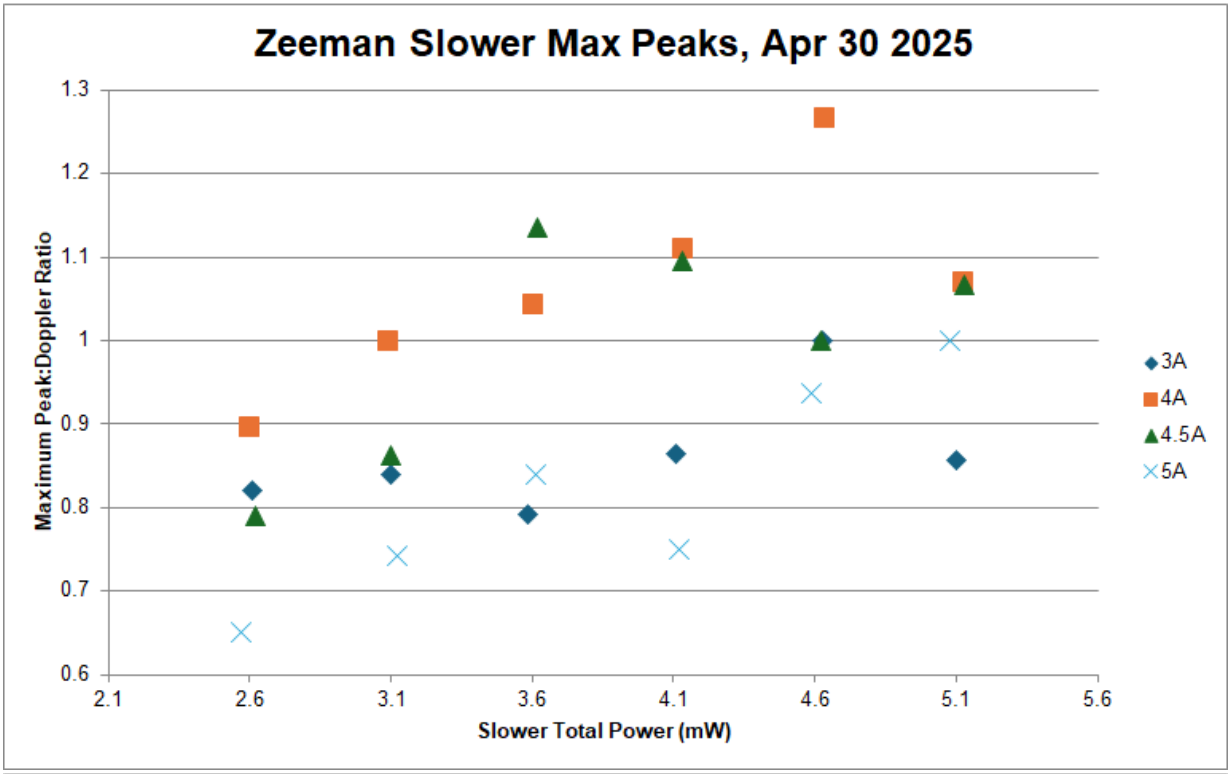
On May 1, we used the low-noise preamp again. The low pass filter setting used was 6dB/oct attenuation above 1kHz and 1×10^3 gain. On this day, slowing at 4.5A gave the highest peak at every power except 4.6mW. The trend of increasing power leading to better slowing also held true.

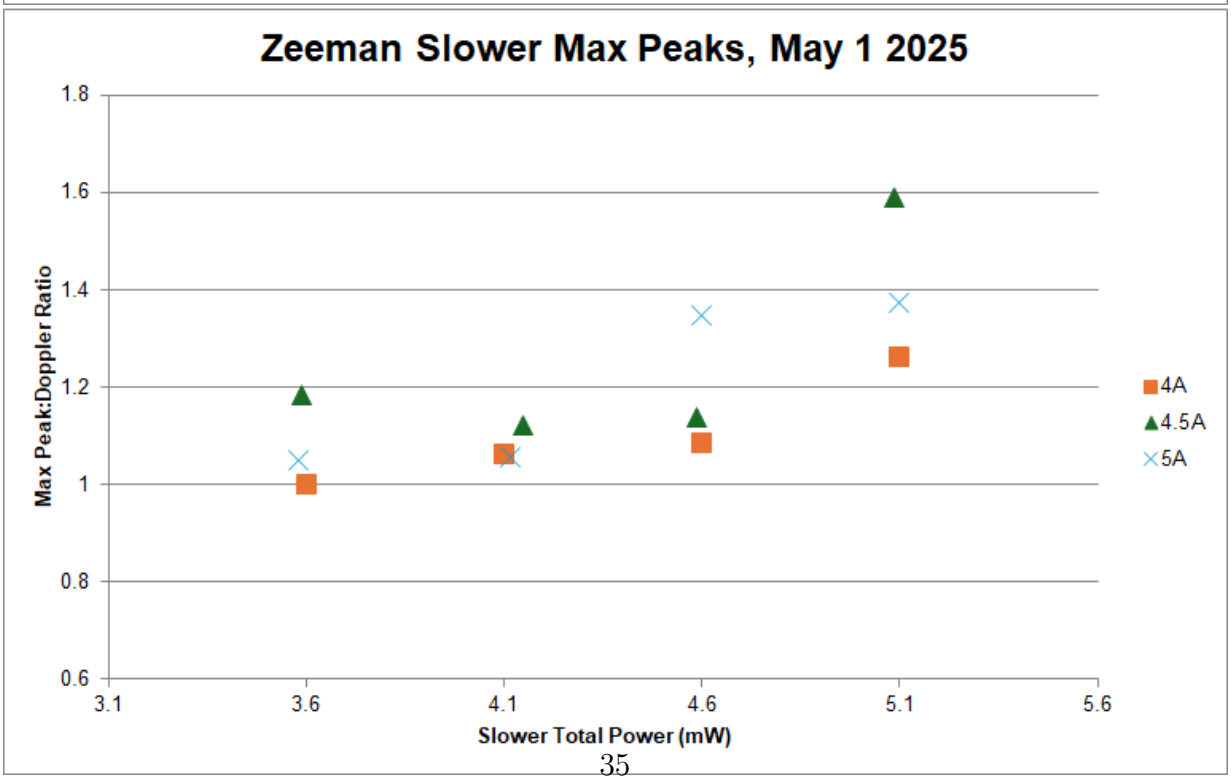
Some general trends can be seen. In varying the intensities, increasing intensity leads to more effective slowing. Despite the fluctuating signal height from day to day, the most effective slower fields were 4A and 4.5A, which makes sense because the slower was designed to work at 4A.

A more comprehensive characterization of the Zeeman slower will be made when optimizing the MOT loading rate, which is the quantity that we are trying to maximize in using a Zeeman Slower. This will include slowing at higher intensities, since the repump laser will be available, and carefully characterizing the fluctuation in the number of slowed atoms over time. When doing this procedure with the Zeeman slower for lithium, we started taking multiple measurements of the same data and doing them in random order to be able to get a more accurate sense of the data and to measure the average fluctuation of the data taken. This will also be done with the MOT loading.


Zeeman Slower Max Peaks, April 28
Maximum Peak:Doppler Ratio
1.1
1.05
1
0.95
0.9
0.85
0.8
0.75
0.7
0.65
0.6
1.5
2
2.5
3
3.5
4
Slower Total Power (mW)
3A
4A


Zeeman Slower Max Peaks, April 29
Maximum Peak:Doppler Rato
1.6
1.5
1.4
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
2.5
3
3.5
4
4.5
5
Slower Total Power (mW)
4A
5A


Zeeman Slower Max Peaks, Apr 30 2025
Maximum Peak:Doppler Ratio
1.3
1.2
1.1
1
0.9
0.8
0.7
0.6
2.1
2.6
3.1
3.6
4.1
4.6
5.1
5.6
Slower Total Power (mW)
3A
4A
4.5A
5A


Zeeman Slower Max Peaks, May 1 2025
Max Peak:Doppler Ratio
1.8
1.6
1.4
1.2
1
0.8
0.6
3.1
3.6
4.1
4.6
5.1
5.6
Slower Total Power (mW)
4A
4.5A
5A

3.2 Controlling & Stabilizing Lasers

3.2.1 Cs Laser Locking

It's important to make sure that the lasers we use keep their frequencies at one specific point to avoid systematic error during experiment. This is done with an optics setup that uses a vapor cell containing a sample of the alkali to which we want to lock the laser. I realigned one of our laser-locking optics board which already contained all of the optics needed for this, set up some of the electronics and locked the laser with the board. Additionally, I instructed and assisted an undergraduate student with another identical laser-locking board.

The first step is understanding saturated absorption spectroscopy, which allows one to pinpoint where a laser's frequency is on-resonance with a transition. We send a laser beam with sweeping frequency through a vapor cell. Measuring the absorption of the laser beam as a function of frequency (and in our case, as a function of time, due to a linear sweep) will give a thermal distribution due to addressing different velocities at different frequencies. The frequency at which the v = 0m/s atoms are resonant is the true transition frequency. The signal is often made up of an overlap of multiple transitions. We send in a second beam from the same laser, overlapped with the first, with greater power from the opposite direction. In any given atom's rest frame, the two laser beams have a frequency difference $2kv$, where k is the laser's angular wavenumber and v is the atom's velocity. When the laser's frequency is on resonance with the atomic transition, the probe and the pump are both exciting the same atoms with a very small velocity. The pump beam pumps all the atoms out of a state that the probe can address, leading to a sharp peak due to the atoms not scattering the probe light. This happens at each transition, showing the frequency of each transition.

We use a proportional-integral-derivative controller, or simply a PID controller, to lock our lasers. PID controllers are devices that take some sort of signal, typically corresponding to a measured physical quantity, and returns the signal with some operations done on it. Those operations are in the name – the returned signal is a sum of the original signal which is multiplied by a variable factor, labelled P for proportional, an integral of the signal over some time, multiplied by variable factor I, and the time derivative of the signal, multiplied by some factor D. This outputted signal can be plugged into a device that can control or effect the aforementioned physical quantity to stabilize the signal at some desired value. Some examples of how we use PIDs in the lab are stabilizing the ambient temperature in the lab, maintaining constant intensities in laser beams where critical, and, as mentioned earlier, locking our lasers.

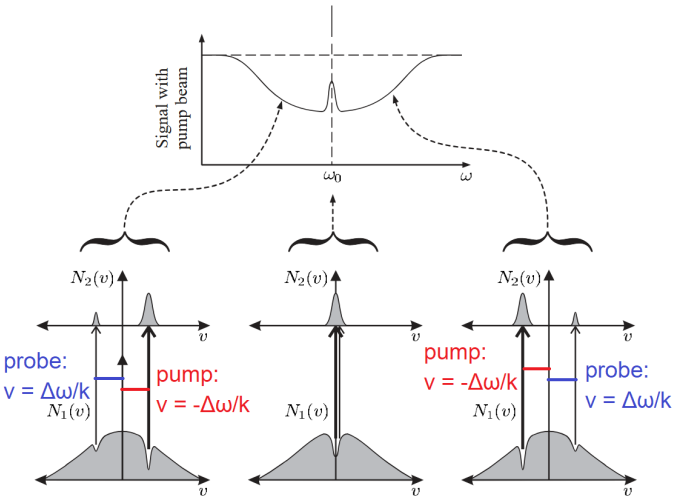


Figure 3.7: The top shows the saturated absorption signal, the atoms' absorption of the probe beam, as a function of frequency. The bottom shows the pump and probe beam's behaviour as a function of frequency. On the left and right, the probe beam and pump beam are addressing different velocity classes due to coming in from opposite directions. In the middle, they're both addressing the v=0 velocity class. At this frequency, the pump beam is pumping atoms out of the state that the probe beam can address, leading to a decrease in absorption. [11]

However, it's difficult to lock a laser simply to the centre of a saturated absorption spectroscopy peak. No matter which way the laser drifts, the PID controller will see a decreasing signal, meaning that the response can only correct in one direction. It's much easier to do so to a signal with a linear slope at the transition frequency. Frequency modulation spectroscopy (FMS) does exactly this. In frequency modulation spectroscopy, frequency sidebands are added onto the probe beam, of frequency $\omega_s = \omega_L \pm \omega_m$. (The sidebands are added with an electro-optical modulator, or EOM, which will be explained later.) These sidebands have a frequency difference ω_m from the main carrier's frequency ω_L on the order of a natural linewidth ($2\pi \cdot 5.6$MHz) and a small intensity compared to the main carrier. The AC and DC signals are separated with a bias tee and the DC signal gives a saturated-absorption spectroscopy signal, as described in the last paragraph. The error signal, which is used to lock the laser, comes from the AC signal. [4]

The source of the AC signal is the beat notes generated by the sidebands beating against the main carrier. Since the sidebands are detuned from the main carrier by ω_m, this is the frequency of the generated beat notes. Finally, the beat notes due to each carrier

have a relative π phase with each other. Since they have the same frequency and opposite phase, the total AC signal has a frequency ω_m, and amplitude equal to the difference in amplitude between the two sidebands after they pass through the cloud of atoms. The amplitude of each beat note depends on which sideband experiences more absorption. In other words, the total amplitude depends on which one is closer in frequency to the transition – if the laser is on resonance with the transition, both sidebands experience equal absorption, and the AC signal is zero. If the laser drifts (or sweeps) down in frequency, the higher sideband experience more absorption than the lower, leading to a signal with the amplitude proportional to their difference in absorption. When $\omega_L > \omega_0$, the lower sideband will experience more absorption, and the signal will have amplitude proportional to their difference but with a π phase shift. This AC signal is amplified and turned into a DC error signal through demodulation.

