The Dissertation Committee for William D. Huntington
certifies that this is the approved version of the following dissertation:

A Neutral Atom Storage Ring and Intense Cold Atom Source

Committee:

Daniel Heinzen, Supervisor

Ernst-Ludwig Florin

Greg O. Sitz

Shyam Shankar

A Neutral Atom Storage Ring and Intense Cold Atom Source

by

William D. Huntington

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

December 2022

...

Acknowledgments

It was a long journey to get to this thesis and there are many people to thank that helped along the way. Foremost amongst them is my loving wife, dancer and astronomer, Susan Huntington. Her support, patience, and love during this part of my life was incredible and I will always be indebted.

Thank you to my parents, Scott and Kelly, for your nurturing yet laissez-faire attitude that allowed me to grow on my own. To my grandparents, "Nana" and "Papa", thank you for your always supportive love of me.

There were many wonderful educators along the way, too many to thank. Most prominent amongst them is Jan Colsby for encouraging me to take AP and honors classes in high school. To my high school physics teacher Mr. Dennison, thank you for your class. To my high school biology teacher Mrs. Martinez, thank you for the wonderful educational opportunities and experiences you afforded me. To my undergraduate physics professor and advisor Dr. Davis, thank you for the expertly taught classes and mentorship you provided.

The machine shop has been an integral part of my research and they have always been professional and willing to go the extra mile to help. Thank you Jack, Kenny, Bryan, Alan, Richard, Flint, Ed and others. I really enjoyed shooting the breeze with y'all.

Perhaps the biggest reason I became enamored with physics is because of books, particularly science fiction. Thank you Isaac Asimov, Lewis Carroll Epstein, Ray Bradbury, Kurt Vonnegut, Richard Feynman, Arthur C. Clarke, Robert Hein-

lein and numerous pulp science fiction authors amongst others for putting ideas to paper and sparking curiosity in myself and countless other scientists.

I wouldn't be the physicist I am today if it wasn't for my tight knit family of fellow physics students at San Diego State University. To Adriana Carmona, Enrique Hurtado and Francisco Santos, thank you for being part of this journey.

I would also like to thank my advisor, Dr. Daniel Heinzen, for his financial support, advice, and guidance.

William D. Huntington

The University of Texas at Austin
October 2022

A Neutral Atom Storage Ring and Intense Cold Atom Source

Publication No. _____

William D. Huntington, Ph.D.
The University of Texas at Austin, 2022

Supervisor: Daniel Heinzen

Presented in this thesis is a proposed design for a continuously loaded permanent magnet meter scale storage ring for confining neutral atoms and molecules. Confinement is generated by neodymium magnets arranged to produce hexapole fields. The ring is in the shape of a "racetrack" with two bending sections of 1 meter radius and is composed entirely of linear segments of magnets. A permanent magnet injection system continuously guides particles from an existing source into the ring using an optical pumping scheme. The dynamics of the design are well explained by charged particle accelerator theory extended to neutral particles. The proposed design is over an order of magnitude larger than previous neutral atom storage rings and provides over 2 orders of magnitude greater trapping depth. An extensive simulation was developed to characterize and optimize the ring and injector system. The design is optimized for our existing source of ^{7}Li atoms, although theory and

simulation indicate it will work well with other paramagnetic species. With our source it is expected to build up a circulating flux of around 5×10^{14} atoms/s and an atom number of about 2×10^{13}. It has previously been suggested that such a design is unstable, but in this work it is shown that an essential stability criteria was neglected.

Also presented is the improvement and characterization of an intense, continuous cold atom beam to be used to load the storage ring. The beam is generated via post nozzle seeding of a supersonic cryogenic ^{4}He jet with hot ^{7}Li atoms. The atomic beam is brought to a focus 176 cm from the nozzle by a 10 cm bore diameter permanent magnet hexapole lens. At the focus the beam is measured to have a flux of $2.3(4) \times 10^{12}$ atoms/s, brightness of $1.8(6) \times 10^{19}$ $\mathrm{m^{-2}s^{-1}sr^{-1}}$, forward velocity of 210(2) $\mathrm{ms^{-1}}$, and longitudinal temperature of 7(3) mK. An improved vacuum design should yield around 10 times higher flux.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Cold atomic and molecular beams are important tools in experimental physics. They are used for precision measurement experiments such as defining the second or searching for time-reversal symmetry violation, dark matter, and spatio-temporal variation of constants [1–5]. They also enable studying ultracold collisions, ultracold chemistry, and quantum gases [6, 7]. Many experiments' Signal-to-Noise Ratio (SNR) are limited by atom/molecular number, flux, and/or density. In a beam, especially a continuous one, only a small fraction of the time-averaged capturable atoms are typically used during measurements. Additionally, many methods of generating cold beams use laser cooling which can require very specific and complicated laser-optical schemes.

This thesis describes two projects motivated by these issues. The first project is the characterization and improvement of a existing novel intense cold continuous source of ^{7}Li atoms. The source is generated without laser cooling and works by seeding hot ^{7}Li atoms from an oven into a supersonic adiabatic expansion of ^{4}He originating from a cryogenic nozzle. The resulting beam is extracted and focused by a Halbach array permanent magnet atom lens [8]. A thorough simulation characterizing the aberrations affecting the atom lens was performed.

The second project is the modeling and design of a meter-scale permanent magnet storage ring, to be continuously loaded by the existing ^{7}Li source. It is shown that the design is applicable to other species as well. The ring is composed entirely of discrete segments of permanent magnet atom lenses. It has previously been found that such a design is unstable [9]. By applying charged particle accelerator theory and carrying out an extensive simulation, I show such a design is stable with the right arrangement of magnets. The design is expected to build up a large circulating flux and atom number under anticipated loss mechanisms.

1.1 An Intense Cold, Continuous Beam

Laser cooling has been the primary method of producing cold beams [10–14]. However, the method is more difficult to apply to molecular and atomic species with complicated energy level structures, though progress is being made [15, 16]. This motivates the use of laser-free methods such as a supersonic expansion where the species expands from an orifice and adiabatically cools. This results in temperatures as low as 200 mK, though one to a few Kelvin is more typical [17–22]. This results in beams with a large forward velocity [23,24]. Another alternative is to seed the species into a cold buffer gas that then effusively escapes through an orifice [25]. Seeding is most commonly done with pulsed laser ablation of a solid target to preserve cryogenic temperatures. The resulting beams have a lower forward velocity of around 40 to 150 m/s and temperatures on the order of one to a few Kelvin.

The source described in this work shares the favorable features of supersonic

and effusive sources while achieving lower temperatures [1]. The cryogenically cooled nozzle produces a ^{4}He jet with a forward velocity of about 210 m/s and milliKelvin temperature thanks to the adiabatic cooling of supersonic expansions [27]. Without cryogenically cooling the nozzle, the forward velocity would exceed 1500 m/s. Seeding before the ^{4}He expansion would result in clogging of the nozzle orifice by frozen ^{7}Li unless only small amounts of ^{7}Li were used. By seeding ^{7}Li into the ^{4}He jet after the expansion we can exploit the low temperatures and the forward directed nature of a supersonic expansion while producing a significant flux of ^{7}Li. Focusing of the ^{7}Li beam increases its intensity and facilitates its separation from the ^{4}He beam which is unaffected and continues to expand. Focusing of neutral atoms has been previously performed [8, 28, 29], but to our knowledge it has not been applied to a seeded supersonic beam.

With this method we have produced a ^{7}Li beam with a flux of 2×10^{12} atoms/s, a forward velocity of 210(2) m/s, a brightness of $1.8(6) \times 10^{19}$ m^{-2}s^{-1}sr^{-1}, and a longitudinal temperature of 7 mK in the moving frame. To our knowledge this temperature is substantially lower than that of any previous seeded jet source. This combination of very low temperature, continuous output, modest forward velocity, magnetic extraction, and high flux distinguishes our work from previous work on seeded supersonic jets and buffer gas sources. A simulation of the seeding dynamics was developed by colleague Jeremy Glick with a complete theoretical model of the Li-He collision cross section from 1 mK to 3000 K. The simulation, in conjunction

[1] The source described in this work was originally designed and built by previous graduate student Michael Borysow [26].

with particle tracing through the atom lens and considering background gas induced losses, explains within 40% the observed flux at the focus. This implies we have a good understanding of the physics from the source to the focus.

An overview of the source physics is given in Chapter 2. The characterization of the source, including simulations of the seeding dynamics and the atom lens aberrations, is described in Chapter 3. To the best of our knowledge no such characterization of atom lens aberrations has been presented before.

1.2 A Meter Scale Neutral Atom Storage Ring

A neutral atom or molecular storage ring can be used to study novel physics [30–36], perform precision measurements [37–39], or amplify the flux of an existing source [40]. A variety of storage rings, or toroidal traps, have been built and demonstrated. Atoms have been trapped in rings using conventional atom trapping techniques of lasers and/or time varying magnetic fields [41–47]. The diameters of these traps are typically less than a mm and are loaded with a pulse of atoms. Trap depths are less than a mK with around 1×10^6 or less atoms. Bose–Einstein condensates (BEC) have successfully been produced with these techniques.

Several molecular storage rings have been demonstrated using static high voltage electric fields [48–50]. This method is only applicable to molecular species with electric dipole moments which precludes the use of ground state atoms. Collections of high voltage wires are arranged toroidally to produce confining hexapole electric fields in which a pulse of molecules are injected. So far only ammonia or deuterated ammonia have been demonstrated. The most robust of these design has

a 50 cm diameter and is composed of 40 individual straight electric hexapoles [49]. The segments allow multiple pulses of molecules to be injected which circulate up to 13 seconds and travel over a mile at a speed of about 124 m/s. Up to 26 packets were injected with half traveling clockwise and the other counterclockwise. The maximum number of stored molecules is about 3×10^7. The lifetime is limited by blackbody-induced transitions out of the trapped state and collisions with background gases. Previous designs had more than an order of magnitude lower lifetime because of a zero field at the center of the hexapole inducing Majorana loss. The transverse trap depth is about 750 mK.

Another approach to trap neutral species is to use static magnetic fields to act on the particle's magnetic moment. This was first successfully carried out in the 1980s with neutrons using superconducting wires to produce a 2 meter diameter toroidal hexapole field [51]. Shortly thereafter a proposal was made to use spatially varying magnetic fields to confine ground state hydrogen atoms in hopes of achieving a BEC [52]. The authors incorrectly claimed that high field seeking particles cannot be trapped in a ring without either spatial or temporal variation of magnetic fields. This error was pointed out by Ketterle and Pritchard when they described the theory of confining high and weak field seeking neutral particles in static toroidal magnetic fields [53].

Weak field seeking particles can be strongly confined in a static magnetic storage ring. However, they are susceptible to loss by spin-exchange when colliding with another particle unless both particles are in a stretched state [54]. High field seeking particles on the other hand are only weakly confined to a small radius thanks

to the centrifugal pseudo force. They are immune to spin-exchange collision losses. Varying the magnetic field in space can be shown to result in strong confining for ground state particles, but the varying magnetic field results in viscous heating. For these reasons, Ketterle and Pritchard suggested that a storage ring to trap high field seeking particles is the most promising design for forming a BEC.

Several static magnetic storage rings have since been demonstrated, all operating with weak field seekers, pulsed loading, and trapping energies of a few mK or less [55–58]. Two have reported the production of a BEC [55, 57]. Spin-exchange losses are avoided by optical pumping into a stretched state. Lifetimes range from 1 to 90 seconds, limited by either collisions with background gas or particle's spin flipping when passing through zeros in the magnetic field. Trap lifetimes of 50 and 90 seconds were achieved by adding a small magnetic field to "plug" the zeros [55, 57]. Diameters range from 1.3 mm to 10 cm and the number of trapped atoms range from 1×10^6 to 5×10^8.

The design described in this work differs from previous designs in several ways. First, magnetic fields are produced by neodymium permanent magnets instead of current carrying coils. Magnets are arranged in a Halbach array hexapole configuration with numerous straight segments defining a racetrack shaped particle orbit. The use of permanent magnets results in a trapping potential of around 550 mK, more than two orders of magnitude higher than existing coil based neutral atom storage rings. This allows loading with larger volumes of phase space and/or hotter collections of particles. Second, optical pumping is used to continuously load the ring. This allows for the build up of a large number of atoms. Using our existing

source for loading and with the anticipated vacuum lifetime of 50 seconds, simulations indicate that a circulating flux of about 5×10^{14} atoms/s and atom number of 2×10^{13} will be generated. This is more than 4 orders of magnitude higher than any existing source. Third, the diameter of the bending portions is 2 meters and the ring encloses an area of about 5 m^2. This is an area nearly 700 times larger than the next largest design [57]. A possible use of a neutral atom storage ring is as a Sagnac interferometer in which the sensitivity is proportional to the enclosed area [59]. The proposed design has the potential to be far more accurate than existing rings.

Charged particle accelerator theory is extended to the case of paramagnetic particles. The theory yields predictions which include particle phase space behavior, density and velocity profiles, stability concerns, and the use of other species. These predictions agree well with an extensive and original 3D particle tracing simulation. In [9] a similar racetrack design was analyzed with 3D particle tracing and was found to be unstable. As explained by this theory, only specific configurations of design parameters result in stability and must be chosen carefully. This was not accounted for in [9].

The Lonestar6 supercomputer at the Texas Advanced Computing Center [2] (TACC) was used to find an optimized design for our ^{7}Li source. Despite being optimized for our species/source combination, the proposed design is shown theoretically and by simulation to be usable with other species/source combination and will achieve nearly the same performance.

[2]The author acknowledges the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this thesis. URL: http://www.tacc.utexas.edu

An overview of the relevant portions of charged particle accelerator theory applied to paramagnetic particles is given in Chapter 4. The simulation of the ring is described in Chapter 5. The proposed design is described and analyzed in Chapter 6. Practical concerns such as vacuum, optical pumping, assembly, and magnet material tolerances are discussed as well.

Chapter 2

Apparatus Overview

This chapter describes an intense cold continuous source of ^{7}Li atoms. The majority of the apparatus was designed and assembled by a previous graduate student, Michael Borysow, and extensive details can be found in his thesis [26]. When I joined the research group, the experiment had been moved from another building and was partially assembled. Changes to the apparatus have been primarily the installation of the hexapole lens, the addition of an observation cell for the focus, a new atomic oven, an improved helium gas manifold and metering system, application of optical pumping, upgrades to the laser and optical system, upgrades to the cryogenic system, temperature control of the nozzle with a heater and feedback loop, and the replacement of LabView with Python for streamlined data collection and analysis.

2.1 Introduction

A schematic of the source is shown in Figure 2.1. Room temperature ^{4}He gas is cooled to about 4.4 K by a cryo-refrigerator while being fed into a small cylindrical copper nozzle. The gas exits the nozzle through a small hole and undergoes a supersonic expansion into a vacuum chamber and cools to about a milliKelvin.