Our error signal already exists in this picture of the AC signal, in the form of its amplitude. Our AC signal can be seen as the product of a RF wave with frequency ω_m and the error signal, which varies as a function of ω_L. The phase of the AC signal switching by π can be interpreted as a negative amplitude, giving us a negative error signal in that region. After amplification with a RF amplifier, the AC signal is demodulated with an RF mixer. A mixer is a device that modulates or demodulates some function by some radio frequency ω – it adds or subtracts the provided frequency to or from the function. For demodulating the error signal, the mixer is provided the AC error signal and an RF signal of the same frequency as the EOM, with a phase adjustment to account for the laser beam's change in phase as it travels through space and the phase attenuation as the signal travels through the electronics to reach the mixer. The signal pulled out is the DC error signal, or the amplitude of the AC error signal. Figure 3.8 shows such a signal found with modulation transfer spectroscopy, but an FMS error signal looks very similar. This signal is used to lock the laser's frequency, which is done by using a PID controller to servo the laser cavity length.

We add frequency sidebands onto our laser beam by sending it through an EOM, or electro-optical modulator. An EOM is made up of a crystal with an oscillating electric field applied to it, that changes the crystal's index of refraction. The oscillating field is created using an LC circuit using the EOM field plates' capacitance and an external inductor. Using the desired frequency, on the scale of 6MHz, and the EOM's capacitance, 14pF, and the formula for the resonant frequency of an LC circuit, $f_0 = \frac{1}{2\pi\sqrt{LC}}$, a required inductance of 50µH was found. An inductor with the closest inductance available is used. The frequency of the sidebands depends on the RF frequency being driven through the EOM (which is generally chosen to be the resonant frequency) and the intensity depends on the power of the RF signal driven through the EOM.

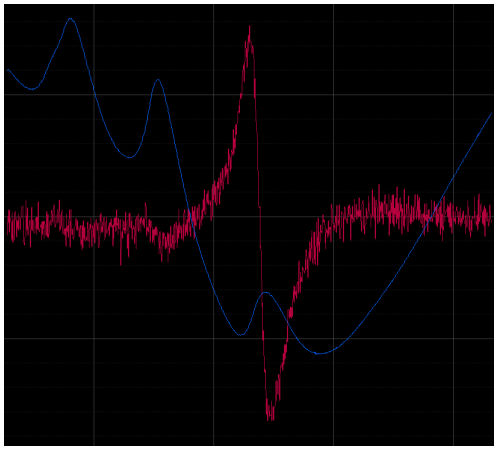


Figure 3.8: Our error signal, found through modulation transfer spectroscopy, which we use to lock our laser. The error signal is overlayed with the Doppler-sensitive spectrum to identify which peak is being locked to.

A second technique that also provides a linear error signal is modulation transfer spectroscopy (MTS). The advantages of this technique of FM spectoscopy are that the signal has less noise and a baseline that is always zero, and that only cycling transitions give a signal with this technique. On the other hand, MTS requires more precise alignment and intensity splitting. With this technique, the pump beam has the sidebands, rather than the probe, and the intensities in each arm are the same. The AC error signal appears through 4-wave mixing (a non-linear interaction) between the pump and probe beams in the atoms. A clear and thorough description of this technique can be found in Daniel McCarron's PhD thesis. [22]

The layout of the optics setup for FMS and MTS can be seen in Figure 3.9. The light comes onto the board through a fibre and is collimated. The beam passes through a half wave plate and a polarizing beamsplitter (PBS) cube to split into the probe and the pump beam. The intensity in each beam can be changed by changing the polarization going into the cube with the half wave plate. Depending on if FMS or MTS is done, different probe:pump intensity ratios are used. FMS uses a ratio of about 5:1, with a total intensity of about saturation intensity. MTS uses a ratio closer to 1:1. The probe beam (the pump beam in MTS) travels through an EOM, is reflected with another PBS cube, and passes through a Cs vapor cell. The pump beam (probe beam) is reflected through a polarizing beamsplitter on the other side of the vapor cell and passes through the vapor cell, overlapped with the probe beam. The probe beam is picked off for FMS by the polarizing beamsplitter cube that reflected the pump beam into the vapor cell, and the pump beam

(MTS probe beam) is picked off for MTS by the PBS cube that reflected the probe beam (MTS pump beam) into the vapor cell. Each beam is coupled into a fibre. The signal is read with a photodiode, split into DC and AC components, demodulated, and used to lock the laser.

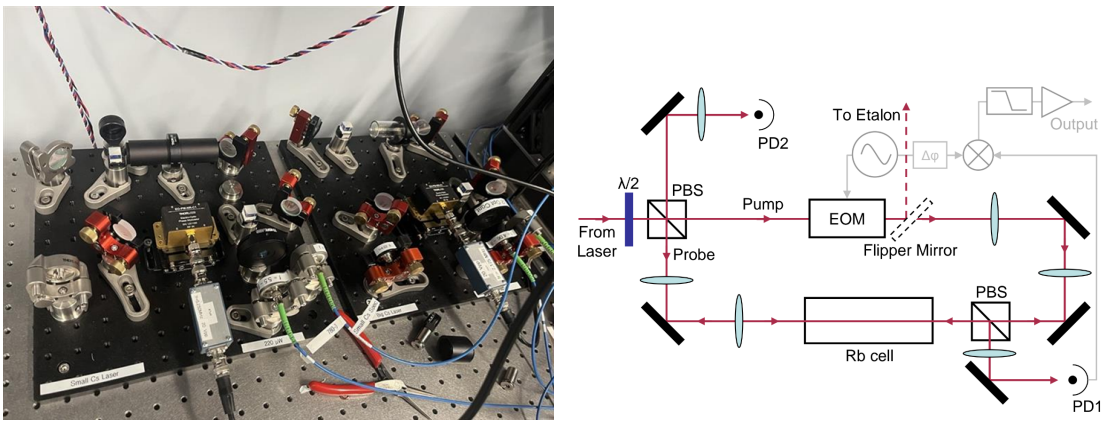


Figure 3.9: On the left, the two Cs spectroscopy boards in the lab. On the right, a simplified diagram of the spectroscopy set-up from Daniel McCarron's thesis. [22] PD1 here is used to observe MTS, while PD2 is used to observe FMS. The flipper mirror is not used in our design.

3.2.2 AOM Boards

In the lab, we use AOM boards to shift the frequency of a laser beam by a small amount as needed. An AOM is an acousto-optical modulator. It's made of a crystal and a piezo-electric transducer that turns an RF signal into a sound wave that propagates through the crystal. The sound wave acts as a moving diffraction grating that adds the sound wave's frequency to the diffracted light. The AOM has a frequency at which the diffraction is most efficient, and diffraction efficiency drops off as frequency deviates further from this centre value.

AOM boards are optics set-ups that take light traveling through a fibre optic, passing it through an AOM, and coupling it into a second fibre. They are set up to pass the light through the AOM twice in what is known as a cat's eye configuration, which allows the changing of frequencies in the AOM without affecting the coupling into the fibre. This shifts the frequency by twice the amount the AOM is set to, which can be useful when a

larger frequency shift is needed. However, in some cases, we use single-pass setups that need to be recoupled whenever the frequency is changed.

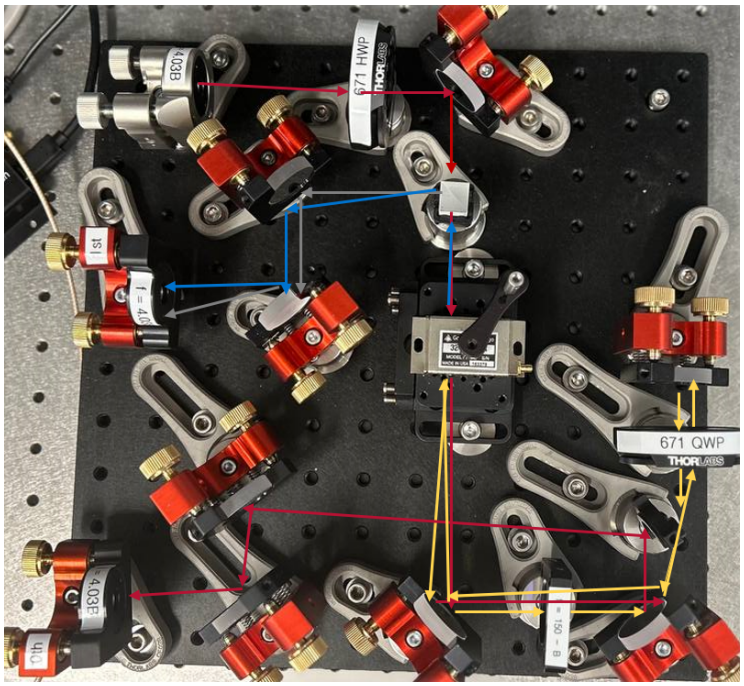


Figure 3.10: The typical AOM board. The laser comes onto the board at the top left. The red arrows represent the undiffracted light, the yellow arrows represent light that has been diffracted once, and the blue arrows represent the light that has been diffracted a second time.

I worked on ensuring that all the AOM boards needed for the Cs optics in the experiment were working. One AOM is needed to shift the frequency at which one of the two lasers are locking. This allows a second AOM board to bring the frequency back up to where it's needed, which is useful for two reasons. The first reason is that it allows for a smaller frequency difference between the transition frequency and any given experiment frequency than would be allowed by the AOM's frequency range. (Eg. A desired MOT detuning of 20 MHz, which is smaller than the 140MHz that is the lower bound at which our AOMs can diffract with sufficient efficiency.) The second is that it allows for easy shut-off of the experiment beam by simply cutting RF signal to the AOM. The reason this is only used with one laser is due to the frequencies the lasers lock on. The first laser, hereafter called the cycling laser, will be locked to the F=4 to F'=5 transition. Since this is the transition on which slowing happens, bringing the laser off-resonance with an AOM board is needed.

Purpose	Source laser	AOM freq. (MHz)	Final detuning (MHz)
Spec. detuning	Slowing	210	$2 \cdot 210 = 420$
MOT slowing	Slowing	200	$-420 + 2 \cdot 200 = -20$
Zeeman slowing	Slowing	170	$-420 + 2 \cdot 170 = -80$
Imaging	Slowing	210	$-420 + 2 \cdot 210 = 0$
MOT repump	Repump	166	$-352.5 + 2 \cdot 166 = -20.5$
Zeeman repump	Repump	143.5	$-352.5 + 2 \cdot 143.5 = -65.5$

Table 3.1: List of AOM boards. The -420MHz shift on the slowing laser AOMs comes from the shift from the spectroscopy board. The -352.5MHz shift on the repump laser AOMs comes from the fact that the laser is locked to a transition that's 352.5MHz lower than the F=3 to F' = 4 transition.

The second laser, which we will call the repump laser, locks to the F=3 to F'=2 transition, which is 352MHz off from the F=3 to F'=4 transition, which is what that laser will be driving. In this case, an AOM board isn't needed since the laser is already detuned.

The slowing laser powers the Zeeman slowing beam, the MOT slowing beam, and the imaging beam. The repump laser powers the Zeeman slower repump beam, the MOT repump beam, and the blue MOT slowing and repump beams. The MOT slowing and repump beams are detuned by 20MHz from the F=4 to F'=5 and the F=3 to F'=4 transitions, respectively. The Zeeman slower slowing and repump beams are detuned from the same transitions by 80MHz. The imaging beam is on resonance with the F=4 to F'=5 transition.

We use many AOM boards that can take up a lot of space, and optics table space is a precious commodity. That is why I built an optics board condo that acts as a sliding drawer for optics boards. This allows optics boards to be stored as if stacked vertically, taking up much less table space per board. The limit on condo height was set by the electronics shelves above the optics table, and so 4 optics boards could fit vertically. At the same time, it's important for the condo to be vibration-proof to keep optics aligned over time and to avoid error due to vibration during experiment. Two different designs were tried, with the second being the final one.

I built the first condo with the side of the condo facing out of the table, and the sliding front of the condo facing sideways, into the table. The logic behind this was that this would save on space, as two condos facing towards themselves could share the room that's

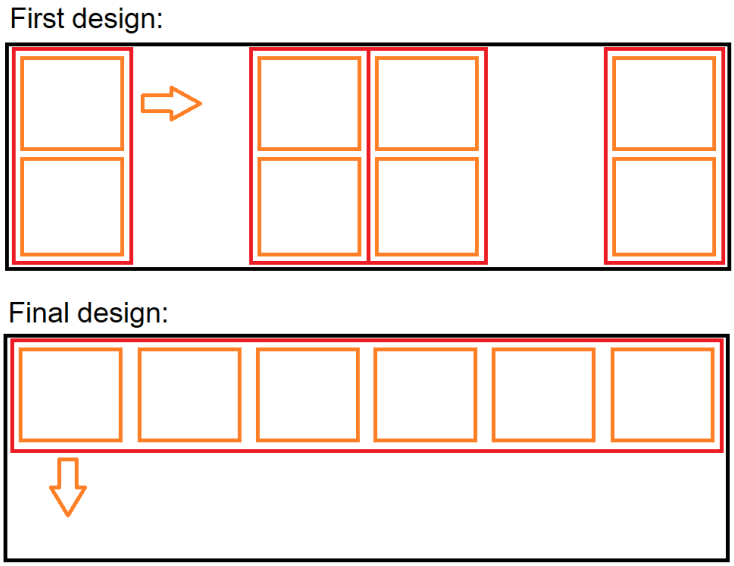


Figure 3.11: The designs for the first and second optics condos. The black rectangle is the optics table, the red rectangles delineate the condos, and the orange squares show the optics boards. The orange arrows indicate which way the boards slide out. It is clear to see that in the initial design, 8 "plots" can be filled with optics boards instead of 6 in the new one.

required to allow the drawers to slide out. In doing so, it would be possible to fit 32 optics boards on the table. The condo outer structure is built mostly out of a heavy aluminum slotted frame scaffolding material with a 1"×1" cross-section, also known as 1"×1" 80-20. The AOM drawers were made out of 1"×1" and 1"×2" 80-20. Heavy material was chosen for these tasks because a structure's resonant frequency is inversely proportional to its mass. The sliding rails were ordered from McMaster-Carr. Acrylic sheets were used as a covering for the top, sides, and rear of the condos to keep dust out. There were several issues with this design. First was that the sideways-sliding drawers made it uncomfortable to work with the AOM boards that were stored in it, even with no other condo on the other side of the gap. The second issue was that the sliding rails used were unstable. The AOM drawers were wobbly even when they were tightly bolted to the sliding frame and the sliding rails were bolted to the outer condo structure. Additionally, each sliding rail was attached to the outer structure with only a single screw, with the rest of the support coming from brackets screwed onto the frame that didn't directly connect to it. This lead to instability in the horizontal level of the rails over time.

The second condo design was built with the intention of fixing these mistakes. The new

design had the drawers sliding out towards the person, with all of the condos in a single row at the back of the optics table. This allows for 24 optics boards to be stored on the table. While this design is less efficient, it's a lot more comfortable to work with, and it still houses more AOM boards than would be possible without the condo. Additionally. some new sliding rails were ordered that were shorter, leading to less wobble within the rails, and they were bolted onto the condo structure with two screws to keep the horizontal leveling more stable. Finally, all 80-20 used in this condo design was 2"×1" 80-20. Using the thicker pieces meant a heavier structure, with even less sensitivity to high frequency vibration. This also came with a second advantage of more precise right angle connections between both two different pieces of 80-20 and of 80-20 to the table, due to the longer lever arm. All of these improvements, along with the fact that in this design all of the condo are connected to one another, lead to a much more stable storage solution for optics boards than in the previous design.

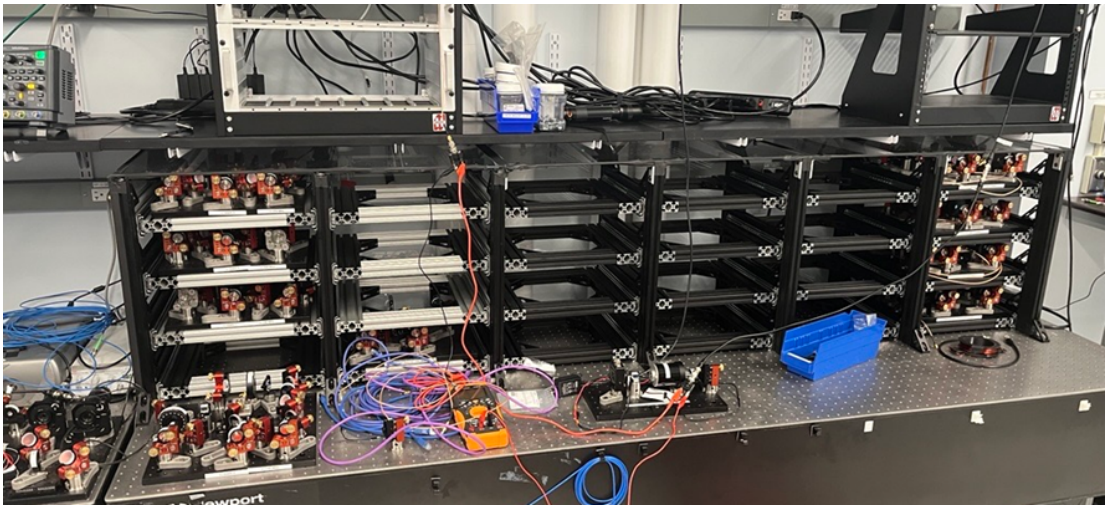


Figure 3.12: The finished optics condo, already containing some AOM boards.

3.2.3 Optics Shutters

Shutters are the tool that allow a beam to be blocked mechanically. While beams can be turned on/off with AOMs, changing the RF power going through an AOM can lead to thermal effects affecting the experiment, so sometimes it's better to use a shutter. Optics shutters are available commercially, but they take up a lot of space, they cause vibration

when opening and closing, and they cost a lot. Additionally, they sometimes need their own specialized controllers to run. We can instead use digital camera shutters that are much smaller and available for cheap on Ebay.

The initial design was provided by Prof. Isaac Chuang. This design had a 3D printed mount and a circuit board with a Raspberry Pi microcontroller and a USB-C port that allowed the shutter to be controlled on a computer using Python. However, the circuit board added bulk to the shutter and running the shutters on Python did not offer any advantage over using analog signals to control a power supply. These shutters require 5V and 60mA to stay open. Initially, we were planning on using Acopian 5V power supplies. However, our 5V digital signal generator provides enough current to power the shutters. Using the digital signal to power the shutters means less setup is required for the shutters and less shelf space is taken up by the power supplies. The design of the 3D mount was also changed by me and then by a different grad student to reflect the removal of the circuit board and change in how it's powered.

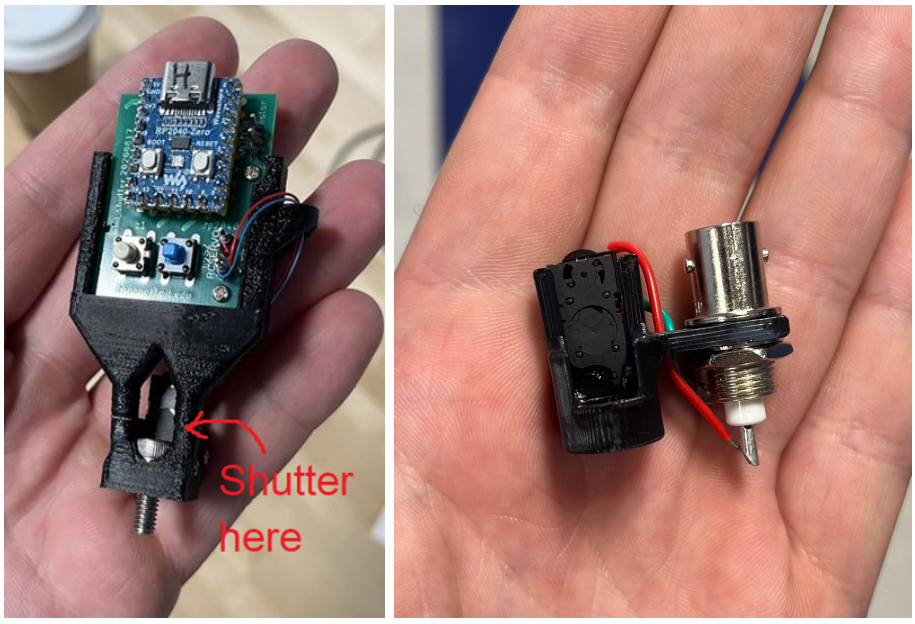


Figure 3.13: On the left, Prof. Isaac Chuang's design for the optics shutter, with the mounted circuit board. On the right, the final design with a BNC port for triggering.

Since the specifications of these shutters aren't available online, some characterization was done. First, the time delay between the trigger signal and the shutter opening was measured and the opening time were measured. These measurements were done by placing

the closed shutter in the path of a laser beam shining on a photodiode and opening the shutter with the trigger signal and photodiode signal displayed on an oscilloscope. The time it took the shutter to open and the time delay between the trigger signal and the onset of opening were measured. The opening time was found to be 1.32ms and the delay was found to be 1.72ms. The dead time, which is the time after closing that the shutter can't open again, was also measured. This was done by sending a square wave signal to the shutter and decreasing the period until triggers were missed. The dead time was defined as half the period and was determined to be 4 ms. The closing time and time delay between the end of the trigger signal and the shutter closing weren't measured because turning off the AOM will always be faster and the heating effects we are trying to avoid occur after the AOM turns back on. This quantity will be characterized in the future if needed.

Finally, the maximum laser intensity the shutters can handle was measured. A fibre coupler with an $f = 4.03$mm asphere was used, with the beam collimated. This creates a spot with a 0.048cm radius, leading to an area of 0.0073cm^2 and 671nm light was used. The shutter was placed in the beam path. A power was chosen and the beam was unblocked, and the time it took for a burning smell to appear was recorded. Before moving to a new power, the beam was blocked for the shutter to cool back down. At 35mW, no scent was noticeable after over 3 minutes. At 40mW, a scent became noticeable after 30s, and at 50mW the smell appeared instantly. Strictly based on this information, it looks safe to use at intensities smaller than 4.8W/cm^2, or 35 mW when using a collimated beam from a f=4.03mm asphere.

Chapter 4

Fr at TRIUMF

The francium part of the FrAg experiment takes place at TRIUMF. We want to use francium because of ^{223}Fr's octupole-deformed nucleus, which has a much higher Schiff moment than other alkali atoms.[9] Because all francium isotopes have unstable nuclei, with the longest-lived having a 22 minute half-life, it needs to be formed through decay of some other radioactive material to have enough to use in experiment. TRIUMF produces francium and already has a facility that laser cools francium, so it's the ideal location for this experiment. In addition to our FrPA experiment, the francium trapping lab is currently home to the FrPNC (francium parity non-conservation) experiment. This experiment looks for parity violation by measuring a forbidden M1 (the lowest energy magnetic dipole transition) transition mediated by the weak force. [35]

Francium's properties as an atom that can be laser-cooled and used in ultracold atoms experiments have not been catalogued to the extent that other alkalis have. This is due to how hard it is to acquire. Additionally, sub-Doppler cooling of francium has never been reported, nor has it been loaded into an ODT. As mentioned in the introduction, in order to evaporatively cool ^{223}Fr, we need to find its scattering length. TRIUMF produces ^{223}Fr in much lower numbers than other francium isotopes, so it's easier to find its scattering length through calculation. The goal is to find the scattering lengths of ^{211}Fr and ^{221}Fr and use those to more accurately calculate ^{223}Fr's scattering length. We plan to find their scattering lengths by doing photoassociation spectroscopy. The goal of my time at TRIUMF was to do photoassociation spectroscopy in ^{211}Fr.

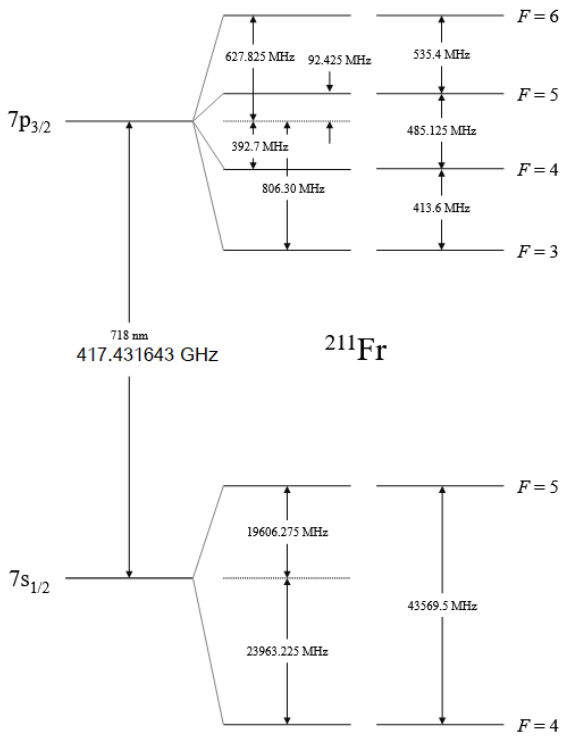


Figure 4.1: ^{211}Francium's D2 line, at 417.431643THz.[31][15]

4.1 Experimental Apparatus

The experimental procedure at the francium trapping lab deviates the most from the configuration described in the introduction. Because of its short half-life (the longest lived isotope has a half-life of 22 minutes), francium does not exist in solid form, so francium atoms can't be obtained by heating up a solid chunk of francium in an oven. At TRIUMF, francium is created by shooting a proton beam coming from a particle accelerator at a uranium carbide target. The francium is ionized and guided to the francium trapping lab using ion optics. Upon entering the vacuum chamber, the francium beam hits and sticks to a piece of zirconium foil that neutralizes the francium ions. The zirconium neutralizes the francium because it has a smaller work function. After a set loading time, the zirconium foil flips up to face into the francium trapping chamber and heats up quickly to emit the now-neutral francium atoms. It flips back down to catch more francium while the last

batch of francium is being experimented on. The presence and intensity of the francium beam is detected with a Faraday cup whenever the foil flips up, which is how issues with the francium beam are diagnosed.