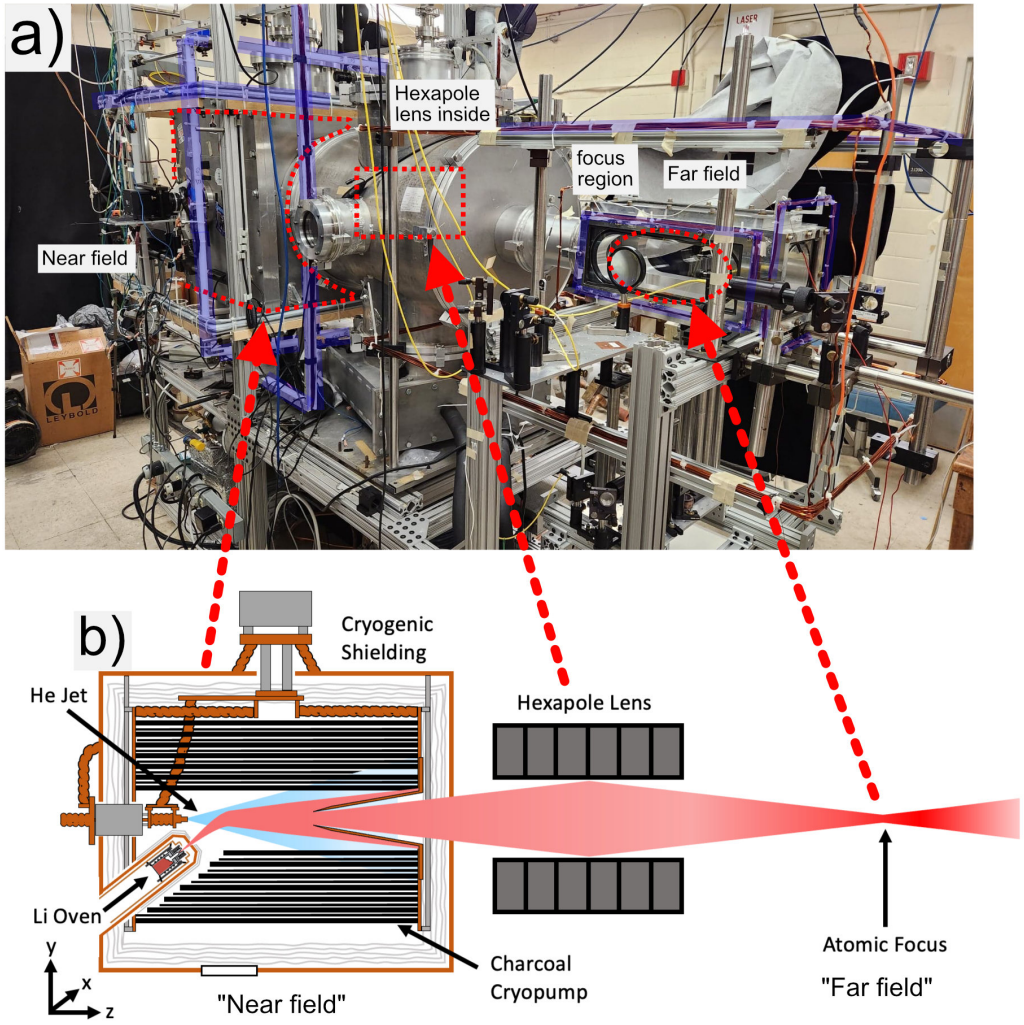

Figure 2.1: Overview of the existing source. As shown, we refer to the region inside the cryogenic shielding the "near field" and the region near the focus the "far field". The camera (FLI ML1001) used for imaging the focus in the far field has been moved to the near field in this photo. (a) A photograph of the apparatus. One coil of each optical pumping Helmholtz coil pair is highlighted in blue. (b) A schematic of the apparatus.

The gas reaches a terminal velocity of about 210 m/s. An effusive lithium beam produced from an oven beneath the nozzle is directed into the helium jet at an angle of 50°. A small fraction of the lithium atoms are entrained in the jet and are cooled to around 10 mK by collisions with the helium gas. The expanding gas is intercepted by a skimmer, which allows a central core of the beam into the following room temperature region of the vacuum chamber. Helium that does not enter the skimmer is pumped away by a charcoal Cryo-adsorption pump [60].

The lithium beam enters a magnetic hexapole (also referred to as sextupole) atom lens, which brings the lithium atoms to a focus by acting on their magnetic moments. In order to produce the highest flux at the focus, the ^{7}Li atoms are optically pumped into the ground state hyperfine level $F = 2, M_F = 2$ (Lithium D line transition and other information is given in Figure D.1 for reference). Without this step roughly half the atoms would be defocused by the lens and lost. It also ensures that all atoms experience the same relationship between magnetic force and magnetic field. Helium atoms are not focused by the atom lens, and therefore continue on ballistic trajectories until they hit a room temperature surface. These atoms are then removed by a large diffusion pump below the atom lens.

2.1.1 Brightness and Brilliance

No single number fully characterizes a beam's distribution in phase space. Besides the obvious metrics of flux (particles per second), intensity (particles per second per area), and temperature, an atomic beam can be further described by

its *brightness* and *brilliance* [1]. These terms are borrowed from the field of optics. Brightness (or radiance) describes the number of particles, per unit time, per unit area in the four dimensional transverse position-angle phase space. The brightness, R, at an infinitesimal point in phase space is given

$$R = \frac{dI}{dA \cdot d\Omega}, \tag{2.1}$$

where I, A, and Ω are flux, area, and solid angle respectively. The average brightness in phase space can be found by integration. In practice this may be challenging because it requires mapping out the distribution of I in phase space. However, in the case of cylindrical symmetry, which occurs often with beams, we can define the average brightness to be

$$\bar{R} = \frac{I}{\pi r^2 \cdot \Delta\Omega}, \tag{2.2}$$

where I is the flux in a circle of radius r and solid angle of $\Delta\Omega$. This is often referred to simply as the brightness rather than average brightness. From the solid angle of a cone we have

$$\begin{aligned} \Omega &= 2\pi(1 - \cos(\theta)) = 2\pi(1 - \cos(\arctan(\Delta v_\perp / 2v_\parallel))) \\ \Omega &\approx \frac{\pi}{4}\left(\frac{\Delta v_\perp}{v_\parallel}\right)^2, \end{aligned} \tag{2.3}$$

where $\Delta v_\perp$ is the full width spread in the transverse velocity, $v_\parallel$ is the nominal longitudinal velocity, and the paraxial approximation was used.

We account for the longitudinal velocity component of phase space with

[1] Unfortunately, the definitions of these terms varies in the literature.

brilliance B. Similarly, this is defined as

$$\bar{B} = R\frac{v_{\|}}{\Delta v_{\|}}. \tag{2.4}$$

2.2 Supersonic Helium Jet

The helium jet is produced by flowing helium gas through a simple sonic nozzle consisting of a 200 μm diameter hole drilled in the thin, flat face of a copper cylinder. Although it is possible that a shaped nozzle could produce better results, a sonic nozzle was chosen for its well-understood properties [61]. A shaped nozzle also has a more forward directed flux which could increase the vacuum load outside of the cryogenic region. Helium flow is metered by an MKS GE50A flow meter. Heat exchangers connected to the two stages of the refrigerator cool the gas so that it reaches a temperature in the range of 4.2 K to 4.6 K inside the nozzle. The number of atoms leaving the nozzle per second is [61]

$$\dot{N}_{He} \approx 0.403\frac{P_0}{k_B T_0}u_0 d^2, \tag{2.5}$$

where P_0 is the stagnation (pre-nozzle) gas pressure, d the nozzle diameter, T_0 is the temperature of the nozzle, and $u_0 = \sqrt{2k_B T_0/m_{He}}$, with m_{He} the mass of a helium atom.

The terminal forward velocity of the jet is [61]

$$v_f = \sqrt{5k_B T_0/m_{He}}. \tag{2.6}$$

Typically this is quite large (about 1750 m/s at room temperature), which is a downside of seeded supersonic jets if a goal is to trap atoms. However our nozzle produces a jet with a much lower forward velocity because it is cooled to cryogenic temperatures.

In an effusive source the density profile follows $n = n_0 \cos(\theta)/r^2$ where θ is the angle from the axis of the nozzle, r is the distance from the nozzle, and n_0 is a constant. In a supersonic source the flux is more forward directed. At distances of about $r > 4d$ the density distribution follows [62]

$$n = n_0 \cos(1.15 \cdot \theta^2)/r^2. \tag{2.7}$$

In order to reach milliKelvin temperatures it is necessary to avoid the formation of helium clusters, since otherwise the heat of condensation is released into the expanding gas [21]. The onset of cluster formation is described by the Hagena parameter [63]

$$\Gamma^* = \kappa \frac{P_0 d^{0.85}}{T_0^{2.29}}, \tag{2.8}$$

where κ is a species-specific condensation parameter. Over a fairly wide range of parameters and gas species, it is found that substantial cluster formation occurs only when $\Gamma^* > 300$. This scaling law has not been tested in our parameter range, but He cluster formation in cryogenic jets with much smaller nozzles than ours has been studied [64–66]. They observed some formation of dimers and trimers with Hagena parameter as low as $\Gamma^* \approx 50$, but heating of the jet generally remained small up to $\Gamma^* \approx 300$. Because of clustering it is not possible to form a cold slow beam of pure

^{7}Li.

Helium has an unusually low κ of 3.85 $\mathrm{K}^{2.29}\mu\mathrm{m}^{-0.85}\mathrm{mbar}^{-1}$ [67, 68], due to the weakness of the He-He interaction and the fact that the helium dimer has only one bound state with an extremely small binding energy of about 1.1 mK [69]. It also has an elastic collision cross-section that increases dramatically as the jet temperature falls [70]. These factors allow helium jets to remain collisional to much larger distances and cool to much lower temperatures than jets of other species.

The nozzle diameter $d = 200$ μm was chosen in order to maximize helium beam brightness without producing an excessively large Hagena parameter. For instance, with $P_0 = 25$ mbar and $T_0 = 4.4$ K the nozzle produces a flow of 200 SCCM and has a Hagena parameter of $\Gamma^* = 290$. So far, we have not seen any clear evidence that our results are affected by helium cluster formation.

In our initial search for the focus we noticed that the nozzle temperature fluctuated. During image acquisition of fluorescence, lasting only a few minutes typically, the nozzle would slowly heat up. The forward velocity would increase and the location of the focus would then drift which was a serious hindrance to locating it. We temporarily solved the problem by flowing helium for long periods of time and allowing the nozzle temperature to equilibriate, but this wasted time and helium. To eliminate this problem we stabilized the nozzle temperature with a 50 watt cryogenic heater clamped to the top of the nozzle and a Lakeshore model 325 PID temperature controller. The stabilized nozzle temperature is a few tenths of a degree above its value with no helium flow, and typically in the range from 4.2 K to 4.6 K. It is important to wait for the nozzle temperature to stabilize before

taking data which requires about 1 minute.

The forward velocity of the jet is about 210 m/s with a He flowrate of 50 Standard Cubic Centimeters per Minute (SCCM) for extended periods of time. As the flowrate increases, the nozzle temperature setpoint may need to be increased to account for the extra thermal load of the flowrate. If not increased the nozzle will heat up past the setpoint and the focus will move. Above 250 SCCM even short runs of 1-2 minutes require the nozzle to held at a higher setpoint resulting in speeds of 217(2) m/s.

As the helium gas expands, its density drops and therefore its collision rate drops. This results in a region of continuum flow near the nozzle, a region of molecular flow far from the nozzle, and an intermediate flow region between these two [61, 71]. In the intermediate region, collisions are no longer frequent enough to maintain local thermal equilibrium so the temperatures parallel and perpendicular to the nominal jet velocity decouple and become unequal. At distances greater than a few nozzle diameters, but still in the continuum flow region, the helium adiabatically cools according to [71]

$$T(z) = 0.287\ T_0 \left(\frac{z}{d}\right)^{-4/3}, \tag{2.9}$$

where T is the helium temperature and z is the distance from the nozzle.

The quality of the nozzle and helium may affect the expansion. We believe we have found this to be the case with our nozzle which had a nicked and pitted opening as shown in Figure 2.2. After polishing with Simichrome the measured

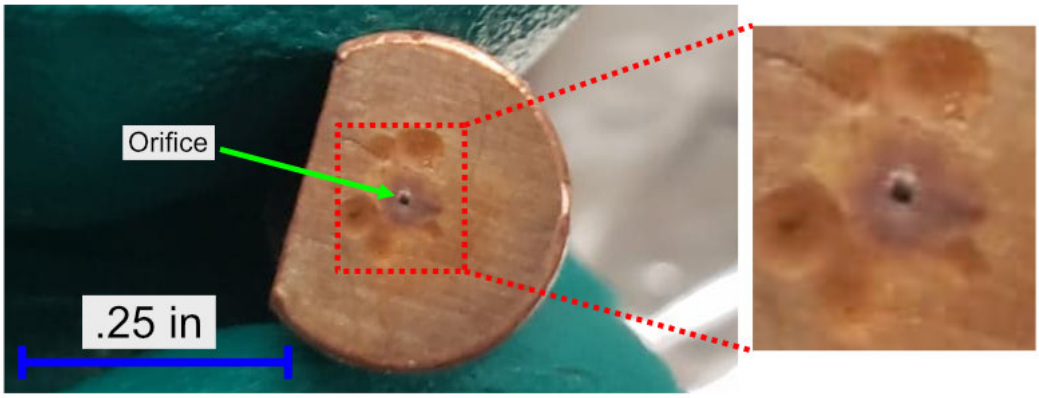

Figure 2.2: Photo of the nicked and pitted front face of the nozzle.

temperature of the focus decreased by 10s of milliKelvin and the location changed by about 18 cm. Inspired by this we also upgraded the He gas delivery system to all stainless steel and added a nitrogen purge system with the goal of reducing impurities entering the expansion. Impurities have been observed to elevate the temperature of a supersonic expansion [72]. Admittedly, it is unlikely that impurities in the He gas would actually make it out of the nozzle. Instead they will likely condense on a surface in the cryogenic portion of the gas manifold.

Adiabatic cooling extends only as far as the continuum flow region. Using the known He-He cross section [70] and density profile of the jet the helium mean-free path at $z = 4$ cm with a flow rate of 50 SCCM is estimated to be 2 mm. Continuum flow therefore likely continues to at least $z = 4$ cm. Thus, according to Equation 2.9, the He gas should reach a temperature of about 1 mK. A similar estimate yields that only a few collisions remain for distances greater than $z = 10$ cm and thus cooling has effectively ceased.

Shock fronts are a common feature of supersonic jets [62, 73]. These form as the supersonic flow, which travels faster than the local speed of sound, is unable to respond to downstream boundary conditions. This results in an over-expansion of the jet, which is compensated for by the formation of shock fronts. These non-isentropic regions add entropy to the jet, heating it and potentially preventing the extraction of the seeded species. When shocks are present, supersonic beam experiments use a skimmer [61] to penetrate through the shock front and extract the central portion of the beam.

The skimmer is located 16 cm from the nozzle with an inlet diameter of 1 in, as shown in Figure 2.3. The rim of the skimmer is sharpened to a knife's edge. At this location the jet is in free molecular flow. The skimmer was placed at this point to avoid skimmer interference effects that can occur with placement closer to the jet and because it was estimated that with the high speed of the charcoal pump no shock would form out to this distance, even at the highest He flow rate. We have seen no evidence of shock fronts.

The capture half angle of the atom lens is about .064 radian. Since this is a relatively small angle, it is important that entrainment is maximized along the centerline of the jet. To facilitate this the nozzle can be displaced and tilted in order to maximize seeding efficiency and alignment to the skimmer. Adjustments are made by three threaded rods that connect the back free flange to the front flange mounted to the chamber. Under vacuum the force on the back flange can cause the bellows to translate in a manner that cannot be corrected by adjustments of the threaded rods. This caused a clear shift of the nozzle from the skimmer centerline.

To compensate this transverse supports were added. One must be extremely careful when making adjustments to not crash the nozzle into any part of the cryogenics. Because adjustments are made against the force of vacuum, it is not possible to feel such a collision occurring.

In order to efficiently capture and extract the seeded lithium, the jet must be of intermediate transverse collisional thickness. If it is too thin, atoms may overshoot the jet, while if it is too thick they may bounce off of it. The transverse collisional thickness varies with distance z as $1/z$, so it is possible to increase or decrease the collisional thickness with a change in the seeding location by adjusting the nozzle. When estimating this location, it is important to take into account the energy dependence of the Li-He cross-section. Fortunately, this cross-section has been calculated to high accuracy [74, 75]. In practice we optimize the seeding by moving the nozzle in or out until a maximum signal is measured.