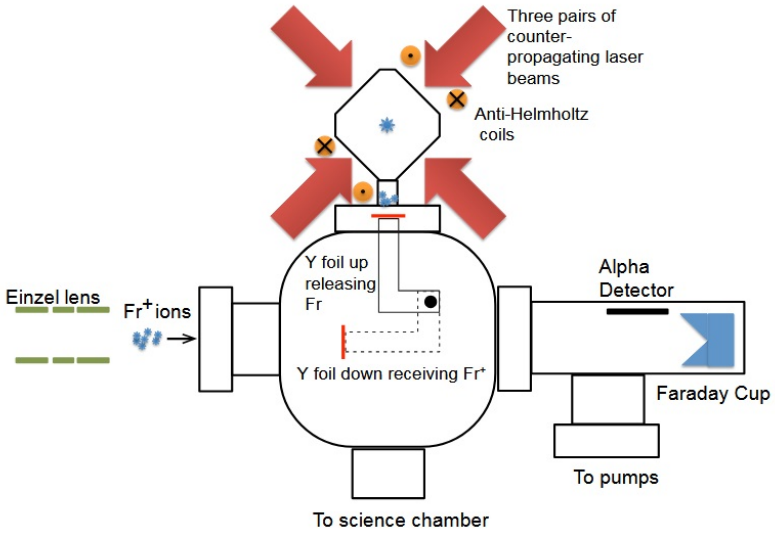


Figure 4.2: The francium ion-catching scheme. The foil, which used to be made of yttrium but is now made of zirconium, catches the francium ions and neutralizes them. The foil flips up to release the francium into the capture chamber, where it's trapped and cooled. While this happens, the Faraday cup detects the francium beam. This signal is used as a diagnostic tool for the francium beam. [35]

Since the amount of francium that can be delivered in this way is small compared to what is desired in ultracold atoms experiments, it's important to maximize the francium that is trapped. The francium is released into a trapping chamber, which contains a MOT and is coated with a polysilicon (silane) that prevents francium from sticking to it. The silane coating keeps the francium bouncing around in the chamber for longer, which makes it more likely to be trapped in the MOT. Since the silane coating can only stick to glass, the whole trapping chamber is made of glass.

After the capture MOT is loaded, the francium is transferred into the science chamber with the help of an unbalanced pushing laser beam. It's captured by a second MOT from which the experiment takes place. The science chamber has a lot more optical access for experiment lasers and imaging optics. The most notable of these for our purpose is the photoassociation beam. The imaging optics include a camera that can show the trapped

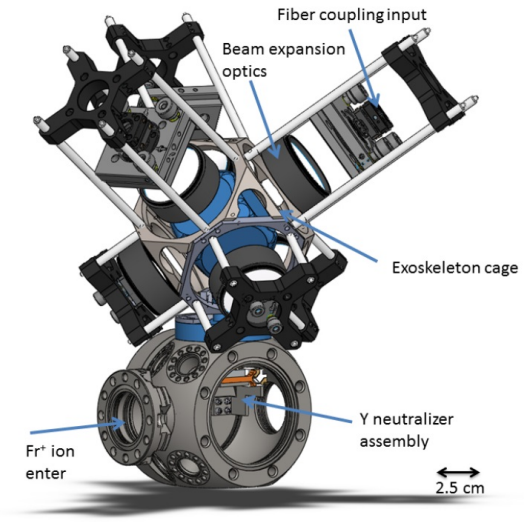


Figure 4.3: The Solidworks model of the foil and capture MOT. The pure glass cell is visible in blue within the exoskeleton cage, which holds the MOT optics in place. [35]

atoms in space and a silicon photomultiplier tube that has single-photon precision, that is used to measure the total brightness of the MOT (after accounting for the fraction of the MOT fluorescence that doesn't make it to the photomultiplier tube).

The experimental apparatus also has two rubidium dispensers, one near the capture chamber and one in the science chamber. Rubidium is used as a stand-in for francium when we don't have a francium beam because it's the alkali with the closest atomic transitions to francium's. The same optics are used for trapping rubidium as for francium. This allows work on parts of the experiment, like the MOT, that otherwise couldn't be worked on without atoms present. Having a dispenser in the capture chamber allows for replication of the shot loaded MOT that is seen during francium beam time. The dispenser in the science chamber allows for a continuously loaded MOT which is easier to work with than the shot loaded MOT.

The light for the MOT in the capture and science chambers came to the experiment chamber from a separate chamber through two sets of 3 fibres. Since there is no need for both MOTs to operate at the same time, they use the same light between them, with one beam being deflected between the two. The slowing and the repump light for each MOT is coupled into a fibre optic, which carries the light into a fibre-splitter. These fibresplitters take the light coming in through the two fibres and couple it out into three new fibres

through evanescent coupling. The trapping light should in principle be evenly distributed over each arm, whereas the repump comes in through one of the arms. However,

4.2 Photoassociation Spectroscopy

Evaporatively cooling francium is important for making FrAg molecules and achieve our long-term science goal of measuring a nuclear EDM. Evaporative cooling requires holding atoms in a trap for relatively long periods of time. A trap's hold time depends largely on the atoms' scattering properties, which can be condensed into the scattering length. A negative scattering length, representing attractive interactions, results in an unsustainably increasing density leading to two-body (and even three-body) loss, which scales with density squared (cubed). On the other hand, a scattering length that's positive but large compared to the length scale governing three-body loss (on the scale of 500 Bohr radii, or roughly 26nm) means simply that the trapped atoms scatter each other out of the trap quickly. The ideal scattering length is one that's positive, but small. Atoms with unfavourable scattering properties can be and have been evaporatively cooled through techniques that change those scattering properties, but the scattering properties need to be known first. We are using a technique called photoassociation spectroscopy to determine francium's scattering length for this purpose.

Photoassociation (PA) is the process in which a photon excites two atoms into an excited molecular state with a high vibrational number. In essence, the system of two colliding atoms is being driven into a higher energy state, as opposed to a single atom, which itself is a system of protons, neutrons, and electrons. A system of two atoms experiences an interaction potential, also known as a molecular potential. The long-distance part of the potential is mostly dominated by attractive van der Waals interactions, whereas at short-range, the atoms are strongly repelled from each other through the Pauli exclusion principle compressing the atoms' electron clouds, leading to Coulomb repulsion. At long range, each interaction potential asymptotes to both atoms occupying single-atom states – both atoms in the $S_{1/2}$ state, one in the $S_{1/2}$ and the other in the $P_{1/2}$, one in the $P_{3/2}$, both in the $P_{1/2}$, and so on. With decreasing interatomic distance, the potential increases infinitely, and between the two extremes there's a potential well with a minimum. The eigenstates of the Hamiltonian containing this potential give us the bound molecular states. As a sidenote, we can extract the scattering length by finding the scattering wavefunction with the ground state (both atoms in the $S_{1/2}$ state) potential at total energy $E = 0$. The reason we cannot simply do this for ^{223}Fr is because the potential at short range is not known precisely enough.

PA is done by shining a laser beam into a trapped cloud of atoms and targeting a specific molecular vibrational state with the laser's frequency. The transition rate is proportional to the position overlap of the two states. Since the interatomic distance is much larger than the deepest part of the interaction potential, we want to target molecular states which have a significant component of their wavefunction at large distance. This is the case with eigenstates near dissociation in energy, where the width of the potential is very large. As a rough heuristic, the amplitude of the wavefunction is proportional to $\frac{1}{\sqrt{p(r)}} \propto \frac{1}{\sqrt[4]{|E-V(r)|}}$, where $p(r)$ is radial momentum, E is the total energy of the system, and $V(r)$ is the molecular potential at position r. The wavefunction amplitude for such states is largest around the very shallow tail at large r. Specifically, it's largest around the classical turning point, where the state's energy is equal to the potential. For this reason, we target high vibrational number molecular states. (The full calculation for the wavefunction at a position r, which comes from the WKB approximation, also has an exponential containing the integral over all accumulated momentum to that r.)

Finally, PA rate is proportional to the dipole moment overlap of the two states and the intensity of the driving beam. The dipole moment overlap is similar to that of the dipole overlap between the asymptotic single-body transition and the saturation intensity increases with decreasing vibrational number. When PA is being driven, the atomic cloud experiences increased two-body loss because molecule formation leads to inelastic collisions and because molecules don't see the trapping light. [17]

The conceptually easiest way to find the scattering length using single-photon PA is by using it to map out the unbound scattering wavefunction. The scattering length is the x-intercept of the scattering wavefunction's asymptote. A high vibrational state's wavefunction is localized at its classical turning point on the potential. This can be approximated as a delta function. Since the transition probability is proportional to the overlap of the ground state wavefunction with the excited state wavefunction, which is the space integral over the two wavefunctions, the probability of transitioning to a certain vibrational state is proportional to the atom's wavefunction at the classical turning point of the probing excited state. In measuring the loss rate over many PA resonances, one can back out the transition rate and therefore the ground state wavefunction over a range in space corresponding to the classical turning points addressed by the PA frequencies. Using this data and numerical calculations for the ground state wavefunction, it's possible to back out the atomic scattering properties. [17]

However, this process requires finding a large number of vibrational states, which are separated by up to several GHz, to within the linewidth of the transition, on the scale of 5MHz. There is an easier way of backing out the atom's scattering length. In 2-photon

PA, a second photon couples the excited molecular state to a high-vibrational state in the ground state molecular potential. This leads to decreased loss in varying ways. When the second laser has a high intensity, the two states are AC stark shifted apart. This means the first laser is no longer on resonance, PA stops happening, and the loss decreases.[17] This has a large linewidth on the second laser, as it is power-broadened. At lower intensity, the coupling causes the formation of a dark state which means that the PA transition stops occurring, which leads to decreased loss. This has a much narrower linewidth. The goal is to find the binding energy, which is the energy difference between the two lasers at the least bound state. The molecular potential at long range is known roughly:

$$V(r) = \sum_{n=6,8,10...} \frac{-C_n}{r^n}$$

where C_n are the van der Waals coefficients for the colliding atoms. However, this is only the long-range potential. One more parameter is needed to calculate the energies of each bound state. Since this parameter is determined by the potential at short range, which is difficult to simulate, it needs to be found experimentally. We can use the newly found binding energy to find that numerical factor to get an accurate potential and solve for the wavefunction. [14]

The vibrational spacings in a molecular potential, which we are trying to find with PA spectroscopy, can be solved for approximately using the Leroy-Bernstein formula. This formula is also known as the near-dissociation expansion. For a potential that looks like

$$V(R) = -\frac{C_n}{R^n}$$

which is the case for molecular potentials at long range, the energy E of a vibrational state with quantum number v is

$$E = -\left[(v_D - v)\left(\sqrt{\frac{\pi}{2\mu}}\frac{\Gamma(1+1/n)}{\Gamma(1/2+1/n)}\frac{\hbar(n-2)}{(C_n)^{1/n}}\right)\right]^{\frac{2n}{n-2}}$$

where $\Gamma(z)$ is the Gamma function, μ is the reduced mass of the system, and v_D is the non-integer v at the dissociation energy D. [20] At long range, the dominant C_n is C_3, which can be calculated using the natural linewidth.[1] $v - v_D$ is unknown, but can only take a value in the range of -1 to 0 for the highest vibrational state v. While the vibrational energies we can solve for aren't accurate enough to solve for the scattering length without doing PA, we used them to decide the general ranges in which to look.

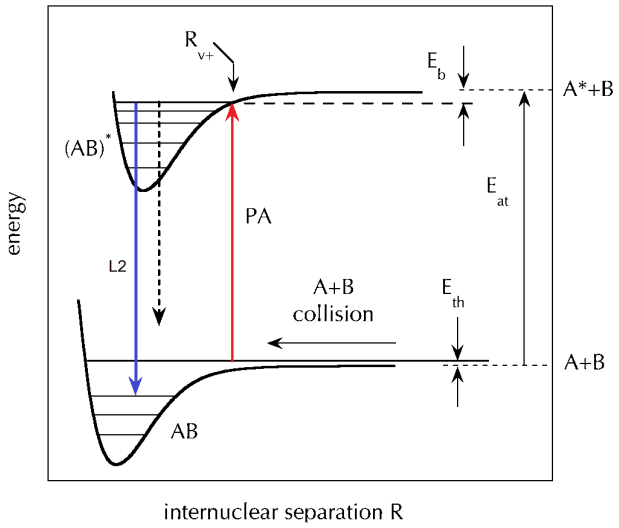


Figure 4.4: 2-photon photoassociation laser scheme. The red arrow labelled PA represents the first PA photon and L2, in blue, represents the second photon in 2-photon PA. The dotted line represents due to from single-photon PA. [21]

We attempted PA spectroscopy in rubidium when we didn't have francium available in order to learn about the conditions needed to see PA. The PA beam diametre was on the order of 0.5mm and the maximum available power was 250mW, however we have used a wide range of powers when looking for PA in rubidium. We only saw PA after the dispenser had a very high current run through it by accident, leading to a MOT that was much larger and denser than what was normally achievable. One issue, but perhaps not the cause for our inability to see PA, experienced with the apparatus was that the brightness of the MOT was fluctuating on the timescale of roughly 5 minutes, with a roughly 10% amplitude in MOT size. Depending on the size of the PA signal, the shot to shot fluctuation due to this effect might be washing out our PA loss. The culprit was the fibre-splitter, which takes the total MOT trapping and repump light as inputs and distributes them into 3 fibres that power each of the MOT beams. We measured polarization fluctuations on the same time scale as the MOT fluctuations, on the scale of 5 minutes. These polarization fluctuations are likely caused by unstable evanescent coupling within the fibre-splitter. This will be solved by building a MOT light distribution optics board that splits, distributes, and couples the light through free space rather than relying on an unreliable fibre-splitter.