The first few Li-He collisions will have high relative energies, for which the cross-section is small. This means that the lithium atoms can easily penetrate into the jet even at locations where the helium jet is thick for He-He collisions. After the first few collisions the relative collision energy drops and the cross-section increases substantially. This results in a pronounced increase in the collision rate, and causes the lithium atom to become entrained with the He flow and to approach the He jet temperature. An implication of this is that the highest energy collisions occur in the outskirts of the jet and most energy is deposited and remains there. This limits heating of the center of the jet which could elevate the base temperature of the lithium beam. A simulation of the source discussed in Section 5 suggests that

this is largely true, though a non-negligible amount of energy may enter the center of the jet from the outskirts.

2.3 Cryogenics and Vacuum

A schematic of the cryogenic region is shown in Figure 2.3. Cryogenic components are cooled by a two-stage Cryomech model PT410 Pulse Tube Refrigerator (PTR). The first stage has a cooling capacity of 40 watts at a temperature of 45 K, and the second stage has a cooling capacity of 1 watt at a temperature of 4.2 K. The cryogenic region is surrounded by an 18x18x18 inch copper heat shield thermally connected to the PTR's first stage. This provides isolation of the cryogenic region from room temperature blackbody radiation. The shield is suspended by thin walled stainless steel tubing from the vacuum chamber's ceiling to minimize thermal conduction. Multiple layers of superinsulation are placed both outside and inside this shield to reduce radiative heat loads [76]. Additional insulation was added throughout the apparatus and found to reduce the temperature in the first stage by about 0.5 K. Helium flows into the cryogenic region through a heat exchanger connected to the PTR's first stage, then through a second heat exchanger connected to the second stage, and finally through the nozzle.

PTRs apply cooling power in pulses which caused an unexpected issue with our data analysis. Fluorescence data in the far field (Figure 2.1) would often be excessively noisy beyond what is expected from shot, background, or camera noise. After investigation we determined the cause was the periodic nature of the PTR [77]. During the cycle the cooling power would vary, which resulted in the nozzle

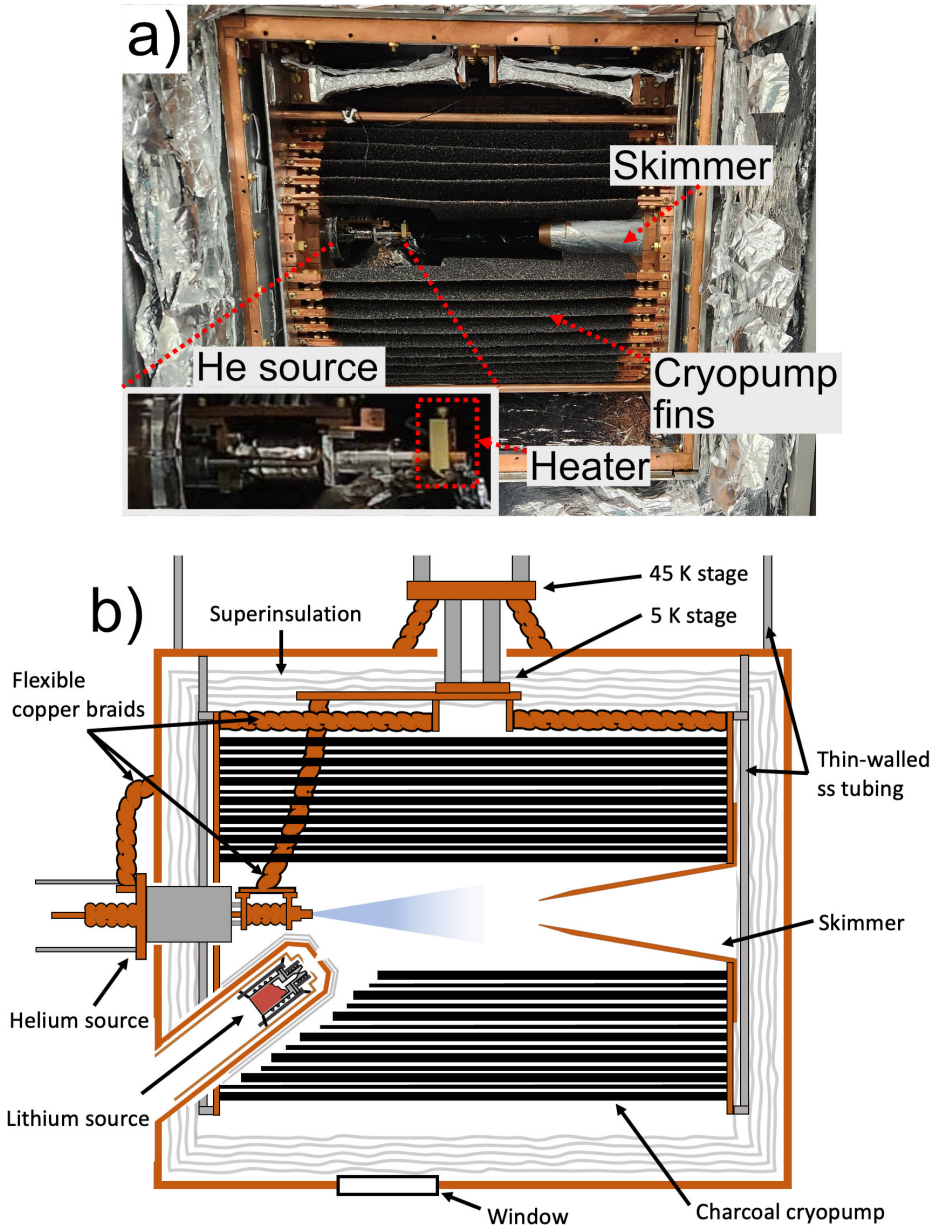

Figure 2.3: (a) Photograph of the cryogenic region. (b) Schematic of the cryogenic region.

temperature varying, and therefore the forward velocity. The effect is small but enough to make data analysis more challenging. The solution was to simply collect data at integer multiples of the period of the PTR, about 700 ms. The discovery of this was somewhat shocking, considering that over 40 lbs of cryogenic copper is present which one would think buffers such an effect.

Thermal connections to the refrigerator cold plates are made by multiple flexible copper braids welded to clamps. Each braid is composed of 2880 36-gauge wires. The braids allow for spatial adjustments to the helium source during operation, and prevent the transfer of excessive strain to the PTR. The first and second stage cold plates are rated to 22 and 11 lbs of transverse force respectively. Unfortunately, during reassembly after the experiment was moved, one of 8 connections to the 45 K cold plate could not be completed without the application of excessive force. To decrease thermal resistance, Apiezon N is applied at all thermal junctions. C10100 copper is used throughout to maintain conduction at cryogenic temperatures [78].

In order to avoid shock front formation and excessive helium background gas pressures, it is necessary to have a very high pumping capacity for the helium gas. This is provided by a toroidal charcoal cryopump. It consists of 60 rectangular 4x16 inch copper fins covered in an epoxied layer of charcoal. These fins are thermally connected to the second stage of the refrigerator, and typically operate at a temperature of 4.0 to 5.5 K. At these temperatures, the charcoal is a highly effective adsorption-based pump for helium gas [60]. The rectangular openings between the fins form an approximately cylindrical pump opening that surrounds the expanding jet. The finned structure provides efficient pumping because most helium atoms

collide with a fin surface that faces another fin. If the helium atom is not adsorbed it will then scatter into the opposite fin. The process repeats until the atom is adsorbed.

Helium atoms have an adsorption energy on charcoal from 80 K to 300 K per atom, varying with the quantity of helium already adsorbed [60]. This is the dominant heat load on the experiment, and sets a limit on the helium flow rate. The pump was designed for a maximum flowrate of 300 SCCM and we have found this to be approximately correct. Flowrates of abut 400 SCCM and higher overload the pumping capabilities.

1.2 to 1.4% of the directed flow of the helium exits the cryogenic region through the skimmer and is pumped away by a CVC PVMS-1000 diffusion pump and a cooled chevron baffle that has an estimated net pump speed for helium of 3,500 L/s. The diffusion pump is backed by a Leybold D90AC rough pump. Oil from the rough pump entering the experiment is a serious concern because it contaminates gauges and increases the base chamber pressure. To help prevent this, argon is bled into the foreline to maintain a pressure of about 100 millitorr, which results in a mean free path of roughly 0.1 mm [79]. Without this the mean free path would be large enough for oil particles to freely travel down the length of of the foreline. A molecular sieve trap containing zeolite 13X is used as well.

To determine the pumping speed of the cryopump we conducted an experiment where we blocked the line of sight between the helium nozzle and skimmer opening. The gas load on the diffusion pump is then dominated by the flow of scattered cryogenic helium background gas through the skimmer opening. Assuming

that the cryopump speed, S_1, is much larger than the skimmer conductance, C_{12}, the pressure in the cryogenic region is

$$P_1 \approx \frac{Q_N}{S_1}, \tag{2.10}$$

where Q_N is the gas flow through the nozzle. The room temperature region ^{4}He pressure is

$$P_2 \approx \frac{P_1 C_{12}}{S_2}, \tag{2.11}$$

where S_2 is the speed of the diffusion pump. In our experiment we measured a linear relationship between P_2 and Q_N of $P_2 = \gamma Q_N$. Solving for S_1 then gives

$$S_1 \approx \frac{Q_N C_{12}}{P_2 S_2} = \frac{C_{12}}{\gamma S_2}, \tag{2.12}$$

The temperature of the ^{4}He gas must be between the 4.4 K and 45 K temperature of the cryogenic region surfaces. Using a value of 15K the skimmer conductance is then $C_{12} = 2.9\sqrt{T/m}D^2 = 36$ L/s where D is the skimmer aperture diameter in cm and m is the molar mass of ^{4}He. Using our measured value of $\gamma = 4.7 \times 10^{-7}$ s/L yields $S_1 \approx 20,000$ L/s. Combining uncertainties in the value of γ and S_2 we estimate this is accurate within a factor of 2.

The estimated background gas pressure and density in the cryogenic region, with our maximum flow rate of 300 SCCM, are 7×10^{-6} torr and 6×10^{-18} m^{-3} respectively. Using the He-He collision cross-section [70] at T = 15 K of $\sigma = 8 \times 10^{-15}$ m^2, the estimated mean free path of a helium atom in the cryogenic region is about 20 cm. With this pressure and mean free path, the jet should transition from

continuum to molecular flow without any formation of a shock front. We have seen no evidence of any effects of shocks in our experiment.

The charcoal has a finite capacity to adsorb helium. We find that we can flow helium continuously at a rate of 50 SCCM for over 12 hours. However as the charcoal fills with helium, the pressure in the cryogenic region begins to increase. We have observed that measured flux decreases with long run times and believe that this elevated pressure is the explanation. As the charcoal continues to fill with helium eventually a runaway condition occurs as elevated pressure promotes thermal conduction which increases charcoal temperature. Eventually the charcoal climbs in temperature by tens of Kelvin and sheds its helium. The experiment can be left running until it is again cooled to its previous temperature, but this takes about 6 hours. The run time can be extended by shutting off the helium flow between data acquisition.

2.4 Lithium Oven

The lithium oven is a two piece design with a separable reservoir and nozzle, both heated independently, as shown schematically in Figures 2.4 and 2.5. Heating is generated by mineral insulated Inconel sheathed heater cables that are wound around and vacuum brazed to the oven. The heater cables are purchased from ARi industries. It is important that the top and bottom can be heated separately so that the top can be held at a higher temperature. This prevents migration of lithium to the top which would, and has, resulted in clogging. The oven can be easily run with only the top heater working, which we were forced to do for about

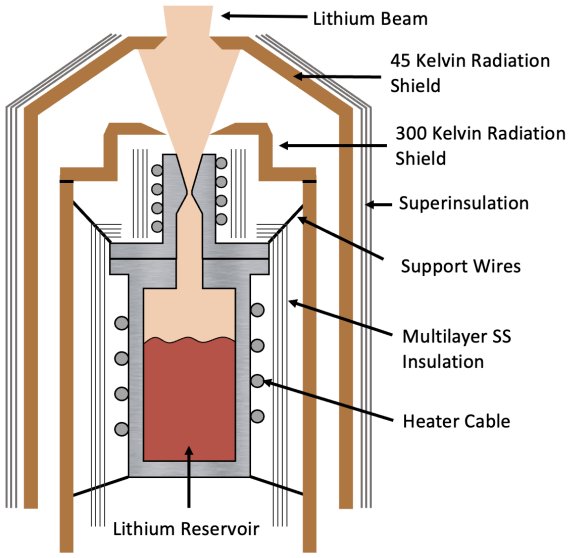

Figure 2.4: Schematic of lithium oven and insulation. The 1050 K oven is insulated from the cryogenics by several methods.

a year after the bottom heater broke after 10 years of age. This turned out to be completely adequate for purposes of characterizing the source, although the total flux was reduced. The delay was quite lengthy and we would recommend having a second oven on hand.

The oven is loaded with about 1.5 grams of lithium wire and can be run at its maximum temperature continuously for several hours. It is constructed of Stainless Steel (SS) and is suspended by SS wires inside a copper cylinder, which is cooled by room temperature water at the base. Multiple layers of thin SS sheets surround the oven providing a degree of radiation insulation. The copper cylinder resides inside an extension into the 45 K shielding. These insulating techniques allow the oven to operate at temperatures up to 1050 K with very minor heating of the cryogenics,

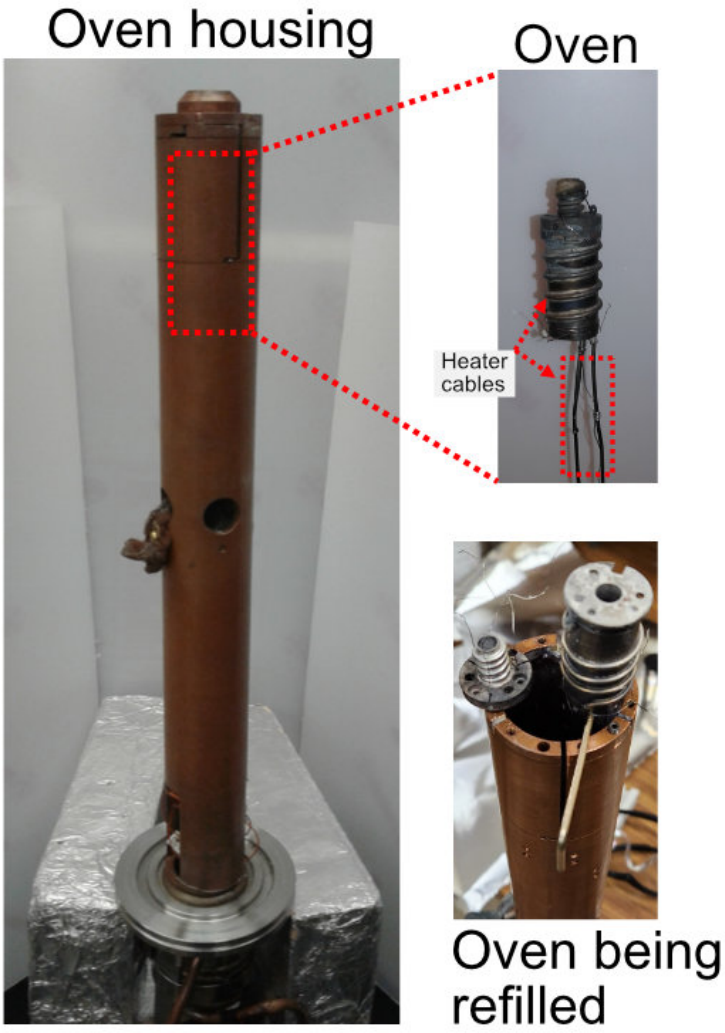

Figure 2.5: Photographs of the lithium oven and the lithium oven housing.

although they are separated by only 5.5 cm. The half angle of the extracted beam from the 45 K shielding is approximately 0.09 radians.

The oven nozzle is a conical design with a 1 mm diameter hole. When heated to 1050 K the Knudsen number is estimated to be 0.14, in the intermediate region between supersonic and effusive. The oven was designed with a shaped nozzle to potentially generate a more forward directed flux, but the extent to which this is occurring in our application is unknown.

Refilling the oven is tedious, requires about 4 hours if everything goes correctly, and presents many opportunities to break components. Care must be taken to not strip any screws that may have seized when the oven was heated. Originally we used boron nitride spray as an anti-seize which performed poorly and resulted in about a 50% chance of a seized screw. We then tried silver plated screws with some improvement. However, it turned out that the boron nitride was expired by a decade and using fresh boron nitride (Momentive Performance Materials Boron Nitride Spray II) worked flawlessly. It is also possible that the different brand was the solution. New socket or flathead head screws are used during every refill to reduce the chance of stripping. 316 SS screws from McMaster-Carr are sufficient. We would originally fill the oven inside a glove bag of argon but found this to be unnecessary as long as the lithium, after two steps of ether rinsing, is transferred to the reservoir quickly. Once in the reservoir it can be protected from oxidation by filling with ether until it is capped.

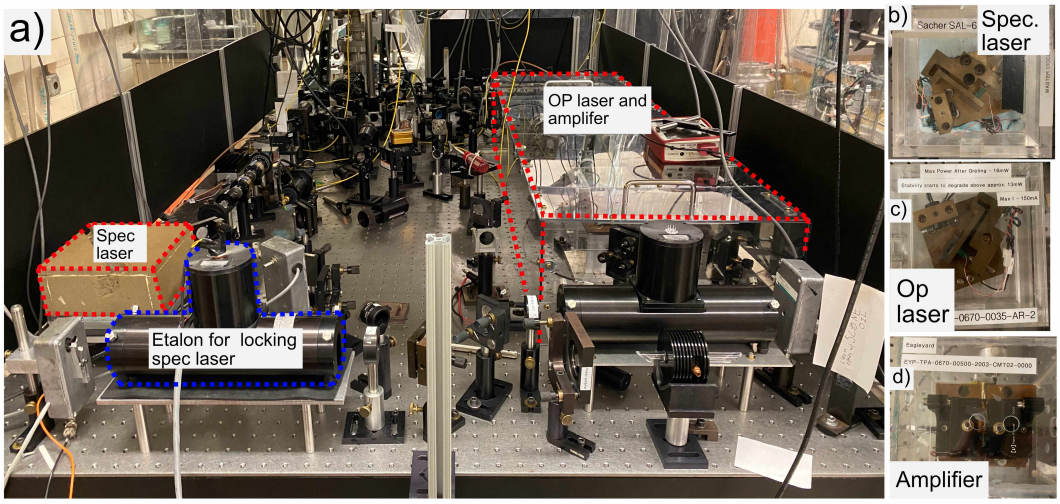

Figure 2.6: (a) Photo of the optics table. (b) The spectroscopy laser. (c) The optical pumping laser. (d) The optical pumping amplifier.

2.5 Laser Optics System

Two lasers are used for spectroscopy and optical pumping operating on the $2S \rightarrow 2P$ transition near 670.8 nm, shown in Figure 2.6. Both are external cavity diode lasers in a Littrow configuration. The spectroscopy laser outputs about 7.5 mW, and the optical pumping laser outputs 500 mW from an Eagleyard tapered amplifier. The spectroscopy laser is typically used by sweeping over a frequency range of about 5 GHz. The laser is frequency locked and controlled by a home-made PI circuit that uses the transmittance through a tunable low finesse Coherent 1906 etalon as the lock point. The frequency is controlled by changing the optical path length of the etalon with an externally applied voltage that modulates a servo controlled glass plate inside the etalon. The original laser-optical system was only the spectroscopy laser without frequency locking and was controlled by directly

changing the piezo voltage.

The spectroscopy laser is calibrated with a separate reference atomic beam produced by an oven housed in a cylinderical chamber. The chamber is pumped to about 1×10^{-6} torr. The oven is typically ran at about 400°C, and the output is collimated by a small slit of width 0.02 inches located 1 inch from the oven output. At this temperature the lithium supply has yet to be exhausted after years of running. A small portion of the spectroscopy laser power is sent through a window into the chamber and retroflected. Retroreflecton eliminates the dependence of the frequency scale on the angle between the laser and the oven output because the frequency of the dip between the two shifted peaks is independent of the angle.

Calibration data is recorded by a PMT. The D2 line of ^{7}Li is used for calibration. The separation between the two ground state ^{7}Li hyperfine levels is 803.504 MHz and is used to set the frequency scaling. The frequency scale is relative to the $2^2S_{1/2}\ F = 2 \rightarrow 2^2P_{3/2}$ transition measured in the reference chamber. Every time fluorescence data is recorded in the experiment the laser is swept over a range that encompasses the D2 line and signal is recorded. During analysis this calibration signal is used to set the frequency scale.

When it eases analysis a repumper laser is added to the spectroscopy laser to prevent unwanted optical pumping. This is done with an Electro-Optic Modulator (EOM) driven at 803.504 MHz, producing two side-bands one of which serves as the repumper. Calibration was originally done with a doppler-free spectroscopy scheme. This was replaced by the above system for the addition of the optical pumping laser.

The optical pumping laser is used to pump atoms in the near field into the

$2^2S_{1/2}$ $|F = 2, m_F = 2\rangle$ state in which atoms experience the ideal magnetic force as explained in Section 2.6. Right hand circular (RHC) polarized light is produced from a Quarter Wave Plate (QWP) and is used to drive σ^+ transitions. The extinction ratio of the QWP through our vacuum windows has been measured to be 100:1 at atmospheric pressure. A vertical magnetic field of 2 G with transverse components less than 20 mG produced by Helmholtz coils is used. Large coils are used to maintain a definite magnetic field direction out to the atom lens. The resulting optical pumping efficiency is around 95%. When performing transverse fluorescence in the far field we attempt to preserve the optical pumping with a similar magnetic field and the use of RHC polarized light.

The laser is locked to the $2^2S_{1/2}$ $F = 2 \rightarrow 2^2P_{1/2}$ $F' = 2$ transition with the reference atomic beam used for calibrating the spectroscopy laser. The locking signal is recorded with a second PMT. The lock is a simple homemade proportional circuit controlling the laser piezo voltage. The maximum angular divergence of atoms that exit the skimmer have a Doppler shift of about 20 MHz. To optically pump all capturable atoms, the laser frequency is broadened. This is done by dithering the current supply of the laser at 800 kHz at an amplitude that widens the laser linewidth to about 25 MHz. Frequency broadening with current modulation has the additional benefit of reducing excess scattered photons as compared to power broadening. This is important at high densities as re-absorption of scattered photons could limit the efficiency of the optical pumping.

Fluorescence spectroscopy in the near field (Figure 2.1) is performed by sending a laser down the center-line of the jet through the skimmer or vertically

through a window in the cryogenic region. The Doppler shift from the longitudinal data is used to determine the atom forward velocity. In the far field, a laser can be used vertically, horizontally transverse, or longitudinally in the observation cell. At certain oven temperatures the beam is sufficiently dense to allow absorption imaging in the far field.

Flourescence data is taken in the experiment with two cameras. The first camera is positioned to image the nozzle region in the near field and the second is located about 176 cm downstream of the nozzle for imaging the focus in the far field (Figure 2.1). A window limits the Field Of View (FOV) of the nozzle to about 4-5 cm past the nozzle depending on its location. In hindsight it would have be beneficial to have extended this FOV to include distances up to the skimmer.

Several different camera brands have been used including Thorlabs, Ximea, FLIR, and FLI. We have found that commercial Complementary Metal Oxide Semiconductor (CMOS) cameras, contrary to their specifications, are actually quite poor choices for imaging when one wants to relate signal to photons. The first reason for this is that these cameras may employ lenslets on the chip surface to help direct light rays into the imaging substrate. This makes it difficult to determine what fraction of photons that enter the camera are actually counted and it is not possible to find data on this effect for many cameras. Second, these cameras may adjust black levels autonomously which has been a source of frustration. The effect can be disabled in some cameras, and in the case of Ximea, required technical support to contact the chip manufacturer. The solution was to send an opaque command to the camera which we were warned could cause other undesired effects though it

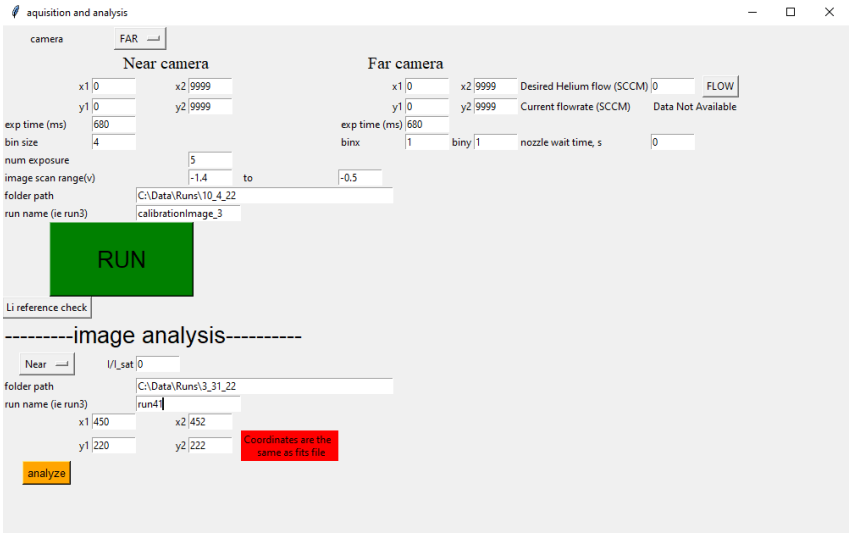

Figure 2.7: An image of the laser/camera control and data analysis GUI.

appears to have worked. In the future it would be best to avoid cameras designed for scanning barcodes and other machine vision applications. The current near field camera is a Ximea brand MC031MG-SY-UB CMOS. The far field camera is a FLI brand Charge Coupled Device (CCD) ML1001 camera meant for scientific or astronomy applications. An advantage of CCD over CMOS cameras is the ability to bin pixels and significantly increase the signal to noise ratio at the cost of losing resolution. This has been instrumental in our experiment.

The cameras and spectroscopy laser are controlled by a homemade Python Graphical User Interface (GUI) shown in Figure 2.7. The cameras are controlled with Python libraries (ximea for the Ximea MC031MG-SY-UB, and pyfli for the FLI ML1001). Voltages and PMT signal are controlled and read with a National Instruments PCI-6251 that is controlled with a Python library (nidaqmx). The GUI is also used to perform temperature analysis with a result shown in Figure 3.2.

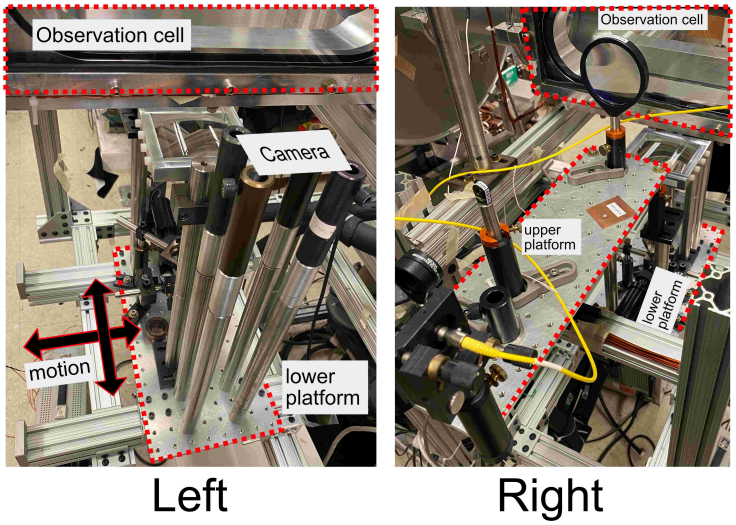

Figure 2.8: Photographs of the far field optomechanic setup. The camera has been removed in this photo. Left and Right refer to the direction when viewed from the nozzle's perspective.

Initial characterization of the far field, primarily finding the focus, was challenging due to an inadequate setup. It took a few iterations to develop a technique that we could use to reliably find the focus in 3 dimensions. The first issue was the camera as previously described. Another issue was the optomechanic setup. We were misusing a optomechanic platform that could move in 3 dimensions over large distance, but would spontaneously shift under the excessive weight. The third issue was we were also using a very narrow laser to induce fluorescence.

The current solution is a sliding upper/lower platform that moves along rails and gliders from the 80/20 brand as shown in Figure 2.8. The lower platform supports a pair of cylindrical lenses that produce a vertically orientated laser with a waist of 2.8(3) cm and 0.59(3) mm in the longitudinal and transverse directions

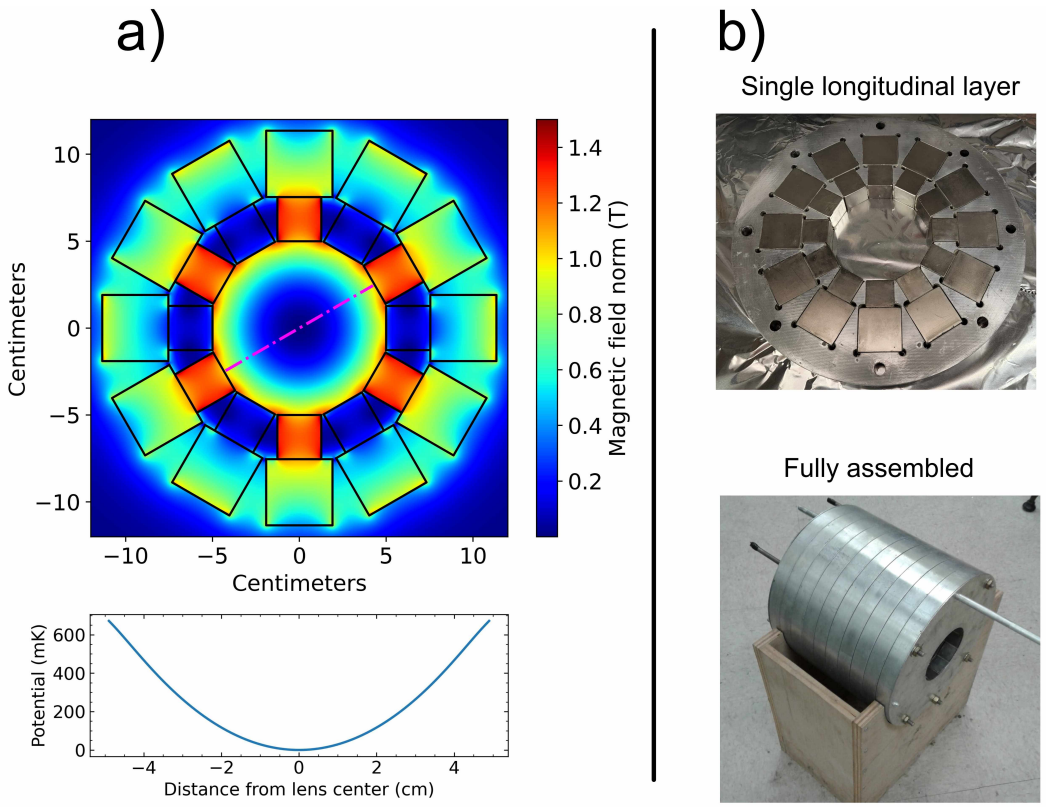

Figure 2.9: (a) Image of the magnetic field and potential of the atom lens. (b) Photograph of a single layer of the lens and the fully assembled lens in a holder.

respectively. The large longitudinal beam waist minimizes the need to move the beam during characterization and eases analysis. The camera resides on the upper platform.

2.6 Magnetic Lens

To collect the expanding lithium beam, a magnetic lens built by an undergraduate student Travis Biles is used. The input face of the magnetic lens is located

72 cm beyond the nozzle. The beam is brought to a focus 176 cm beyond the nozzle for a forward velocity of 207 m/s. The helium atoms are unaffected by the lens and continue expanding. We defined the focus as the point with the minimum FWHM. This choice is explained in Section 3.2.