We had 56 hours of francium beam time. We started with getting a francium MOT in the capture chamber and transferring the francium into the science chamber MOT, which

was achieved within 3 hours. These steps involved tuning the frequency of the MOT laser and the positions of the magnetic coils. Next, we optimized the push beam to maximize the transfer. This was done by increasing the optical intensity of the push beam to decrease the brightness of the MOT, and using that as a guide to optimize its overlap with the MOT. The intensity was changed back to the optimal power afterwards. Finally, we adjusted the frequency of the push beam to optimize the MOT transfer.

Because two-body loss scales with density squared, a higher two-body loss rate generally means a denser MOT, which is desirable for seeing PA. (PA is a source of increased two body loss, meaning that an increased density will likely lead to a higher PA signal.) Therefore, we did a survey of different MOT settings to find the highest two-body loss rate. We tried two different laser power settings and two different frequency settings: 1mW and 4.6mW measured power in each arm, 42MHz and 53MHz for the frequency. The beams are 2cm in diameter but the power was measured over the face of a 0.95cm diameter powermeter head, giving intensities of 1.42mW/cm^2 and 6.57mW/cm^2. Francium's D2 transition saturation intensity is 2.67mW/cm^2, which gives us saturation parameters of 0.53 and 2.46. Francium has a slightly larger natural linewidth than other alkali atoms, at $2\pi \cdot 7.57$MHz. [15] This gives us detunings of 5.5Γ and 7Γ. We chose 1mW and 42MHz because this combination of settings gave the largest loss coefficient. We didn't change the repump intensity throughout the experiment and were sending in a constant 11mW with one of the MOT beams. The next thing we changed was the photo-association beam intensity. The PA beam had a waist of roughly 0.5mm and a maximum power of 172mW. This results in an intensity of about 88W/cm^2, which is very large compared to the saturation intensity of any single-atom transition. Although photoassociation can saturate at much higher intensities than the one we used, we used a range of smaller intensities to avoid saturating this transition.

The PA rate was measured through the MOT's fluorescence. Upon being captured in the science chamber, the MOT's fluorescence was measured as a function of time with a stack of lenses imaging the emitted light onto a silicon photomultiplier tube, which is a high sensitivity photodetector. As soon as the MOT is loaded, the atoms start escaping so the number of trapped atoms starts decaying. In a MOT, the two significant modes of loss are two-body loss and single-body loss. At the beginning of the curve, the two-body loss is the dominant source of loss, so we can find the two-body loss coefficient without seeing any single-body loss. The two-body loss rate is proportional to density squared, and the trapped atom density is described by $n(t) = \frac{n_0}{1+\beta n_0 t}$, where β is the two-body loss coefficient and n_0 is the initial density. (In contrast, the single-body loss rate is linear in density, leading to an exponential decay. This loss rate takes over later in the MOT's lifetime, when density is small.) We expect PA to look like two body loss – we expect to see PA by seeing an increased two-body loss coefficient when the PA light is on. Towards this

end, we initially characterized the two-body loss coefficient without any PA light. During our search for PA, we compared the loss coefficient with PA light present to it without PA light to look for increased loss when the PA light was on.

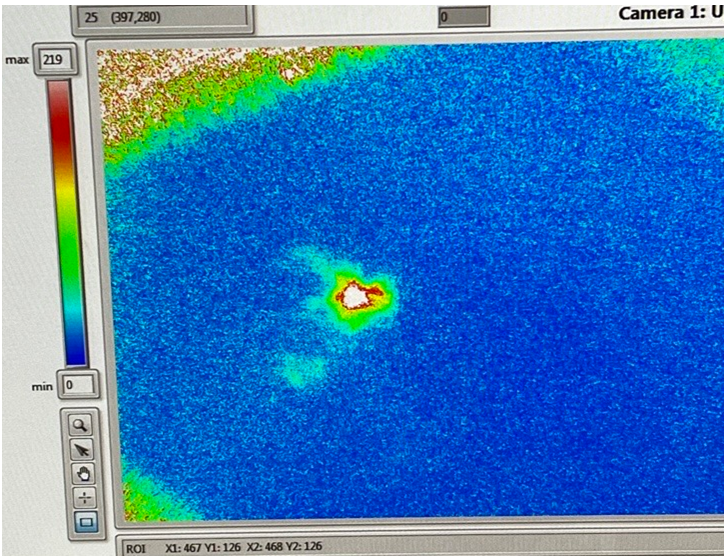


Figure 4.5: The francium MOT near the beginning of the beam time.

We started doing PA while chopping the science chamber MOT beams (meaning flashing them on and off), with the PA beam flashing on when the MOT light turns off. This was done because we suspected that the MOT beams can suppress photoassociation. It has been shown that beams detuned blue from the atomic transition can suppress photoassociation, and while the MOT beams are red-detuned from the atomic transition, they are blue-detuned from the vibrational state transition, so we wanted to avoid this possibility. However, at a certain point we turned off the MOT chopping because having the MOT off for any time leads to increased (non-photoassociative) atom loss during the measurement.

We looked for photoassociation in the potentials that asymptote to $P_{1/2}$ and $P_{3/2}$, the D1 and D2 excited states. Photoassociation lines are the most dense within several gigahertz red of the atomic transition, so we tried to look for photoassociation at such detunings on both transitions. However, on the D2 transition, the laser would hop modes over a large range where we would have expected to see the densest grouping of modes. When sweeping the PA laser, we took steps on the order of 10MHz and broadened the PA laser by frequencies larger than the step size to make sure that no resonances were missed. The broadening was achieved by frequency modulating a double-passed AOM. We scanned

over ranges of a few hundred megahertz to a few gigahertz at various detunings between 3 and 55 GHz from the D1 and D2 lines, given by the calculations mentioned earlier.

Unfortunately, we did not see any signs of photoassociation during the experiment. An analysis of the MOT suggests that it wasn't dense enough to observe PA. The atom number was measured by flashing a short duration, high intensity burst of near-resonant light, and measuring the emitted fluorescence. This measured fluorescence is equal to the product of the number of atoms in the MOT, the scattering rate, the burst time, and the fraction of emitted photons which were measured. With the intensity and detuning, the scattering rate can be calculated. The fraction of emitted photons that are seen by the photomultiplier tube is equal to the fraction of the solid angle captured by the imaging lens multiplied by the PMT's sensitivity. Using these quantities and the measured fluorescence, one can estimate the number of atoms in the MOT. The MOT's diameter was estimated knowing the camera's pixel size and the imaging optic amplification. Our estimate was 2.9×10^5 atoms in total and a density of $5.5 \times 10^9 \mathrm{cm}^{-3}$. Despite the MOT's low density, it was experiencing higher two-body loss than expected. A theorist collaborator estimated a two-body loss coefficient on the order of $10^{-11}\mathrm{cm}^3\mathrm{s}^{-1}$ to $10^{-10}\mathrm{cm}^3\mathrm{s}^{-1}$, whereas our measured coefficient was $5.2 \times 10^{-10}\mathrm{cm}^3\mathrm{s}^{-1}$. This implies that the MOT light was causing increased two-body loss.

One of the main forms of two-body loss that is present in a MOT is caused by light-assisted collisions, which temporarily excites a pair of colliding atoms into a more attractive potential. They accelerate in this potential, and upon decay back into a less attractive potential, have enough energy coming out of the kinetic interaction to escape the trap. Another two-body loss mechanism is hyperfine-structure changing collisions, which involves one or two atoms in the upper ground hyperfine state flipping into the lower ground state upon collision, turning that energy into kinetic energy and shooting both atoms out of the trap. This loss mechanism is enabled by the presence of atoms in the upper hyperfine state, which is affected by the MOT settings. While we tried to maximize the two-body loss to improve the density, it's possible that we instead increased the two-body loss coefficient of one of these mechanisms, thereby making the measurement less sensitive to a photoassociation signal. We decided to use the loss to optimize the density because our only other diagnostic tool was looking at the MOT's fluorescence. The MOT gives off more fluorescence when it has more atoms, but increasing the scattering rate without changing the number or density will also increase the fluorescence. Figuring out what was causing the increase in fluorescence would take more time than we could afford. The solution to this is to optimize the density using absorption imaging, which involves sending a beam through the trapped atomic cloud and measuring how much of the light was scattered as a function of position. While the absorption rate does depend on the beam's properties,

the beam's properties don't depend on those of the MOT, meaning it's possible to know how the number and density change as the MOT settings change.

The next step is to trap francium in an ODT before doing photoassociation in it. Because the ODT does not work by driving atoms to an excited state, we won't see as much two-body loss in an ODT as we do in a MOT, which means that the atoms in the ODT can be trapped more densely. A compression step or sub-Doppler cooling step would also help with atom number in the ODT. Finally, absorption spectroscopy would be really useful because of the ability to optimize the atom number and density.

Chapter 5

Ag at the University of Chicago

The silver part of the francium silver experiment is worked on at the University of Chicago by Prof. David Demille's group. We want to use silver in this experiment because its high electron affinity makes alkali-silver molecules much more polar than bialkali molecules.

The experimental apparatus at Chicago follows the atomic experiment structure laid out in the introduction quite closely. (It uses a commercial oven design due to the high temperature the oven needs to run at 900°C.) After coming out of the oven, the atoms pass through a spin-flip Zeeman slower, which is different from the two kinds of Zeeman slower described in the laser cooling chapter. (This Zeeman slower's magnetic field passes through zero.) The silver is trapped in a MOT, and the goal in the long run is to sub-Doppler cool it, load it into an optical dipole trap, and evaporatively cool it. Sub-Doppler cooling had not been achieved with silver at the time that I visited Chicago, so one of the goals was to push this boundary.

The two goals of my time there were to do photoassociation spectroscopy with silver and to sub-Doppler cool it using gray molasses. Photoassociation spectroscopy is important because it's the method we use to find the scattering length. The scattering length informs us about whether we can evaporatively cool the silver without large amounts of loss. This is discussed more in the TRIUMF chapter. Sub-Doppler cooling silver is important because its low ground state polarizability makes an ODT's trapping potential more shallow for silver. A lower temperature means more silver can be trapped in an ODT. Gray molasses was chosen because the D2 excited hyperfine states are too close in number of linewidths to each other for polarization gradient cooling, the most common choice, to be used. Unfortunately, due to issues with the magnetic coils and lasers, I was not able to achieve

either of these goals. However, I worked on various other projects that were also important to the experiment in the longer term.

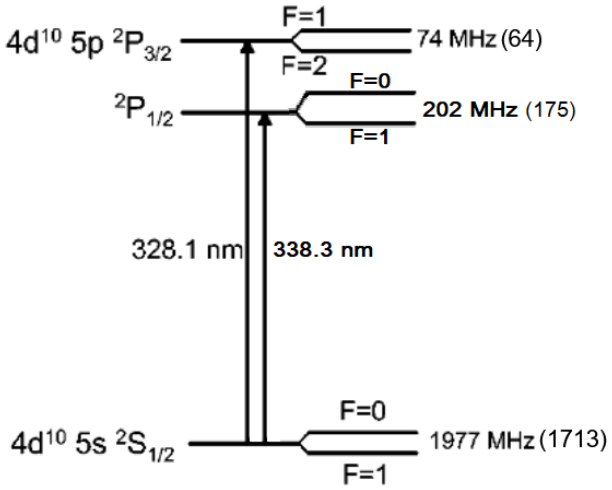


Figure 5.1: The D1 ($S_{1/2} \rightarrow P_{1/2}$) and D2 ($S_{1/2} \rightarrow P_{3/2}$) lines lines of ^{107}Ag and ^{109}Ag. The values in the brackets represent the values for ^{107}Ag. [18][36]

I worked on the MOT magnetic coils and then on the MOT optics and finally regained a MOT after it had been lost due to the aforementioned issues. After the lasers had to be taken apart, I worked on an electronics box that interfaces between the analog control system and power supplies that would power magnetic coils in the experiment. Finally, I worked on some magnetic coils that will cancel out external magnetic fields where the atoms will be trapped.