The lens is a permanent magnet Halbach array hexapole configuration as shown in Figure 2.9. It is formed of two concentric layers with cuboidal magnets of dimensions 1x1x1 in and 1.5x1.5x1 in. Longitudinally each layer is 1 in long. The field strength is about 0.9 tesla at a radius of 4.5 cm, which is approximately the usable aperture before field quality sharply decreases. The inner layer magnet grade is N4017 and the outer layer is N4012. Simulations indicated that the inner grade may be exposed to magnetic fields that could demagnetize the material, so the higher grade N4017 was used there. The original design was 10 layers long, however when installed it was found that the focus was occurring too close to the magnet face to appear in the observation cell. The lens was reduced to 6 layers and the focus appeared in the observation cell. It is worth noting that the process of removing or adding layers is quite tedious and a bit dangerous because of the strength of magnetic fields at play.

For later use we will review the physics of atom-magnetic field interactions and magnetic materials. A neutral alkali atom in a magnetic field has ground state potential energy

$$V = -\frac{\Delta E_{hf}}{2(2I+1)} + g_I \mu_B M B \pm \frac{\Delta E_{hf}}{2}\sqrt{1 + \frac{4M}{2I+1}x + x^2}, \tag{2.13}$$

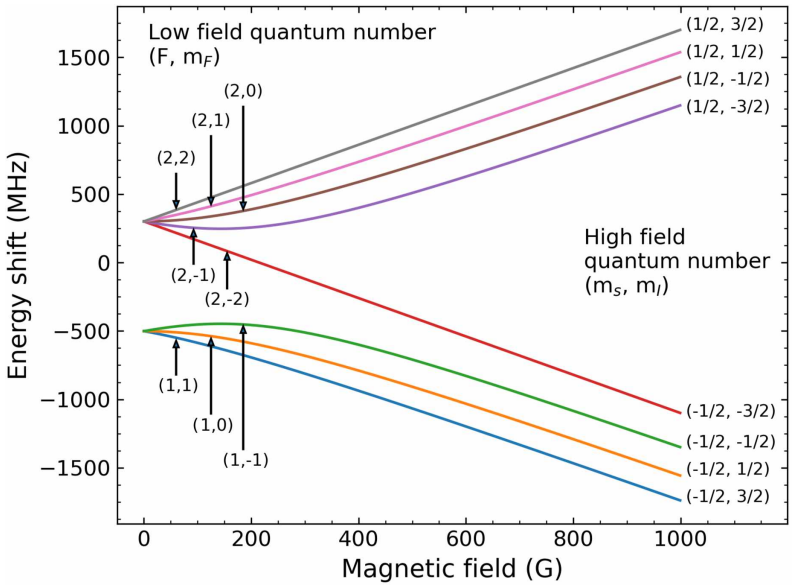

Figure 2.10: Potential energy shift of ground state ^{7}Li atom in a magnetic field. Above a few hundred gauss the atom-field interaction becomes dominated by the electron-field interaction.

where

$$x = (g_J - g_I)\frac{\mu_B B_0}{\Delta E_{hf}}, \tag{2.14}$$

I is the nuclear spin quantum number, J is the total electron angular momentum quantum number, g_i is the nuclear g-factor, g_j is the Lande g-factor, M is the total atomic angular momentum (F) projection quantum number, ΔE_{hf} is the hyperfine splitting at zero magnetic field, μ_B is the Bohr magneton, and B_0 is the norm of the magnetic field [80]. The plus-minus outside the square root corresponds to the sign of the electron spin projection m_s. This is known as the Breit-Rabi equation. A Breit-Rabi diagram for ground state ^{7}Li is shown in Figure 2.10.

Notice that the lines become nearly straight by 500 gauss, while the fields

inside a permanent magnet lens reach values of nearly 10,000 gauss. To first order in a permanent magnetic lens it is reasonable to expand Equation 2.13 to get

$$V \approx \pm \mu_B B_0, \tag{2.15}$$

where g_I was dropped because $g_I \ll g_J$, we neglected terms that do not depend on B_0, and used $g_J \approx 2$. In this regime the interaction between the atom and the magnetic field is dominated by the electron, so the Breit-Rabi diagram simply becomes two sets of parallel lines with the slopes dictated by the z projection of the electron. We optically pump our atoms into the $2^2S_{1/2}$ $|F = 2, m_F = 2\rangle$ state which experiences a linear restoring force everywhere (Figure 2.10). The force on the atom is then simply

$$\vec{F} = -\nabla \left(\pm \mu_B B_0 \right) = \mp \mu_B \nabla B_0. \tag{2.16}$$

Equation 2.16 only holds under the adiabatic approximation which assumes that the magnetic field vector changes direction slowly. If the magnetic field vector changes too fast then an atom may undergo a Majorana spin-flip and transition into a high field seeking state and be lost. This can occur near the vicinity of a magnetic field zero-point where the magnetic field vector changes direction rapidly. The timescale is set by the period of Larmor procession. To avoid this effect, traps will typically have an additional magnetic field of a few gauss added to "plug" the zeros.

To achieve *radial* lensing we need a restoring force that is linear in radial

displacement and thus a magnetic potential second order in radial displacement. To see this recall that a lens works by deflecting rays (or atoms) by an amount proportional to the ray's displacement from the center of the lens. To find the right magnetic field we look to the multipole expansion of 2D magnetic fields. In polar coordinates this is conveniently expressed as [81]

$$\begin{aligned} B_\theta &= \sum_{n=1}^{\infty} r^{n-1} \left(a_n \sin(n\theta) + b_n \cos(n\theta)\right) \\ &\text{and} \\ B_r &= \sum_{n=1}^{\infty} r^{n-1} \left(-a_n \cos(n\theta) + b_n \sin(n\theta)\right), \end{aligned} \tag{2.17}$$

where θ is the polar angle, r is the radius, and a_n and b_n are the amplitudes. $n = 1$ are the dipole terms, $n = 2$ the quadrupole terms, $n = 3$ the hexapole terms, and so on. a_n is known as the "skew" amplitude and b_n the "upright" amplitude. This expansion is used extensively in charged particle accelerator theory as will be discussed in Chapter 4. Equation 2.17 can also be rewritten with a trigonometric identity in a form where the variables a_n and b_n become A_n and ϕ_n where ϕ_n describes the orientation of multipole n and A_n the amplitude.

The magnetic norm of a single multipole is

$$B_n = |B_{rn}\hat{r} + B_{\theta n}\hat{\theta}| = K_n r^{n-1}, \tag{2.18}$$

which is more conveniently expressed by

$$B_n = B_p \left(\frac{r}{r_p}\right)^{n-1}, \tag{2.19}$$

where B_p is B_n at radius r_p, and r_p is typically the magnet bore radius and also known as the pole face radius. With $n = 3$ (hexapole) we get what we were looking for, namely

$$B_0 = B_p \left(\frac{r}{r_p} \right)^2, \tag{2.20}$$

which with Equation 2.16 gives

$$\vec{F} = \mu_B \vec{\nabla}_r B_0 = \mp 2 B_p \mu_B \frac{r}{r_p^2} \hat{r}, \tag{2.21}$$

a linear restoring force for a low field seeking atom and thus provides focusing.

We should stop and consider an important distinction between Equations 2.17 and 2.19. The magnetic multipole expansion is an orthogonal basis and any magnetic field can be expressed by a linear combination of Equation 2.17. Turning this around, we can decompose any desired realistic magnetic field into its multipole components. This property is exploited in accelerator magnet design to eliminate unwanted magnetic fields by applying a canceling multipole term. Unfortunately, Equation 2.19 does not possess the same properties.

A multipole of order n can be produced interior to an annular cylinder if the magnetization in the cylinder follows

$$\hat{M}_n = \cos(\theta(n+1))\hat{x} + \sin(\theta(n+1))\hat{y}, \tag{2.22}$$

which can be approximated by discrete magnets [82] as shown in Figure 2.11. The resulting field is only a pure multipole for infinitely long magnetic material. Truncation introduces fringing fields that cause the multipole to not only be imperfect

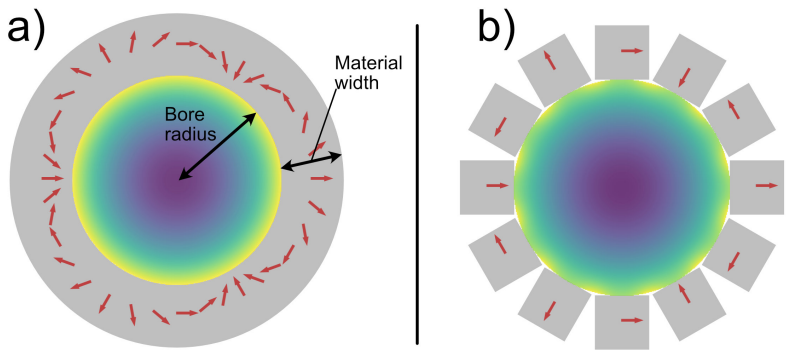

Figure 2.11: (a) The magnetic field norm of a hexapole field produced by an ideally magnetized material. (b) The same field but produced by a discrete approximation to the continuously varying magnetization. Near the edges the break down of the approximation can be seen.

in the 2D plane, but also introduces longitudinal effects.

2.6.1 Permanent Magnets

Neodymium permanent magnets, or NdFeB magnets, are a common choice for constructing Halbach arrays. They produce strong magnetic fields and are relatively cheap and robust. To a decent approximation they are a *hard* magnetic material as opposed to iron which is *soft*. This means that once magnetized they tend to maintain their magnetization even when opposing magnetic fields are introduced.

A common tool for understanding the behavior of a magnetic material is the B-H curve. An example B-H curve for neodymium is shown in Figure 2.12. The "knee" is the region where the hard material properties of the magnet begin to rapidly break down and magnet is at risk of demagnetization. The other terms will

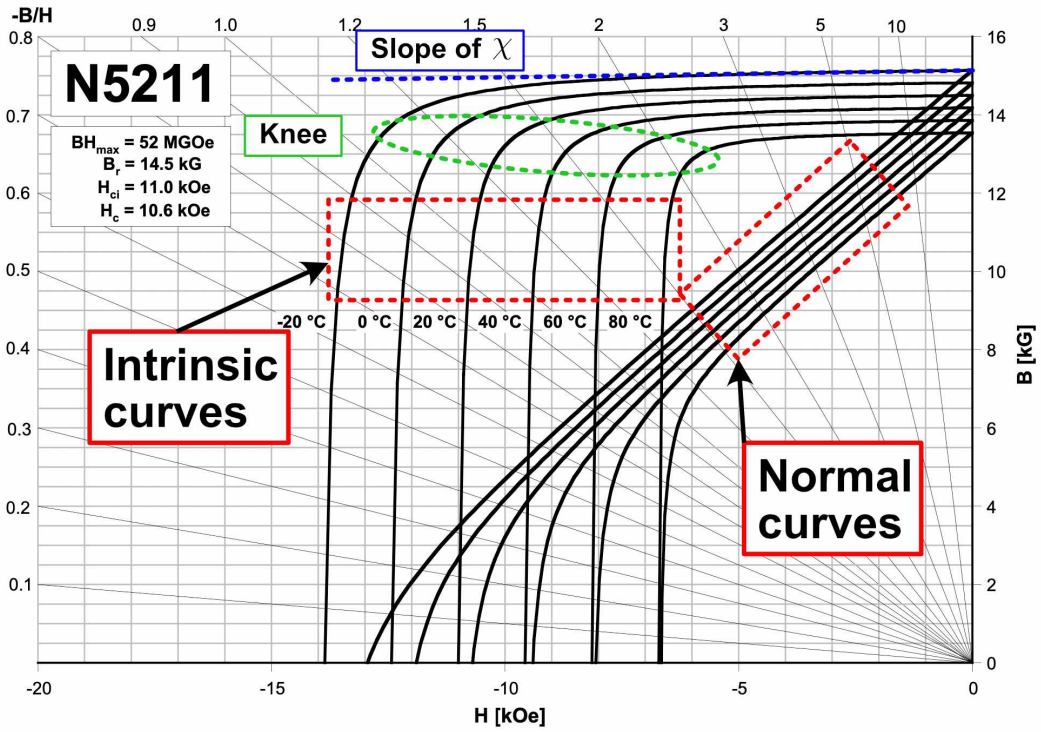

Figure 2.12: B-H curve of N5211 for demagnetizing fields [83]. Both the normal and intrinsic curves are plotted.

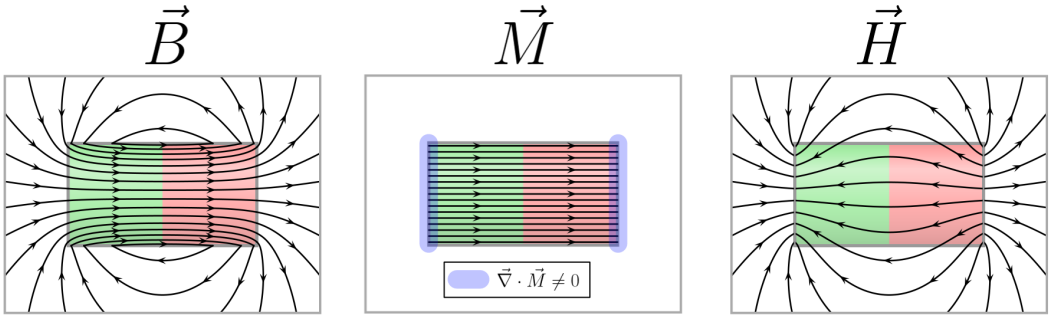

Figure 2.13: Plots of $\vec{B}$, $\vec{M}$, and $\vec{H}$.

be explained shortly.

Recall that the relation between $\vec{B}$ and $\vec{H}$ is

$$\vec{H} \equiv \frac{\vec{B}}{\mu_0} - \vec{M}, \tag{2.23}$$

which with Maxwell's equations give

$$\begin{aligned} &\vec{\nabla} \times \vec{H} = \mu_0 \vec{J}_{\text{free}} \\ &\text{and} \\ &\vec{\nabla} \cdot \vec{H} = -\vec{\nabla} \cdot \vec{M}, \end{aligned} \tag{2.24}$$

where $\vec{J}_{\text{free}}$ is free current and we ignored electric fields. Notice that not only is $\vec{H}$ produced by currents, but it is also produced by discontinuities in $\vec{M}$. A magnetized bar has no free current, but does have $\vec{\nabla} \cdot \vec{M} \neq 0$ at its end. This produces an $\vec{H}$ that is in the opposite direction to $\vec{B}$ and $\vec{M}$ inside the material as shown in Figure 2.13. For a constant $\vec{M}$ this results in reducing the strength of $\vec{B}$. As the magnet's length is increased the average strength of $\vec{H}$ falls. A magnet formed into a toroid will have no discontinuties in magnetization and there will be no $\vec{H}$.

The "normal" curve shown in Figure 2.12 is the total measured magnetic field as a function of $\vec{H}$ while the intrinsic curve is defined as the measured magnetic field without the contribution of $\vec{H}$. From Equation 2.23 this gives for the intrinsic curve

$$\vec{B}_{\text{intrinsic}} = \mu_0 \vec{M}(\vec{H}), \tag{2.25}$$

where we have included that $\vec{M}$ is a function of $\vec{H}$.

The intrinsic curve thus describes the magnetizing behavior of $\vec{H}$. In a neodymium permanent magnet an opposed $\vec{H}$ works to demagnetize the magnet as seen in Figure 2.12. We made no mention of the origin of $\vec{H}$ such as wether it originated from an external coil, another magnet, or from the magnetic material itself. Because the permanent magnet produces its own opposing internal $\vec{H}$ (Figure 2.13), it can partially demagnetize itself. This is known as *self demagnetizing* and depends on the geometry of the magnet.

Away from the knee of the intrinsic curve the slope can be approximate by a straight line with slope χ (Figure 2.12). Thus $\vec{M}$ can be written as

$$\vec{M} \approx \vec{M}_0 + \chi \vec{H}. \tag{2.26}$$

Now if we add $\vec{H}$ back to the intrinsic curve to get the normal curve, then we have over the region away from the knee

$$\vec{B} = \mu_0 \vec{H} + \mu_0 \vec{M} = \mu_0 \vec{H}\,(1+\chi) + \vec{M}_0 = \mu_r \mu_0 \vec{H} + \mu_0 \vec{M}_0, \tag{2.27}$$

where μ_r is the relative permeability. The reader may notice what appears to be some notational abuse. Typically μ_r and χ are defined by $\vec{M} = \chi \vec{H}$ and $\vec{B} = \mu \vec{H}$. This is intentional because we have approximated the B-H curve as a perfect hard material with the additional contribution of a linear soft material.

In permanent magnets this relative permeability is known as the recoil per-

meability. It is a tensor with typical values in neodymium permanent magnets of

$$\mu_r = \begin{bmatrix} 1.05 & 0 & 0 \\ 0 & 1.15 & 0 \\ 0 & 0 & 1.15 \end{bmatrix} \tag{2.28}$$

for magnetization aligned along the x axis [84].

For reference, the effect of self demagnetization of a cube is to reduce the total magnetization by about 1.6% while the difference in magnetization between the two neodymium grades N52 and N50 is about 2.5%. For many applications the recoil permeability is ignored, while attention is only given to remaining away from the knee. However, in a permanent magnet storage ring where particles may make hundreds or thousands of revolutions, this may not be reasonable.

2.6.2 Magnet Imperfections

In reality magnetic material have prescribed tolerances. A permanent magnet has three key tolerance parameters: dimensional deviations, magnetization angle deviations, and magnetization strength deviations. Dimensional deviations are deviations in the size of the magnet. Magnetization strength deviations are deviations in the mean magnetization, $\langle|\vec{M}|\rangle$, of the magnet. The magnetization angle deviation is the tilt of the magnetization vector from the intended direction. Table 2.1 gives the tolerances for high quality magnets (as opposed to novelty grade magnets).

Type	value	unit
Dimensional	$\sim$.002	$\pm$ inches
Magnetization angle	$\sim$3	+ degrees
Magnetization strength	$\sim$1	$\pm$ %

Table 2.1: Standard high quality NdFeB magnet tolerances. This was assembled from vendor spec sheets and private communications. The values can vary and some vendors can provide nearly arbitrary tolerances at greater cost.

Chapter 3

Source Characterization

This chapter reviews recent work on characterizing and improving the existing source. Experimental measurements as well as simulated results are included.

3.1 Near Field

Characterization of the near field sheds light on the seeding dynamics and its efficiency. Performing fluorescence spectroscopy in this region allows us to measure the velocity spread of the lithium atoms as well as the density within the capturable solid angle of the magnet. These results can be compared with simulations and give crucial insight into the conditions that maximize extracted flux.

Transverse and longitudinal spectral profiles for a range of flow rates are shown in Figure 3.1. Laser intensity was below 1 μW/cm^2 for all measurements. Measurements were taken 4 cm beyond the nozzle. The fluorescence is from the $2^2S_{1/2}\ F = 2 \rightarrow 2^2P_{3/2}$ transitions without a repumper and very low laser power to limit optical pumping. The lineshapes are asymmetric, which indicates lithium atoms which still have a relatively high velocity compared to the expected velocity distribution of the jet. The asymmetry in the transverse data indicates there is an unthermalized velocity component upwards through the jet. Within the viewing

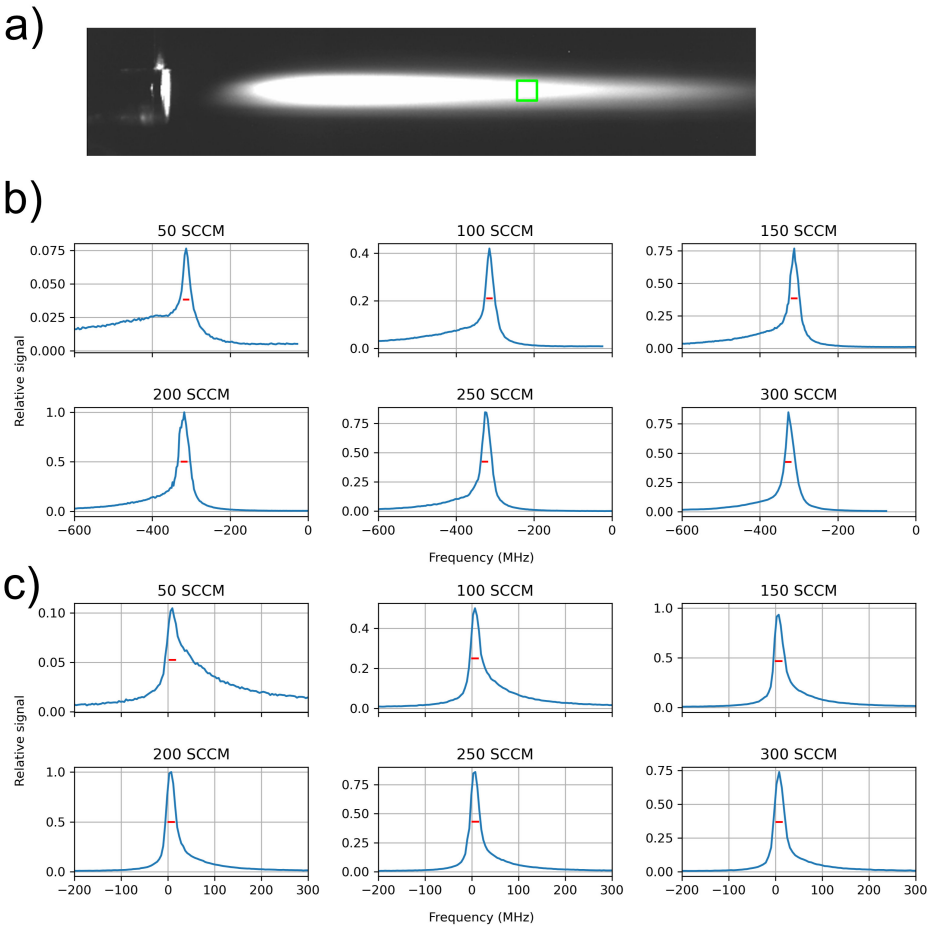

Figure 3.1: Near field fluorescence image and data from transverse laser. (a) Image of longitudinal fluorescence in the near field up to the limit of the FOV. The green square indicates the region that corresponds to the fluorescence profiles and is enlarged for visibility. (b) Longitudinal fluorescence profiles for various flow rates in the near field. (c) Transverse fluorescence spectra in the near field. The red line depicts the natural linewidth (5.9 MHz) at the peak value for reference. Notice for higher flow rates the profile narrows indicating better thermalization at the measurement point.

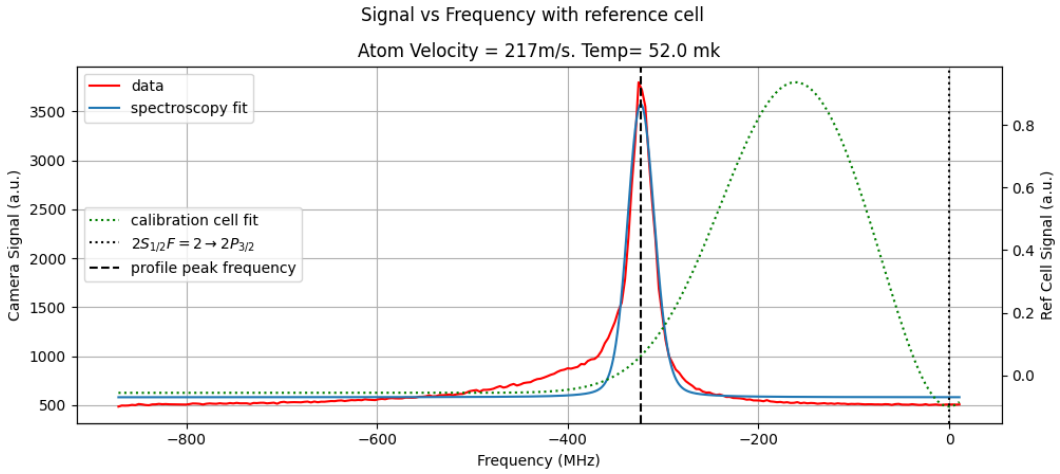

Figure 3.2: Data analysis example applied to 250 SCCM flow for longitudinal data. The fit is to the entire profile. Agreement is reasonable for except for the asymmetry.

region the width of the profile decreases further from the nozzle. Collisions likely continue to occur past 4 cm, the limit of the viewing region, reducing the asymmetry and width of the velocity distribution. As mentioned it would have been beneficial if the FOV were larger so that this collisionality could be measured.

Because the lineshapes are asymmetric their temperature is not well described by a Maxwell-Boltzman distribution. The atoms are not in thermal equilibrium, for which case ascribing a temperature to them is dubious. However, we can consider the full distribution to be a thermalized portion, which results in the sharp central distribution, and an unthermalized distribution resulting in the long tails. Fitting the distribution then gives some measure of the temperature of the thermalized portion.

The width of the distributions for He flowrates above 100 SCCM is larger but

not significantly larger than the natural linewidth. A simple Gaussian distribution fit would then overestimate the temperature. Accounting for the natural linewidth results in a Voigt profile which is the convolution of a Lorentzian and Gaussian. Furthermore, because the hyperfine transitions are not well resolved and overlap a multi-Voigt fit must be used where each transition is appropriately weighted. Laser linewidth broadening, if modeled as a Gaussian, is easily added the multi-Voigt by the convolution property $\sigma_{\text{Voigt}} = \sqrt{\sigma^2_{\text{thermal}} + \sigma^2_{\text{laser}}}$. A similiar property exists for Lorentzian distributions. σ_{laser} is less than a MHz and makes little difference for temperatures above about 10 mK. Applying this analysis to the 250 SCCM longitudinal profile shown Figure 3.1, the narrowest profile, yields a temperature of about 50 mK. There is non-negligible absorption occurring in the longitudinal which further complicates the analysis. The fit and characterization produced by the experiment control and analysis GUI is shown in Figure 3.2. The same analysis yields around 20 mK for the transverse. In the longitudinal at the measured density, some absorption is anticipated which may result in an elevated measured temperature.

Two parameters we can vary to maximize signal are the He flow rate and the seeding distance. The seeding distance is the distance from the nozzle to the centerline of the lithium beam from the oven and is determined by analyzing fluorescence images with zero helium flow. We found that a flow rate of 150 SCCM and a seeding distance of 1.7 cm maximized signal. This contrasts with the far field where the values are 0.7 cm and 50 SCCM. This will be discussed in the next section.

Flow Rate (SCCM)	Measured Density ($\times 10^8$ cm^{-3})	Simulated Density ($\times 10^8$ cm^{-3})
50[a]	2.4(4)	2.3(4)
50	1.1(2)	1.0(2)
100	2.6(5)	2.4(4)
150	4.5(8)	3.2(5)
200	4.0(7)	3.1(5)
250[b]	3.7(7)	2.6(4)
300[b]	3.1(6)	2.1(3)

Table 3.1: Near field average ^{7}Li density within the capturable solid angle of the skimmer compared to simulation. a: Seeding distance of 0.7 cm. b: The nozzle temperature set point must be increased at higher flow rates such that the terminal velocity is 217(2) m/s instead of 210(2) m/s.

The average lithium atom density n over volume V is given by

$$n = \frac{\hbar\omega\Phi(w)}{\sigma(w)} \left(\iiint_V l(\vec{r}) I(\vec{r}) dV \right)^{-1}, \tag{3.1}$$

where $l(\vec{r})$ is the spatial profile of the atomic beam normalized to a height of one, $I(\vec{r})$ is the laser intensity, $\sigma(w)$ is the broadened cross section, and $\Phi(w)$ is the number of scattered photons per second. For greatest signal to noise the resonant frequency w_0 is used. Φ accounts for the collection fraction and efficiencies of the imaging system. The spatial profile is determined by fitting the fluorescence signal assuming cylindrical symmetry. $\sigma(w)$ accounts for frequency broadening from the natural linewidth, laser linewidth, Doppler broadening, and hyperfine structure [80, 85]. This is calculated numerically by

$$\sigma(w) = \zeta \frac{3}{4} \lambda^2 A_{21} s(w), \tag{3.2}$$

where A_{21} is the Einstein rate coefficient, ζ is a factor that depends on the polarization, and $s(w)$ is the *measured* spectral profile with area normalized to 1 Hz.

Average atomic density values for various He flow rates are given in Table 3.1 along with a comparison with the predicted value from a simulation described in Section 3.2.1. The oven temperature was 800 K. The forward velocity and seeding distance are 210(2) m/s and 1.7 cm respectively, except where indicated. To make the comparison, we needed to determine the flux leaving the oven into the near field. This was done by turning off helium flow and sweeping the laser over a wide range while taking images of the fluorescence from the oven. The density profile, and therefore flux, of the oven could then be determined with Equation 3.1. This was measured to be $1.4(2) \times 10^{14}$ s^{-1}.

The average density within the capture angle of the skimmer includes atoms which have not entirely thermalized with the jet. Since atoms with a high relative velocity to the jet are not focused by the hexapole magnet, it is not expected that the entirety of the atoms that escape through the skimmer will survive to the atomic focus.

3.2 Far Field

In the far field the atomic beam is brought to a focus in the observation cell by the hexapole lens. We consider this focus to be the usable portion of our source. Defining exactly what the focus is and how much of it to consider usable is somewhat arbitrary. In terms of the focus there are at least two options: the minimum COLC, or similar concept, and the minimum FWHM. We elected to use

the minimum FWHM because the COLC is not applicable, as previously explained.

Since one of our goals was to use this source to load a storage ring, we looked to the constraints that the storage ring may impose to decide how much of the focus is usable. From Chapter 4 the bore radius of a 1 m bending radius magnet is about 2 cm. With this in mind we decided to define a 1 cm diameter circle at the focus as the usable portion of the source.

This choice is also reasonable in a more general sense. Some authors have used the entire output of their source to compute figures of merit such as flux, but this is misleading because only a fraction of that output can be expected to be used in most situations. A 1 cm diameter circle located tens of centimeters from the magnet defines a reasonable size in both position and momentum space.

Unlike in the near field, the far field signal is maximized for He flowrates of 50 SCCM and a seeding distance of 0.7 cm. The flow rate is lower because the negative effects of pressure in the far field outweigh the positive benefits of the greater density shown in Table 3.1. As will be explained in Section 3.2.1, the smaller seeding distance results in more efficient seeding at this lower flow rate. The pressure issue will be discussed in Section 3.3.