5.1 Silver MOT

Silver is an interesting atom in ultracold atom physics because it's so similar to alkali atoms despite not being one. This is because of its atomic structure. Silver has a $[Kr]5s^24d^9$ atomic structure, and one of the $5s^2$ electrons moves to the 4d orbital to fill it. This results in a single valence electron in the 5s shell, which looks like the atomic structure of an alkali atom. As a result, silver can be laser-cooled like any other alkali atom. However, it also has some differences. For example, its D1 and D2 transition frequencies are much larger than those of any other alkali atoms, sitting at 885.9THz and 913.5MHz respectively. [18] [36] These correspond to wavelengths of 338nm and 328nm respectively. In contrast, the

shortest wavelengths of any alkali atom belong to sodium, sitting at 589.8nm and 589.2nm for the D1 and D2 lines respectively. (The longest wavelengths belong to Cs, at 894.6nm and 852.3nm. [34])

Silver also has a much larger saturation intensity and natural linewidth than alkali atoms. Silver has a D2 saturation intensity of 87mW/cm^2[36], whereas alkali atoms have saturation intensities near 2.5mW/cm^2. Silver has D1 and D2 natural linewidths of $2\pi \cdot 21.5$MHz and $2\pi \cdot 23.4$MHz,[5] whereas any alkali atoms have natural linewidths near $2\pi \cdot 6$MHz. Silver has a much lower ground state polarizability, which is a measure of the dipole force it experiences, compared to most other alkali atoms. Finally, silver has a high electron affinity. Molecules that contain silver have much larger electric dipole moments than any bialkali molecules. This is the main reason we want to use silver.

Some of these qualities make trapping silver harder than trapping any alkali atom. As seen in the laser-cooling chapter, the frictional coefficient α in a MOT is proportional to the saturation parameter $s = I/I_{Sat}$. A higher saturation intensity means that more light is needed in a MOT. Using the typical saturation parameter used in a MOT, $s = 1$, means a silver MOT needs 87mW/cm^2 in each arm. We use beams 1cm in diameter, meaning that we need almost 70mW in each arm for cooling alone. A natural linewidth of $2\pi \cdot 23.4$MHz and an excited D2 hyperfine state splitting of 74MHz means that the pumping rate into the ground state dark to cooling light is 0.4MHz, or roughly every 2.4µs. 4mW of repump light is shined in to deal with this.

Having the cycling transition frequency in the UV with such a high saturation intensity means that the MOT has to use optics made for high power UV light. The most useful tool that is affected by this is the fibre optic. Regular fibre optics are destroyed by high powers of UV light, and fibre optics made for UV light aren't commercially available, so all the light coming out of the doubling cavity has to travel through free space. This makes the optics set-up much more sensitive to alignment, particularly the alignment closer to the doubling cavity. Additionally, such high power in the UV makes the laser itself expensive and high-maintenance. For example, the UV causes colour centres appear in the doubling crystal, which decrease lasing efficiency and grow over time. This makes translating the crystal every few months necessary to allow the colour centres to relax out.

Silver's much larger natural linewidth means that its Doppler temperature, given by $T_D = \hbar\Gamma/2k_B$, is 561µK, or over half a millikelvin. Additionally, since its hyperfine splittings are so small compared to its natural linewidth, it does not seem to experience PG cooling in a MOT. As a result, a silver MOT is much hotter than a MOT of any alkali atom.

Finally, silver's high-frequency transitions means that an ODT looks shallower to it

compared to other alkali atoms. At large detuning δ compared to natural linewidth Γ, the ODT's trapping potential, mediated by the dipole force, scales as $U \sim \alpha I \sim I/\delta$, where α is the polarizability, which is inversely proportional to detuning δ. Since the highest powers available are generated at 1064nm, the detuning for silver is very large, resulting in a smaller polarization and therefore a smaller trapping potential. (Silver's transition frequency of 885THz gives a detuning of 600THz from 1064nm. This is 6 times larger than in rubidium, 3.6 times larger than in lithium, and 8.5 times larger than in caesium.) In order to load an ODT with a number of silver atoms comparable to what's achievable with alkali atoms, we need to get the silver denser and to a lower temperature before loading into an ODT. Because the MOT cannot get even the silver atoms as cold as other alkalis, a strict sub-Doppler cooling step is necessary. Because of its independence from hyperfine splitting and demonstrated efficacy in Li, Λ-enhanced gray-molasses was chosen. (This cooling mechanism is described in the laser-cooling chapter.)

The silver apparatus at the University of Chicago uses Bitter coils to generate the MOT magnetic field. The benefit of Bitter coils is that they are water-cooled more effectively than coils made of wrapped hollow wire, which is the more commonly used option. Unlike magnetic coils made of wrapped wire, Bitter coils are made of overlapped thin copper discs. There are holes drilled into the disc every 45 degrees, with a notch cut out between two holes on each disc, forming a C-shape. The discs are stacked on top of one another, each rotated by one eighth turn from the one below it. The overlapping part of the coils between the two cutouts, have a copper plate connecting them. There's three oval rubber gaskets between the discs, with two overlapped sets of holes within each gasket. Screws pass through the overlapped holes, keeping the coils stacked together. The water cooling water also travels through the holes, kept in the coil by the rubber gaskets. One screw, made of copper, carries current from the bottom of the coil back to the top, in order to have the two connections at the same spot. The design used for this coil can be seen in Figure 5.2 on the right. [26]

Around when I first arrived in Chicago, 3 bitter coils broke in the span of a couple weeks after working for almost a year - the bottom one, then its replacement, and finally the top one. This resulted in the coils spewing water over the optics and vacuum chamber. The fault point was the same in each of them, the current-carrying screw. The suspected cause was too much stress being put on the screw. Using a screw made of a stronger copper alloy solved the problem, with the coils not leaking again. However, it's possible some sort of corrosion also contributed, due to the fact that after the first time the coil broke, the second and third coils broke in short succession after having coolant sitting in them without flowing for a couple days to a week. Additionally, the applied stress should have not been enough to break the screw on its own. I took part in draining all of the coolant

out of the cooling system and replacing it. In doing so, we noticed that the coolant filter wasn't in place, which also contributes to the hypothesis that it was caused by corrosion. Upon noticing this, we installed the filter. I was also involved in cleaning, reassembling, and replacing the Bitter coils throughout this process.

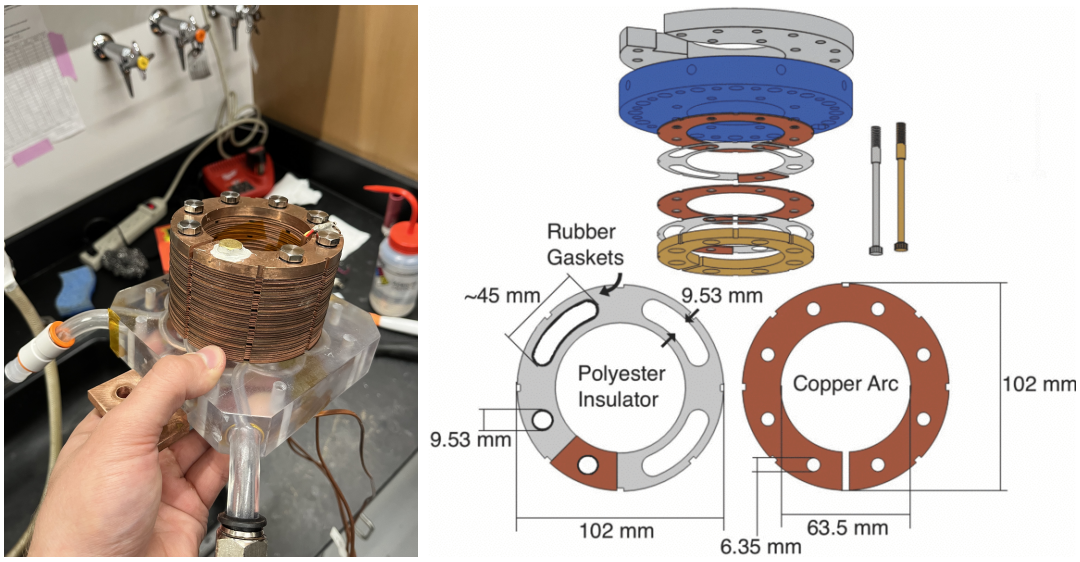


Figure 5.2: On the left, the replacement lower Bitter coil for our MOT. On the right, the Bitter coil design from Prof. Cheng Chin's lab.[26]

After the Bitter coils were replaced and working again, we started trying to get our MOT back. The MOT beams were roughly 1cm in diametre, 87mW/cm^2 in intensity (roughly saturation intensity), and were detuned red from the cycling transition roughly by 1.5Γ (35MHz). There is typically 4mW of repump light in total. Aligning the MOT involved centring the incoming beam on the first viewport and walking it to overlap the reflection off the viewport overlapped with the beam source. Next, the beam was walked to make sure that it came out of the opposite viewport centred. Finally, the reflected beam that makes up the opposing MOT beam was centred on both viewports. The tools used include UV-reactive targets that mount onto the viewports to centre the the beams on each viewport and pinholes that allow more accurate overlapping of reflected beams. The MOT beam alignment was done in all 3 dimensions. The vertical was the most difficult because the vertical axis had bucket viewports, which are sunk 6" into the vacuum chamber in order to accommodate the magnetic coils. Additionally, the presence of optics made it difficult to see into the top or bottom of the viewports. We achieved a silver MOT with a density of roughly 8×10^{10}cm^{-3} and temperature of roughly 1mK.

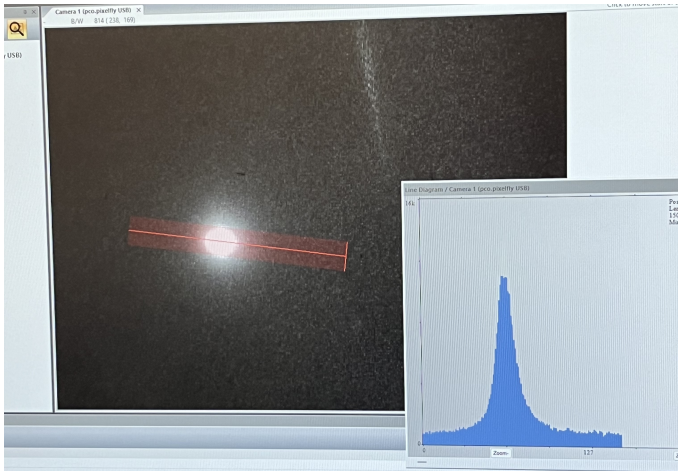


Figure 5.3: The silver MOT we achieved. This MOT had a temperature of 1mK and a density of $8\times10^{10}\text{cm}^{-3}$.

The next steps were to start working on gray molasses or photoassociation spectroscopy in silver. This would have involved setting up the optics needed for these processes. However, the crystals in both the doubling cavities were deteriorating due to misalignment and had to be taken out for replacement. They hadn't been realigned because realigning the doubling cavity would mean having to realign all of the optics in the silver apparatus. The crystals weren't replaced until after I left Chicago.

5.2 Compensation Coil Design

I designed a pair of magnetic coils that is meant to cancel out unwanted external magnetic fields in the centre of the chamber, where the atoms will be. These coils are meant to be mounted onto the optics table, and are specifically meant for the two horizontal dimensions in the chamber. They will be mounted outside of the vacuum chamber, each one 13cm away from the centre of the chamber. Each pair of coils will share an axis and the two pairs' axes will be perpendicular to each other, allowing control in the two dimensions. The third axis is done separately because of the complications related to mounting a coil vertically.

These coils had several constraints to work around. The first was a low inductance to allow a quick shutoff. The back EMF (electro-motive force) that arises from a shut-off

of a magnetic coil is equal to $\varepsilon = L\frac{dI}{dt}$. The MOSFET switches that we use have a spec sheet shut-off time of 100µs, and the maximum current we can drive is 10A. The MOSFET switches can handle a maximum voltage of 100V. Setting 100V as the maximum EMF and using 10A/100µs as the rate of change of the current, we find 1mH to be the maximum inductance the pair of coils can have. 0.75mH was chosen for the maximum inductance the coils can have for safety.

The second constraint was a high enough field at low currents. The maximum current we can use is set by the power supplies we're using, and is 10A. The reference scale for the field strength is the Earth's magnetic field, which is the most important residual field we are trying to eliminate. The Earth has a magnetic field between 0.25G and 0.65G. However, we want to be able to compensate for other magnetic fields that may arise from other experiments or other labs. These fields are typically also on the same scale as the earth's magnetic field. We aim to be able to reach 5G (roughly 10× the Earth's magnetic field) at the centre of the chamber at the maximum current (10A). This gives a large leeway for error in building the coils and it allows us to nullify external fields while running much less than the full 10A through the coils.

The last constraint was the geometry of the experimental apparatus. The coils were to be placed as close to the vacuum chamber as possible in order to maximize the magnetic field at the centre of the chamber and to not get in the way of the optics surrounding the chamber. This places each coil 13cm away from the centre of the chamber. At this position, the vacuum chamber has protruding gate valves at 45 degrees to the axes along which the magnetic coils will be oriented that set a constraint of the maximum width of the coils. The height of the table sets a constraint on the height of the coils, as the coils' centre needs to be level with the centre of the chamber.