Characterization of the far field is similar to the near field. In the far field we also make use of absorption imaging for more reliable results at higher density values. This is only appreciably present at lithium oven temperatures above 900 K. Absorption imaging is done by sending a $\gtrsim$5 cm diameter horizontal laser through the focus and into the camera. A lens is used before the camera to fit the laser onto the CCD chip. An absorption image is taken in 3 steps. First a dark image is taken

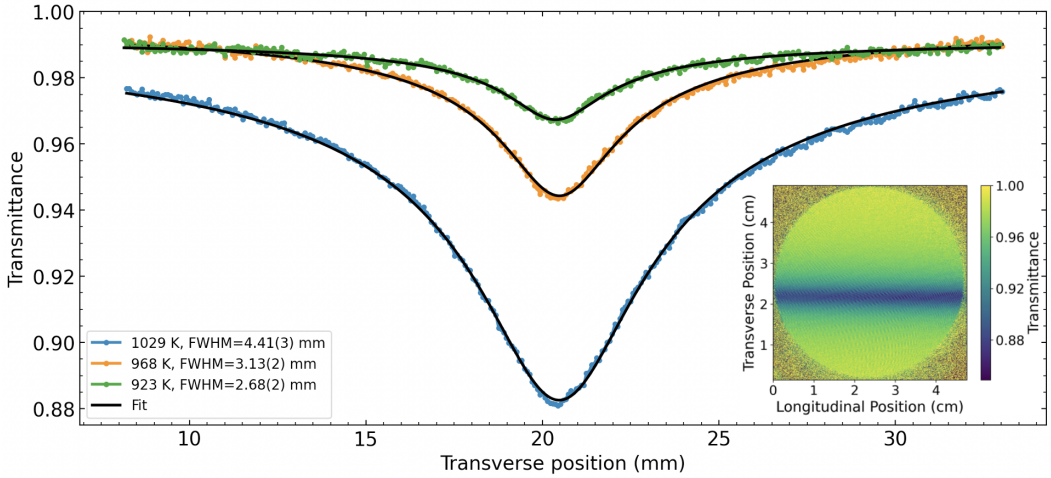

Figure 3.3: Absorption spectroscopy results at the atomic focus. The data is fit to the Abel transform of a Lorentzian profile to determine the FWHM of the atomic beam. In the bottom right corner is an example absorption image.

with the laser blocked and no He flow. Next the laser is unblocked and an image is taken. Finally the He flow is turned on and a third image is taken. The dark image is subtracted from both images and the absorption image is produced by dividing the two images pixel wise.

Figure 3.3 shows absorption profiles from 3 absorption images and an example of an absorption image. A repumper beam is present and the laser power was well below the saturation intensity. To aid in analysis the curves are fit to the Abel transform of a Lorentzian [86]. An Abel transform is used to transform a radial function into its projection from the side, which is exactly what our absorption image is. The Lorentzian was used because of its reasonable fit and simple analytic formula for the Abel transformed version. The radial Lorentzian is $\propto /(r^2+(\Gamma/2)^2)$ and the Abel transformed version is simply $\propto /\sqrt{r^2+(\Gamma/2)^2}$. Absorption is very

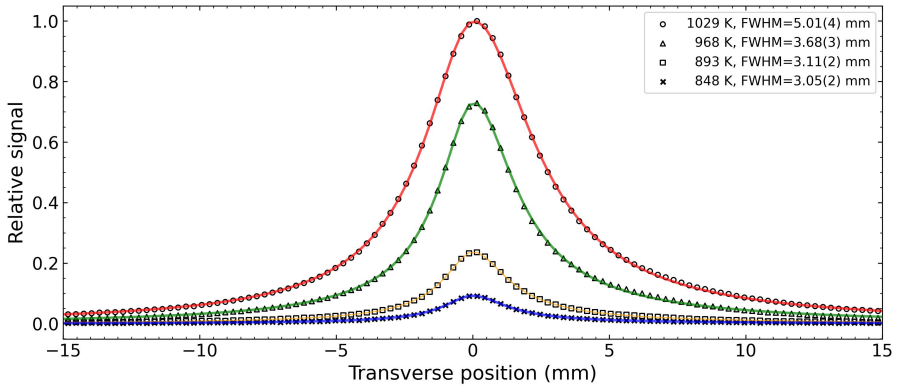

Figure 3.4: Spatial profile at the atomic focus for different oven temperatures fit to a q-Gaussian distribution. The data is from fluorescence spectroscopy and is normalized to the peak signal at the highest oven temperature.

weak off resonance so fluorescence measurements are used to determine temperature and the total broadened cross section.

Figure 3.4 shows similar data but for fluorescence. Here we have found the best fit is given by a q-Gauss [87, 88]. It is a Gaussian-like distribution, but has larger variable-sized tails. A Lorentzian profile also works, but not as well. We have no motivation for using the q-Gauss distribution except that it fits well, is not overly complicated, and has only a few parameters. At maximum oven temperature, around 1030 K, and laser intensity well above saturation the focus is bright enough to be photographed as shown in Figure 3.5.

The FWHM from fluorescence data as a function of distance from the nozzle is given in Figure 3.6 over a range of 14 cm. The focus occurs approximately 176 cm from the nozzle. The temperature of the lithium oven was 820 K, and the beam

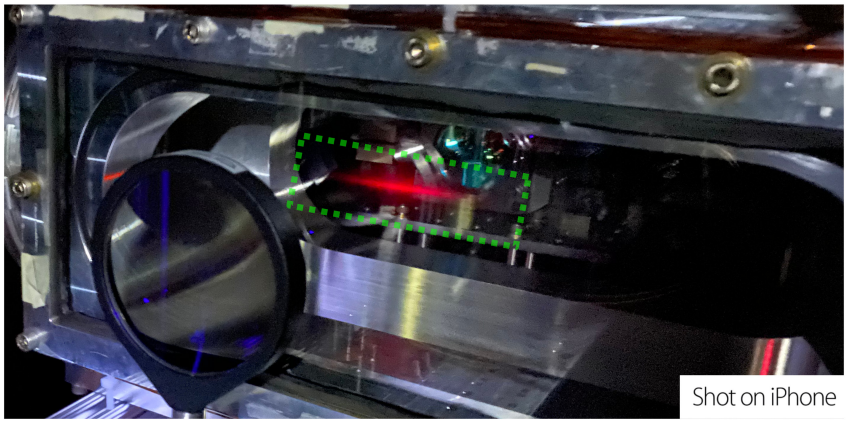

Figure 3.5: A photograph of the focus taken at the highest oven temperature and high laser power. The focus is inside the the green dashed outline and is clearly visible even with a camera flash.

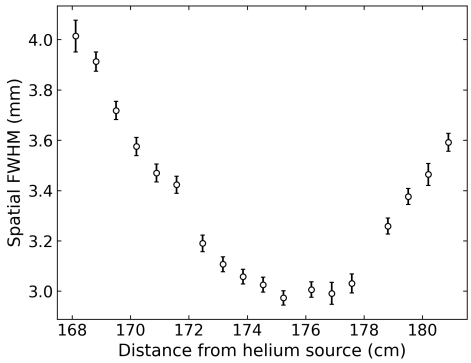

Figure 3.6: FWHM of the focused beam vs distance from the nozzle.

had a forward velocity of 207(3) m/s. The viewing region is about 7 cm, though about 6 cm is usable, so the laser and camera platform had to be moved several times which results in the discontinuities. The beam is fairly flat over a large region which contributed to the difficulty of finding it.

For both fluorescence and absorption imaging, the FWHM increases with increasing oven temperature. The near field spatial profile follows the same trend which suggests that this effect is related to the seeding dynamics. We currently do not have an explanation for the effect. One possible explanation is the additional heat load from the oven because of the increased flux and atom temperature. While the size of the focus increases with oven temperature the location of the focus remains unchanged. The lithium atom temperature appears to increase with oven temperature, but the increase is on the order of our resolution.

If our theory that collisions continue past our observation point in the near field is true then the temperature in the far field should be lower. We can measure this temperature transversely and longitudinally. Quantifying the temperature transversely is somewhat subtle. One method is to measure it using the entire velocity distribution. However, this is not quite correct because that distribution is not an equilibrium distribution and overestimates the temperature of the atoms. For example, consider a 1 μK cloud of atoms launched from the nozzle that travels through the lens and comes to a focus. At the focus the velocity distribution may be significantly more broad than the initial thermal distribution because of the geometric effect of focusing the atoms. Measuring the temperature using the entire distribution would then significantly overestimate the temperature. On the

other hand, there may be times where it is helpful to think of this distribution as characterized by an effective temperature.

Another method of measuring the temperature is to model the effect that the lens has on the velocity profile. This effect can be expressed as the convolution of a geometric profile from the lens with the multi-Voigt profile described in Section 3.1. The geometric profile, not accounting for lens aberrations, is

$$n(v_\perp) \approx \sqrt{1 - (v_\perp / v_{\perp 0})^2}, \tag{3.3}$$

where $v_\perp$ is the transverse velocity and $v_{\perp 0}$ is the maximum transverse velocity exiting the atom lens. Deconvolving the geometric profile from the data results in the original multi-Voigt profile. In practice numeric deconvolution is an ill-poised problem. Instead a convolution of the geometric effect with a parameterized multi-Voigt is used to determine the temperature by curve fitting. Additionally, due to the substantial geometric aberrations in the lens, atoms traveling through the lens at larger radii do not contribute significantly to the flux within the 1 cm diameter circle. To account for this $v_{\perp 0}$ is added as a parameter in the fit.

Longitudinally the temperature is measured with a laser tilted at 13° to avoid depumping the optically pumped atoms and potential slowing effects. Because the beam is tilted, some of the transverse velocity distribution will appear in the longitudinal measurement. To minimize this the temperature is measured 5 cm past the focus.

Figure 3.7 shows far field longitudinal and transverse profiles with the under-

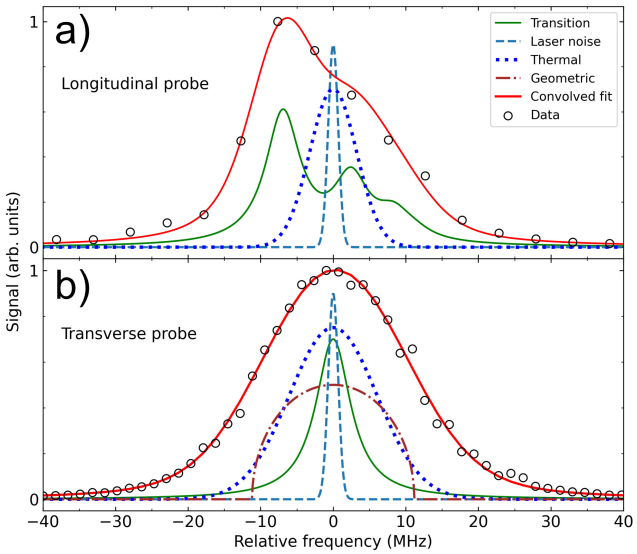

Figure 3.7: Data, fits, and underlying convolving profiles for fluorescence data at the atomic focus. Heights are arbitrary. The transverse and longitudinal temperature can be extracted. (a) Longitudinal fluorescence profile. All 6 hyperfine transitions are present. (b) Transverse fluorescence profiles. Optical pumping is applied resulting in a single transition. The geometric effect of atom lens focusing is included.

lying convolving profiles. The transverse profile was measured with optical pumping in the near and far field with the only allowed transition being $2^2S_{1/2}\,|F=2, m_F=2\rangle \rightarrow 2^2P_{3/2}\,|F=3, m_F=3\rangle$ so only a single Voigt profile is present. The probe laser power was kept below 1 μW/cm^2 such that an atom scattered less than one photon on average to limit unwanted transitions from magnetic field or polarization imperfections. The spectral profile is measured over a 1 cm tall region. The longitudinal laser is scanned over the D2 line and a repumper is present yielding 6 total transitions. The measurement is at the center of the beam over a few pixels.

The measured transverse temperature is about 15 mK, though values up to 20 mK still result in a reasonable fit. The measured longitudinal temperature is about 7 mK and varies by a few mK for different runs that are otherwise the same. These temperatures are much lower than the near field values which substantiates our assumption that further cooling is occurring past our limited near field viewing region. The profiles are also now symmetric. However, there may be some effective cooling from velocity filtering by the lens. From the transverse profile the FWHM of the velocity distribution is 13.4(1) m/s.

3.2.1 Monte-Carlo Simulation of Seeding

Three different simulations have been carried out to model the seeding dynamics. The first two were conducted by the previous graduate student Michael Borysow using a model based on the available theoretical cross sections at the time and assuming isotropic scattering. The third simulation, briefly reviewed here, was developed by colleague Jeremy Glick and made use of newly available fully quan-

Quantity	Fluorescence	Absorption
FWHM (mm)	5.01(4)	4.41(3)
Density^a (cm^{-3})	$3.4(6) \times 10^8$	$2.64(9) \times 10^8$
Intensity^a ($\mathrm{cm}^{-2}\mathrm{s}^{-1}$)	$7(1) \times 10^{12}$	$5.6(2) \times 10^{12}$
Flux ^{b} (s^{-1})	$2.3(4) \times 10^{12}$	$1.54(6) \times 10^{12}$
Brightness^a ($\mathrm{m}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1}$)	$1.8(6) \times 10^{19}$	-
Brilliance^a ($\mathrm{m}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1}$)	$4.3(4) \times 10^{20}$	-
$T_\perp$ (mK)	< 20	-
$T_\parallel$ (mK)	7(3)	-
v_{terminal} (m/s)	210(2)	-

Table 3.2: Parameters of the source at the focus. ^{a} Peak value at focus. ^{b} Value in 1 cm diameter circle at focus

tum mechanical differential cross sections from theoretical calculations [74, 75]. The simulation is written in Python and makes use of the Numba library described in Section 5.6.1.

The simulation works by time stepping lithium atoms through the jet. Lithium atoms originate from the oven with initial phase coordinates sampled from an effusive distribution. At each time step there is a probability of collision with the background He based on the local density, collision cross section, and relative velocity.

An example of some selected simulated trajectories are shown in Figure 3.8. Because small scattering angles are most probable at high relative collision energies, it is possible for particles to undergo many collisions without being deflected into the beam and becoming thermalized. At the other extreme it is possible for a particle to experience many collisions in the jet periphery and bounce off the jet. Both situations are shown in Figure 3.8. This effect can also occur for the lithium beam

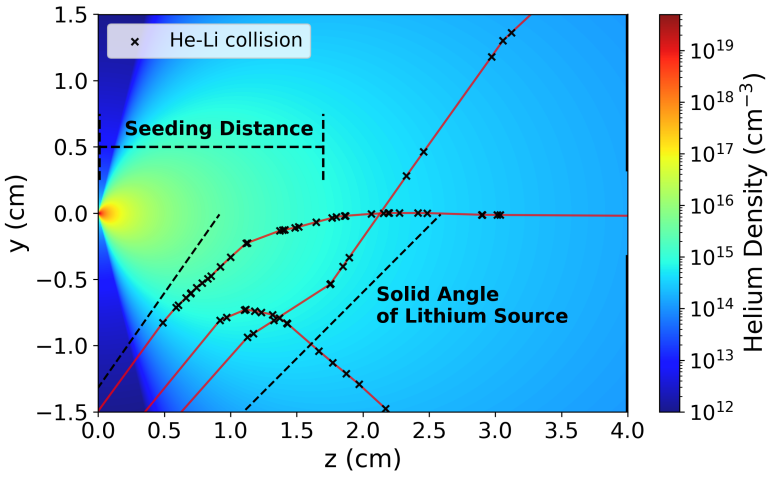

Figure 3.8: Example simulation results for 200 SCCM of He flow. Three distinct types of particles are shown: one which is deflected by the jet, one seeded and within the capturable solid angle of the skimmer and one which undergoes collisions but passes through the jet.

as a whole, as shown in Figure 3.9. For a low flow rate the lithium beam mostly pass through the jet, while for a higher flow rate it mostly bounces off the jet. An intermediate value is optimal.

As mentioned, the optimal He flow rate and seeding distance differ dramatically between the near field and far field. The difference in flow is explained by vacuum considerations, but the seeding distance difference was a source of confusion. To understand the cause, the He flow rate and seeding distance were varied in the simulation. The maximum average density occurs approximately 1.7 cm from the nozzle with a flow rate of 150 SCCM which matches near field observations. With a flowrate of 50 SCCM, a seeding distance of 0.7 cm is optimal instead of 1.7 cm which matches observation. This is shown in Table 3.3.