I simulated the magnetic coils to find the magnetic field and inductance of a pair of coils as a function of their spatial dimensions, the number of turns, and the current running through the coil. The code was written in Mathematica, and used the Radia package. It was created by a different student for a different project, but the code required a lot of debugging in order to simulate these coils correctly. The general structure of the code is as such: it first generates the spatial dimensions of the magnetic coil using the desired inner dimensions, the number of wraps per layer, and the number of wrapped layers, generating a rectangular cross-section. Next, it calculates the inductance and resistance of the generated structure, and finally it calculates the magnetic field of the coils using Biot-Savart law. It's important to note that to get an accurate inductance it needs to be measured and not simulated, as it's been observed that a coil's actual inductance is usually 3 to 6 times smaller than what is found in simulation. The code that estimates the inductance uses the equation for a finite cylindrical solenoid instead of a square coil. This is a lot easier to

simulate and we design our coils to work with more current than we need anyway. As a result, we have a large leeway in setting a max current. The inductance simulation is still useful because it gives us an order of magnitude measurement for the inductance, which we use to inform the initial design of our coil.

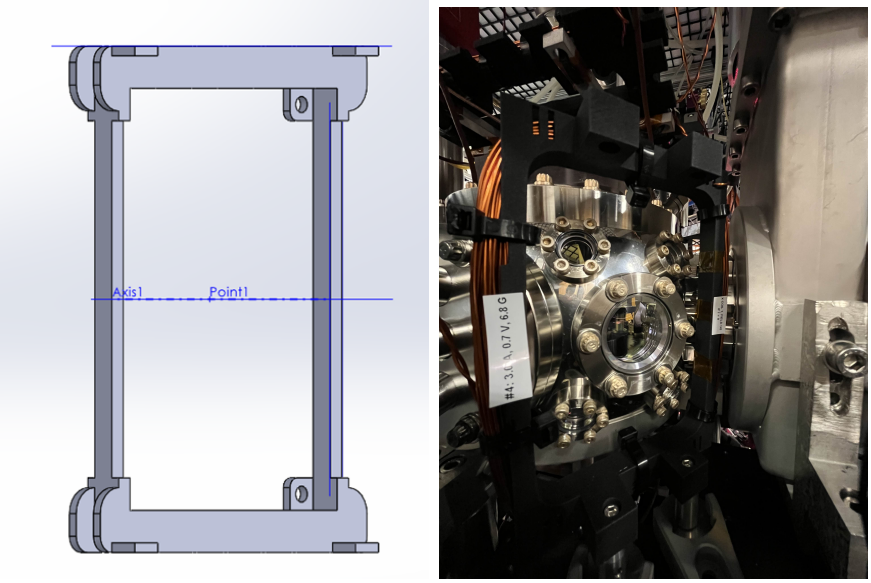


Figure 5.4: On the left, the magnetic coil mount for the compensation coil I designed in Solidworks. On the right, the final design of the magnetic field compensation coil, mounted onto the the experimental apparatus. The metal gate valve on the left sets the spatial constraint on this magnetic coil.

The coils I designed have inner dimensions of 5"×7". This design was iterated upon after I left Chicago, and the number of layers and the number of turns per layer in my design was lost, however they were similar to those in the final design. The final design of the coils, which was completed by a grad student in Prof. Demille's group, has dimensions of 4"×6", with 6 rows of 6 turns. The coil I designed was too large to fit the gate valve, so the next iteration made it smaller. The pair of coils generate a magnetic field of 0.69G/A at the centre of the chamber, and each coil has a simulated inductance of 3.26mH. As mentioned earlier, the coil's measured inductance is generally 3 to 6 times smaller than what is calculated in the simulation. The true inductance will be measured, and the maximum current will be calculated with the equation used to find the coils' desired maximum inductance earlier.

Let's look at the dependence of inductance and magnetic field on the coil's size. The

inductance of an solenoid is proportional to its cross-sectional area and the number of turns squared: $L \sim N^2R^2$. In finding the behaviour of its magnetic field as a function of size, we can approximate the square coil as a finite loop. The diameter d is a stand-in for the edge lengths l_1 and l_2 of a roughly square coil, with d being equal to the diameter of a circle with $l_1 \times l_2$ cross-sectional area. I can justify using one variable to stand in for two because the ratios of the edge lengths we used are on the order of ~ 1. The magnetic field of a round wire loop at some point along its axis z, as a function of its radius, is given by

$$B_z = \frac{\mu_0 NI}{2} \frac{R^2}{(R^2 + z^2)^{3/2}} \tag{5.1}$$

and goes as $\sim R^2/z^3$ when $R \ll z$ and $\sim R^{-1}$ when $R \gg z$. B_z decreases with increasing z, so the ideal would be to have the coils as close to the centre of the chamber as possible. The closest the magnetic coils can be to the centre is 13cm, giving us our fixed z. Our radius is constrained on the scale of 6-8cm so unfortunately, we're working in the regime where $R \sim z$, which means we cannot use either approximation. This function increases until a maximum at $R \simeq 18.4$cm, so we expect increasing size will monotonically increase magnetic field within the size range we have available. B_z is linear in R at $R = 5.64$cm, and remains roughly linear until around $R \simeq 7.5$cm. As R passes ~ 8cm, it becomes less linear, with second derivative reaching 10% of the first derivative at $R \simeq 7.75$cm and continuing to rise. Since the second derivative is negative, the field's gains in R diminish with increasing R. In contrast, the inductance increases with radius quadratically. B_z scales linearly and L scales quadratically with N. This means that, when limited by inductance, trading off size for more loops leads to gains in magnetic field. However, this relationship is not certain to hold with a square coil.

Finally, I used Solidworks to model a mount for the magnetic coil. It was designed to accommodate the cross-sectional dimensions of the magnetic coil and to make it easy to wrap a coil with those dimensions without having issues related to the coil spreading out. This was done by adding a lip at the top and bottom that fits the correct number of coils in width, and having the lip be tall enough that it allows the correct number of those layers of coils. The sides of the mount don't have a lip to allow the coil to be as wide as possible. Additionally, the top and bottom lips have feet with holes that allow the coil to be screwed or clamped onto the optics table. There are also cutouts with holes that allow the coil to be mounted onto a vertical bracket. The coil mount was also remade when the coil dimensions were changed, but the design used for the new coils was based on mine.

5.3 Current Controller Box

To power the experimental apparatus's magnetic coils, we use Delta Elektronika ES150 power supplies. These coils include the coils used for the MOT and the compensation coils. These power supplies take in and output different signals through a single DSUB port, whereas the control system uses BNC signals. Therefore, we need to interface between the two. I designed and built a current control box that does exactly this. The current controller box also contains a MOSFET switch that runs in series with the magnetic coils and allows a quick shut-off of the coils. Quick shut-off of the magnetic coils is important for the experiment, as certain stages of the cooling procedure cannot be done with a magnetic field present. (One example of this is the sub-Doppler cooling stage.)

The MOSFET switch (or MOSFET relay) allows for shut-off times of less than 100µs, whereas most of the experimental phases (such as the MOT phase and sub-Doppler cooling phase) last for milliseconds to seconds. In contrast, the power supplies take 7ms to fall from 15V to 0V. The longer the magnetic field takes to shut off, the more atoms are lost.

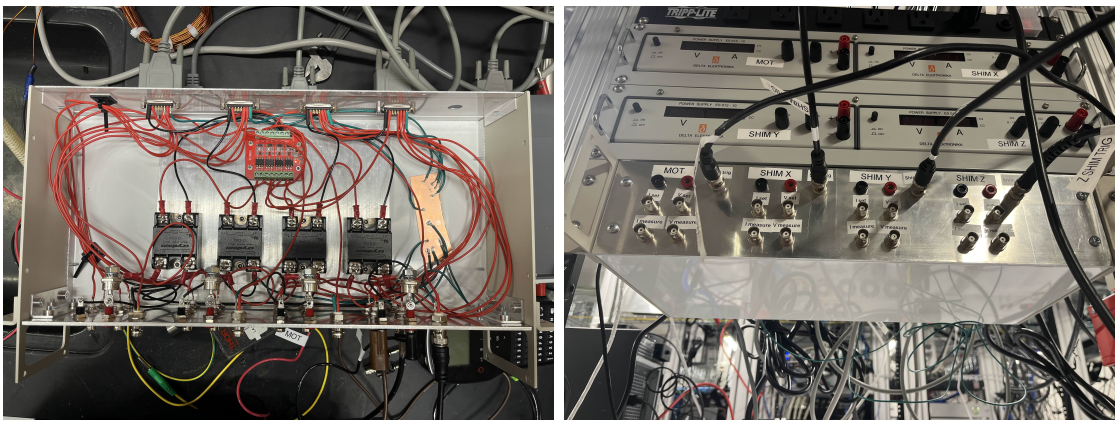


Figure 5.5: On the left, the inside of the current supply box. The rear panel is further from the viewer. On the right, the mounted box with the 4 power supplies and a labelled front panel.

The rear panel has DSUB wire connectors with wires soldered to the pins corresponding to the desired channels. These were the set current, set voltage, measure current, and measure voltage channels. The DSUB wire connectors plugged into a DSUB cable which plugged into the power supply. Each of the wires connects to a BNC port in the front panel. The front panel contains 4 BNC ports for the 4 desired channels, banana cables

ports to put the switch in series with the magnetic coils, and a fifth BNC port to control the switch per magnetic coil, and the box is built to support 4 coils. There are 4 15-pin DSUB connectors in the rear panel. We use 16AWG copper wires to accommodate a maximum current of 10A. The inside of the box and the front panel can be seen in Figure 5.5.

The MOSFET switch is glued to the bottom of the inside of the box. Two wires connect the positive and negative connections of the MOSFET to the banana cable ports in the front panel, with a diode in the circuit that prevents back EMFs that arise from the sudden shutoff (described by $\varepsilon = L\frac{dI}{dt}$). Such backwards EMFs are often quite high, and can damage the MOSFET switch. The MOSFET switch actuator connections from all of the switches are connected to wires which run to a 4-channel MOSFET relay, which provides the correct voltage (12V) to turn the switches on and off. This circuit can be seen in Figure 5.6. Each of the switch BNC ports activate the appropriate MOSFET relay channel. The MOSFET relay is powered by one of the pins from one of the power supply DSUB connectors.

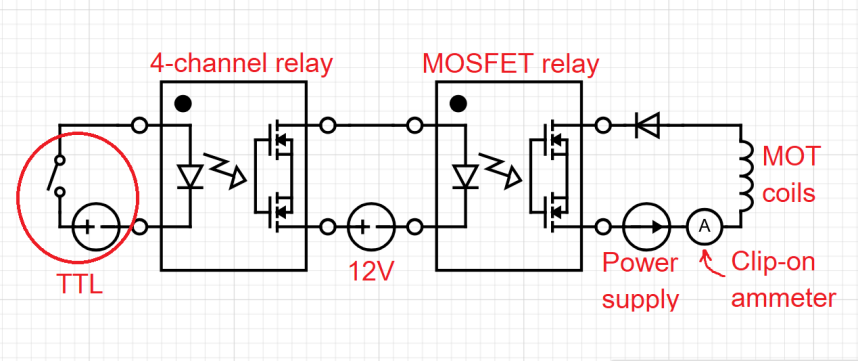


Figure 5.6: The circuit for the MOSFET relay (or MOSFET switch). The clip-on ammeter that measured the shut-off time is hooked up in series with the magnetic coil.

The box is a 2U rack-mounted box and the holes for the BNC, banana cable, and DSUB ports in the front and back panels were modeled in Solidworks and cut with a waterjet. The ports for each magnetic coil and power supply are grouped together on the front panel. The wires were soldered onto the BNC, DSUB, and banana cable ports, and connected to the MOSFET switches with crimped on ring connectors and screws. 16AWG wires were used to be able to handle the full 10A that the power supplies can output. The MOSFET relay has direct wire connectors to which the wires attached. Finally, all of the power supply infterfacing BNCs and DSUB connectors share a common ground, along with the MOSFET relay.

A MOSFET switch was chosen for its fast shut-off and the important time scales are the untrapped atom cloud expansion rate and the length of the experiment phases. I characterized the shut-off time by running current through some magnetic coils and turning it off with the MOSFET switch. The current inside the coil, as a function of time was measured with a clip-on ammeter, and the current in the coils was plotted on an oscilloscope alongside the shut-off trigger. The clip-on ammeter's position is labeled in Figure 5.6, and the measurement results are seen in Figure 5.7. I found that the shut-off starts 330µs after the trigger and takes 90µs to fully shut off. The shut-off time was measured at 5A, which is the current we expected to run through the first set of magnetic coils this box was meant to control, a pair of MOT coils for a different apparatus. The shut-off start time was measured at 2.5A, which was arbitrarily chosen. This value is not expected to change with current as the delay depends entirely on the time it takes the signal triggering the shut-off to reach the MOSFET.



Figure 5.7: On the left, the measurement of the time delay between the shut-off signal, in blue, and the magnetic coil shut-off, in purple. The shut-off starts 330µs after the trigger. On the right, the time the magnetic coil shut-off takes, in blue. The current takes 90µs to shut off.

The MOT phase of the experiment is typically on the scale of milliseconds to seconds, so the shut-off start time can easily be accounted for by sending the trigger 330µs early. We want the shut-off time to be short compared to the time it takes to lose a significant number of atoms after turning off the MOT. In other words, we want most of the atoms to stay within the trapping volume of the MOT beams during the shut-off time. The atom cloud's radius as a function of time is given by $\sigma_x^2 = \sigma_0^2 + (k_B T/m)t^2 = \sigma_0^2 + \sigma_v^2 t^2$. Using a temperature of 1mK and the atomic mass of ^{109}Ag in the expression for the radial expansion rate of the cloud $\sqrt{k_B T/m}$ gives us an expansion rate of 27.5cm/s. Putting in the the size

of the MOT beams (0.5cm radius) as the maximum expansion radius and the initial size of the MOT cloud (0.05cm radius), we find 18ms as the maximum expansion time. The coil shut-off time (90µs) is significantly smaller than this expansion time, meaning that we can switch between phases of the experiment without significant loss in between. In 90µs, the cloud expands 0.000061cm, which is a negligible amount compared to the relevant length scale, the MOT beam size.

Chapter 6

Conclusion

A brief overview of our long-term experimental goals is to make molecules with octupole-deformed francium and silver atoms. Putting an octupole-deformed nucleus in a polar molecule will result in a much more sensitive Schiff moment. Using ultracold molecules means a lot more control over the system that would allow you to eliminate systematic error.

Throughout this thesis, I have tried to shine light on the francium silver experiment and ultracold atoms experiment as a whole. In the first chapter, I went over our goals for the francium silver project and some past notable experiments. In the second chapter, I went through some of the theory behind laser-cooling that's relevant to my work. In the third chapter, I gave a brief overview of the lab, the experiment goals, and what I did at waterloo. I talked about the steps needed to Zeeman slow Cs. This involved building different diagnostic optics set-ups, characterizing the atomic flux coming out of the oven and the velocity distribution reaching the MOT region, and looking at a range of parameters in beginning to optimize the Zeeman slower and preparing to load a MOT. I also talked about controlling and stabilizing lasers at Waterloo. This involved setting up the laser locks, designing parts for and characterizing homebuilt shutters, putting together AOM boards, and building an optics board condo that can store optics set-ups without ruining alignment through vibrations. In the fourth chapter, I talk about the unique apparatus at the francium trapping facility at TRIUMF. I also go through the theory behind photoassociation, and our first attempt at it in francium. In the fifth chapter, I talked about what I accomplished at Chicago. Even though the magnetic coil leaked and one of the lasers broke while I was there, I got a chance to work on the silver MOT and I worked on some other important parts of the experiment.

The next steps at Waterloo are loading the Cs MOT and finishing optimizing the Zeeman slower. Next, we'll compress the Cs MOT and try sub-Doppler cooling in the form of a blue MOT, which has never been done before. The lithium MOT will also be built soon, and the lithium Zeeman slower will also be optimized in the same way. The second experimental chamber for experiments involving molecules will also be built soon. This will include a new design for the oven layout, with the two ovens in line with each other to allow more optical access in the main chamber.

Prof. Demille's group achieved grey molasses working in silver after I left Chicago. The next steps there are refining gray molasses and trying other sub-Doppler cooling techniques and attempting PA in silver. These are both important steps for silver because of how hard it is to trap silver in an ODT. They are also working on an offline francium source, which may reduce the need to wait for francium beam time in the future.

The next steps at TRIUMF are to set up the ODT laser in the francium trapping lab. Our next francium beam time is in September so we will continue working towards seeing PA. Because of how little is known about francium, this is also an important step towards making francium silver molecules.

References

[1] C. Amiot, O. Dulieu, R. F. Gutterres, and F. Masnou-Seeuws. Determination of the $cs_2 0_g^-(P_{3/2})$ potential curve and of $Cs6P_{1/2,3/2}$ atomic radiative lifetimes from photoassociation spectroscopy. *Phys. Rev. A*, 66:052506, Nov 2002.

[2] V. Andreev, D. G. Ang, D. DeMille, J. M. Doyle, G. Gabrielse, J. Haefner, N. R. Hutzler, Z. Lasner, C. Meisenhelder, B. R. O'Leary, C. D. Panda, A. D. West, E. P. West, X. Wu, and ACME Collaboration. Improved limit on the electric dipole moment of the electron. *Nature*, 562(7727):355–360, October 2018.

[3] J Baron, W C Campbell, D DeMille, J M Doyle, G Gabrielse, Y V Gurevich, P W Hess, N R Hutzler, E Kirilov, I Kozyryev, B R O'Leary, C D Panda, M F Parsons, B Spaun, A C Vutha, A D West, E P West, and ACME Collaboration. Methods, analysis, and the treatment of systematic errors for the electron electric dipole moment search in thorium monoxide. *New Journal of Physics*, 19(7):073029, jul 2017.

[4] G C Bjorklund, M D Levenson, W Lenth, and C Ortiz. Frequency modulation (FM) spectroscopy. *Applied Physics B*, 32(3):145–152, November 1983.

[5] J Carlsson, P Jönsson, and L Sturesson. Accurate time-resolved laser spectroscopy on silver atoms. *Zeitschrift für Physik D Atoms, Molecules and Clusters*, 16(2):87–90, June 1990.

[6] Daniel Comparat. Improved leroy–bernstein near-dissociation expansion formula, and prospect for photoassociation spectroscopy. *The Journal of Chemical Physics*, 120(3):1318–1329, January 2004.

[7] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, 6(11):2023–2045, Nov 1989.

[8] J A Devlin and M R Tarbutt. Three-dimensional doppler, polarization-gradient, and magneto-optical forces for atoms and molecules with dark states. *New Journal of Physics*, 18(12):123017, dec 2016.

[9] V. V. Flambaum and A. J. Mansour. Enhanced nuclear schiff and electric dipole moments in nuclei with octupole deformation. *Physical Review C*, 111(5), May 2025.

[10] Timo Fleig and David DeMille. Theoretical aspects of radium-containing molecules amenable to assembly from laser-cooled atoms for new physics searches. *New Journal of Physics*, 23(11):113039, November 2021.

[11] C.J. Foot. *Atomic Physics*. Oxford master series in atomic, optical, and laser physics. Oxford University Press, Oxford, 2005.

[12] Isabella Fritsche. *OSub-Doppler Cooling of Fermionic Lithium*. Masters thesis, University of Innsbruck, 2015. Available at https://www.ultracold.at/theses/2015-fritsche.pdf.

[13] B. Graner, Y. Chen, E. G. Lindahl, and B. R. Heckel. Reduced limit on the permanent electric dipole moment of ^{199}Hg. *Phys. Rev. Lett.*, 116:161601, Apr 2016.

[14] A. Guttridge, Matthew D. Frye, B. C. Yang, Jeremy M. Hutson, and Simon L. Cornish. Two-photon photoassociation spectroscopy of csyb: Ground-state interaction potential and interspecies scattering lengths. *Phys. Rev. A*, 98:022707, Aug 2018.

[15] Timothy Hucko. *Measurement of the 7s 8s M 1 transition in laser-trapped francium*. Phd thesis, University of Manitoba, 2023. Available at https://mspace.lib.umanitoba.ca/items/c8f3e869-7635-4178-8732-a8d416d9eefe.

[16] Michał Śmiałkowski and Michał Tomza. Highly polar molecules consisting of a copper or silver atom interacting with an alkali-metal or alkaline-earth-metal atom. *Phys. Rev. A*, 103:022802, Feb 2021.

[17] Kevin M. Jones, Eite Tiesinga, Paul D. Lett, and Paul S. Julienne. Ultracold photoassociation spectroscopy: Long-range molecules and atomic scattering. *Rev. Mod. Phys.*, 78:483–535, May 2006.

[18] T. Karaulanov, B. Park, Dmitry Budker, and A. Sushkov. Pressure broadening and shift of d1 line of ag by he, ar and n2. *Physical Review A*, 86, 01 2012.

[19] Jacek Kłos, Hui Li, Eite Tiesinga, and Svetlana Kotochigova. Prospects for assembling ultracold radioactive molecules from laser-cooled atoms. *New Journal of Physics*, 24(2):025005, feb 2022.

[20] Robert J. LeRoy and Richard B. Bernstein. Dissociation energy and long-range potential of diatomic molecules from vibrational spacings of higher levels. *The Journal of Chemical Physics*, 52(8):3869–3879, 04 1970.

[21] Ch Lisdat, N Vanhaecke, D Comparat, and P Pillet. Line shape analysis of two-colour photoassociation spectra on the example of the csground state. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 21(3):299–309, December 2002.

[22] Daniel John McCarron. *A Quantum Degenerate Mixture of ^{87}Rb and ^{133}Cs*. Phd thesis, Durham University, 2011. Available at https://etheses.dur.ac.uk/3242/.

[23] William D. Phillips, P. D. Lett, R. N. Watts, C. I. Westbrook, P. L. Gould, and H. J. Metcalf. Observation of atoms laser-cooled below the doppler limit. In *Annual Meeting Optical Society of America*, page FA2. Optica Publishing Group, 1988.

[24] Sara Rosi, Alessia Burchianti, Stefano Conclave, Devang S Naik, Giacomo Roati, Chiara Fort, and Francesco Minardi. Λ-enhanced grey molasses on the D2 transition of rubidium-87 atoms. *Scientific Reports*, 8(1):1301, January 2018.

[25] Tanya S. Roussy, Luke Caldwell, Trevor Wright, William B. Cairncross, Yuval Shagam, Kia Boon Ng, Noah Schlossberger, Sun Yool Park, Anzhou Wang, Jun Ye, and Eric A. Cornell. An improved bound on the electron's electric dipole moment. *Science*, 381(6653):46–50, 2023.

[26] Dylan Sabulsky, Colin Parker, Nathan Gemelke, and Cheng Chin. Efficient continuous-duty bitter-type electromagnets for cold atom experiments. *The Review of scientific instruments*, 84:104706, 10 2013.

[27] M. S. Safronova, D. Budker, D. DeMille, Derek F. Jackson Kimball, A. Derevianko, and Charles W. Clark. Search for new physics with atoms and molecules. *Rev. Mod. Phys.*, 90:025008, Jun 2018.

[28] Andrei D Sakharov. Violation of cp invariance, c asymmetry, and baryon asymmetry of the universe. *Soviet Physics Uspekhi*, 34(5):392, may 1991.

[29] G. Salomon, L. Fouché, P. Wang, A. Aspect, P. Bouyer, and T. Bourdel. Gray-molasses cooling of 39k to a high phase-space density. *Europhysics Letters*, 104(6):63002, jan 2014.

[30] V. Sanchez-Villicana, S. D. Gensemer, K. Y. N. Tan, A. Kumarakrishnan, T. P. Dinneen, W. Süptitz, and P. L. Gould. Suppression of ultracold ground-state hyperfine-changing collisions with laser light. *Phys. Rev. Lett.*, 74:4619–4622, Jun 1995.

[31] S. Sanguinetti, R. Calabrese, L. Corradi, A. Dainelli, A. Khanbekyan, E. Mariotti, C. de Mauro, P. Minguzzi, L. Moi, G. Stancari, L. Tomassetti, and S. Veronesi. Accurate measurements of transition frequencies and isotope shifts of laser-trapped francium. *Opt. Lett.*, 34(7):893–895, Apr 2009.

[32] L. I. Schiff. Measurability of nuclear electric dipole moments. *Phys. Rev.*, 132:2194–2200, Dec 1963.

[33] Franz Sievers, Norman Kretzschmar, Diogo Rio Fernandes, Daniel Suchet, Michael Rabinovic, Saijun Wu, Colin V. Parker, Lev Khaykovich, Christophe Salomon, and Frédéric Chevy. Simultaneous sub-doppler laser cooling of fermionic ^{6}Li and ^{40}K on the D_1 line: Theory and experiment. *Phys. Rev. A*, 91:023426, Feb 2015.

[34] Daniel A. Steck. Cesium d line data. Available online at http://steck.us/alkalidata (revision 2.3.3, 28 May 2024).

[35] M. Tandecki, J. Zhang, R. Collister, S. Aubin, J. A. Behr, E. Gomez, G. Gwinner, L. A. Orozco, and M. R. Pearson. Commissioning of the francium trapping facility at triumf. *Journal of Instrumentation*, 8(12):P12006, dec 2013.

[36] G. Uhlenberg, J. Dirscherl, and H. Walther. Magneto-optical trapping of silver atoms. *Phys. Rev. A*, 62:063404, Nov 2000.

[37] Sudhang Varshney. *Optimizing Photoassociation Signals in Small Magneto-optical Traps and Evaporative Cooling Schemes for Shallow Optical Dipole Traps*. Phd thesis, The University of British Columbia, 2023. Available at https://open.library.ubc.ca/soa/cIRcle/collections/ubctheses/24/items/1.0437301.

[38] M. Weidemüller, T. Esslinger, M. A. Ol'shanii, A. Hemmerich, and T. W. Hänsch. A novel scheme for efficient cooling below the photon recoil limit. *Europhysics Letters*, 27(2):109, jul 1994.