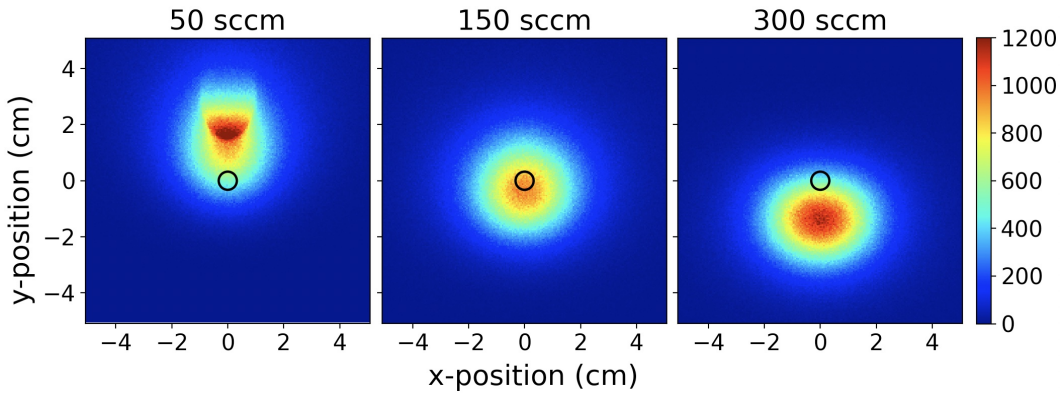

Figure 3.9: 2D histogram of 5×10^6 simulated particles and the positions of particles that pass the xy plane 4.1 cm from the nozzle. The black circle indicates particles that are within the capturable solid angle of the skimmer projected to the xy plane. The number above each plot gives the helium flow rate.

Another source of confusion that was cleared up by the simulation is the asymmetry and long tails of the lithium velocity distribution. Initially it was expected that any particle that makes it out to 4 cm in the jet would be entirely thermalized. However, the simulation predicts that many particles can reach this point without an adequate number of collisions to be fully thermalized. Many particles can also reach this location with large velocities.

This explained another unexpected observation. Transverse fluorescence in the far field results in a spatial distribution that initially follows a Lorentzian, but with much heavier tails over the entire 7 cm viewable region. This made it impossible to define a figure of merit such as Circle Of Least Confusion (COLC). It also made it appear that only a small fraction of thermalized atoms were reaching the focus which significantly disagreed with expected results. We thought perhaps the signal was from the lens performing much worse than anticipated near the bore face and

Flow Rate (SCCM)	50*	50	100	150	200	250	300
Efficiency (%)	0.11	0.04	0.21	0.37	0.43	0.25	0.20

Table 3.3: Simulated results for the percentage of atoms that arrive at the atomic focus from the lithium source. *Seeding distance of 0.7 cm. All other results are for 1.7cm.

failing to focus atoms. We tested this by placing an aperture on the lens input and saw no difference in the size and shape of the tails. We now believe that this effect occurs from the fairly large population of lithium atoms which are not thermalized yet escape through the skimmer. Their high velocities prevent focusing and they appear as a nearly constant background signal.

3.2.2 Modeling the Focus

The jet seeding simulation can be combined with a particle tracing simulation, described in Chapter 5, to predict the shape and location of the atomic beam focus. To do this the jet seeding simulation is extended to 10 cm past the nozzle to account for further ^{4}He-^{7}Li collisions that are expected to occur. Beyond this distance and out to the magnet it is estimated that there is less than one collision on average per lithium atom for a helium flow of 50 SCCM. The phase space coordinates of the atoms at 10 cm from the nozzle are used as initial conditions for particle tracing through the lens.

This model predicts a focus location that is very close to the observed location but about half the observed FWHM. The disagreement is still not well understood, but one possible explanation is heating in the jet. The seeding simulation

does not account for heating but a simple model can be used. In this model the base temperature of adiabatic cooling is limited to a minimum value. The results of this model and the adiabatic model are shown in Figure 3.10. There is much better agreement for a temperature of 3 mK, but there is no physical motivation to choose this value although it is not far off from the measured temperature. For now we can say that small amounts of heating is detrimental to the focus and possibly explains the difference in FWHM.

The effect of heating is explained by a virtual source [89], an imaginary source of atoms located at the nozzle plane that would produce the same spatial and velocity distribution at a given distance from the nozzle neglecting collisions. The phase space distribution of the virtual source is produced by projecting atoms back to the plane of the nozzle. Higher temperatures produce larger velocity distributions which in turn result in a larger virtual source size. With the virtual source as the object, our lens then produces a larger image for higher jet temperatures.

In addition to the location and FWHM of the focus, the model can be used to predict the seeding efficiency of atoms arriving within a 1 cm diameter circle at the focus. These results are presented in Table 3.3. Accurate measurements of the flux leaving the oven shielding are performed at lower temperatures to avoid absorption effects. At an oven temperature of 800 K, we measure a flux of $1.4(2) \times 10^{14}$ s^{-1} leaving the oven shielding, while a flux of $3.8(7) \times 10^{10}$ s^{-1} is measured at the focus. This gives a total efficiency of 0.03% and accounting for a 2.5 times loss from background pressure, discussed below, agrees within 40% of the simulated results.

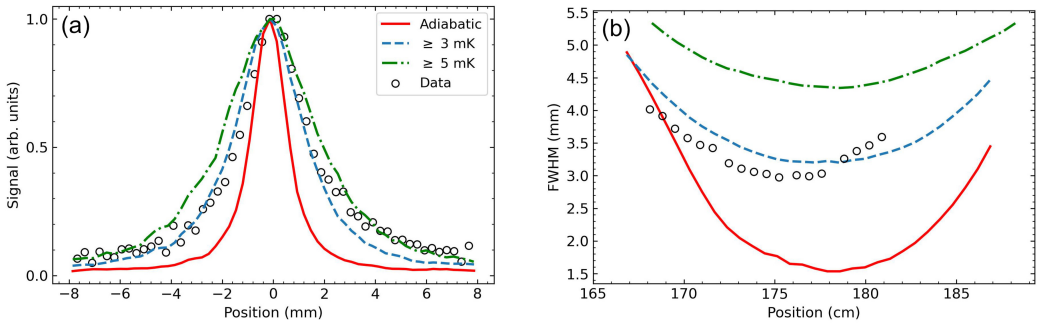

Figure 3.10: Simulated focus results versus flourescence data. (a) Measured transverse profile vs simulated transverse profile. (b) FWHM versus distance from nozzle for data and simulation. In the legend adiabatic refers to a pure adiabatic expansion, while 3 and 5 mK refer to introducing a temperature floor in the expansion.

3.3 Effects of Background Gas Pressure

A major design concern is sufficient vacuum pumping to reduce scattering between lithium atoms and background gas helium atoms to negligible levels. We tested whether this condition was satisfied with a simple experiment. While flowing helium through the nozzle at 50 SCCM we introduced additional helium into the far field through a valve and recorded the signal at the focus. If collisions are happening they should be explained by Beer's law

$$\frac{dT}{dz} = n\sigma T \rightarrow T = T_0 e^{-n\sigma L}, \tag{3.4}$$

where n is the density of background helium, σ is the He-Li cross section that moves a lithium atom out of the focus, L is the distance from the skimmer aperture through to the focus, and T is the transmittance. This depends on a roughly constant pressure throughout the far field, which was validated with MolFlow+ [90]. The

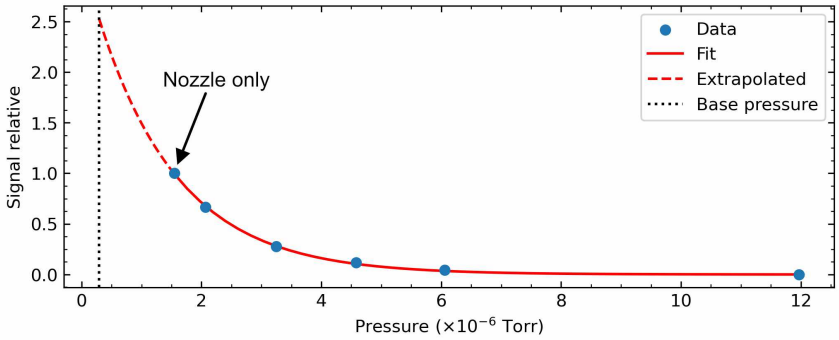

Figure 3.11: Relative signal at the focus versus chamber pressure. The first data point is with on nozzle flow and the following data point include helium metered into the chamber.

distance from the nozzle to the focus is not used because the ratio of cryopump speed to skimmer conductance is so great that there should be hardly any pressure increase in the near field. Replacing n with the ideal gas law, collecting constants besides He pressure P_{He} into β, and replacing the transmittance with signal S gives

$$S = S_0 e^{-\beta P_{\mathrm{He}}}. \tag{3.5}$$

The results of measurement and fit are given in Figure 3.11. The total chamber pressure is fit instead of He partial pressure, but the results are identical as long as the base chamber pressure is used as reference rather than 0 pressure. The base chamber pressure includes almost no helium partial pressure.

Our measurements show that even at our lowest operating flow rate of 50 SCCM, about 60% of the lithium atoms are lost. The flux of our source would be about 2.5 times higher at 50 SCCM if the far field pressure could be reduced by an order of magnitude. If this low pressure were maintaned up to the maximum

seeding efficiency flowrate of 200 SCCM, then from Table 3.3 the total flux increase would be about $2.5 \times (0.43/0.11) \approx 10$ times.

3.4 Atom Lens

The performance of the atomic focus is ultimately limited by the performance of the atom lens. In the ideal case the source, or the object, would be a point source expansion with zero temperature. An ideal lens would then image another point source of zero temperature and with an angular divergence set by the magnification. In reality this would not occur because of the following aberrations:

1. Magnetization deviating from the ideal relation, Equation 2.22, because of the use of discrete magnets which reduces the magnetic field purity.

2. Truncating the lens produces fringing fields, both inside and out of the lens, reducing the field purity and introducing longitudinal forces.

3. Permanent magnets deviating from the ideal hard material approximation. Note that in the following sections the soft material nature of permanent magnets is accounted for by default in the Halbach array lens. The aberration will be explored by disabling this feature.

4. Atoms must slow down when entering the lens to conserve energy as they climb the initial potential. They then speed up as they leave. This does not occur by the same amount for all atoms.

5. If all atoms have the same speed, then there is a distribution of longitudinal velocities $\propto \cos\theta$ resulting in some atoms spending more time in the lens.

6. Material imperfections introduce additional multipoles and reduce field purity.

7. Assembly misalignments of individual magnets introduce additional multipoles and reduce field purity.

8. Gravity deflects atoms away from the optical axis.

3.4.1 Ideal Lens

We will first use an ideal lens model to understand the impact of the some of the lens aberrations as well as the relation between temperature and spot size. The aberrations we will explore are 4, 5, and 8. An ideal lens will give the best case scenario of these effects without the impact of other aberrations. The lens is modeled as a purely transverse restoring force that is only applied within the bore. The parameters of the ideal lens are chosen to most match the values of our existing lens. To review, the existing lens is 72 cm from the nozzle, has a bore radius of 5 cm, a length of 6 in and a pole face field strength of about 1 tesla. The ideal lens is matched by varying the pole face strength to achieve the same focus. This yields a value of 1.0145 tesla. The focus is defined to be the point with the minimum D_{90}, which from simulations is 188.9 cm past the nozzle.

This appears to conflict with our previous statement that the focus is 176 cm past the nozzle. However, we experimentally defined our focus to be the minimum FWHM, which is not at the same location as the minimum D_{90}. There are two

Aberration	Location (cm)	D_{90} (mm)	Note
None	188.9	<.001	
Gravity (8)	188.9	<.001	Vertical shift of -0.209 mm
Velocity distribution (5)	188.8	0.2181	
Energy conservation (4)	186.2	1.125	

Table 3.4: Focus location and size for various aberrations in an ideal lens.

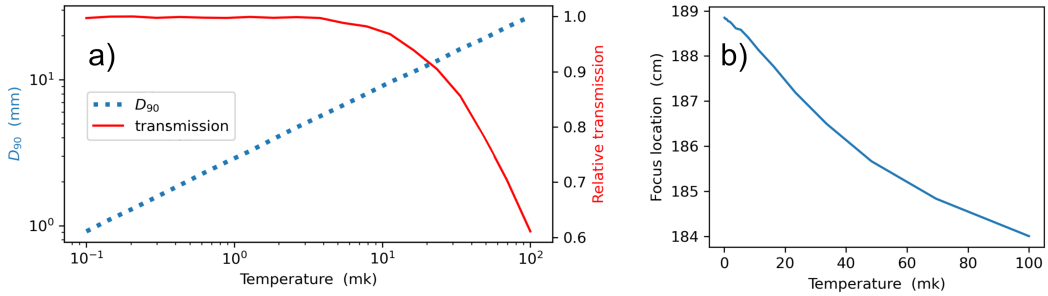

Figure 3.12: Focusing behavior of an ideal lens vs particle temperature. Particles originate from a point source at the nozzle. (a) Focus size and transmittance vs temperature. Transmittance is relative to the transmittance at zero temperature. (b) Focus location vs temperature.

reasons for this discrepancy. First, the measured FWHM was for a flowrate of 207 m/s, while simulations here are done for 210 m/s. Second, the seeding is not as simple as a point source and the object is not necessarily positioned at the nozzle.

Results for the aberrations are shown in Table 3.4. The effect of each aberration is calculated with the other aberrations disabled. The energy conservation aberration was introduced by adjusting the longitudinal velocity of a particle when entering and exiting the ideal lens based on energy conservation. Unsurprisingly the effect of introducing gravity is to simply shift the focus.

Let us also consider the effects of imaging a point source with non zero

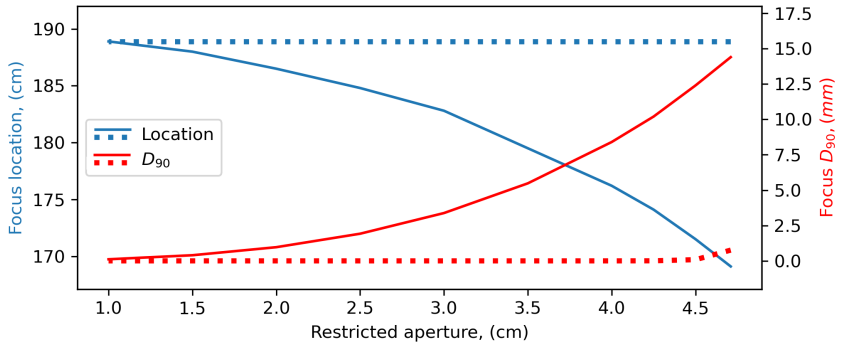

Figure 3.13: Performance of an ideally magnetized lens vs restricted aperture. The solid line indicates the results from the full lens. The dashed lines indicate the results when using only an inner plane free of fringe field effects. The small uptick of the inner plane model occurs because the ideal magnetization model breaks down close to the bore face.

temperature. To include the effects of temperature, the point source's velocity components are perturbed by sampling from a Maxwell-Boltzmann velocity distribution. Results are shown in Figure 3.12. The relative transmission decreases with increasing temperature indicating that some velocity filtering is occurring. The temperature of our source is a minimum of about 5 mK longitudinally and a maximum of about 20 mK transversely which corresponds to a D_{90} of about 6-12 mm. Again, this neglects any seeding dynamics.

3.4.2 Ideally Magnetized Lens

Next we will consider a lens of continuous magnetic material, modeled by 48 wedges, with magnetization according to Equation 2.22 (see Figure 2.11). This is the ideal case of a physically realistic hexapole lens. Similar to the previous section where we tuned the magnetic field of the ideal lens to match our current lens, here

we tune the width of the magnetic material. This gives a width of 0.9963 inches. Here we will investigate the truncation aberration 2. We will disable aberrations 5 and 8. Aberrations 1, 3, 6, and 7 are not applicable.

First the full lens is characterized for a baseline. It is common in optics to restrict the aperture of a system to reduce the aberrations, so we will apply the same concept here. Note that no material dimensions of the lens are changing, only a restriction to the bore aperture for particles is being applied. Next we will remove the truncation aberration by modeling the force from the lens as coming from the central plane of a very long lens where the effects of fringe fields have fallen to zero. The distance over which the force is applied must be tuned to reproduce the same focus location. This gives a length of about 15.10 cm. Results are shown in Figure 3.13. The effects of fringe fields dominated the previously explored aberrations and are comparable to a temperature of 30 mK at full aperture.

3.4.3 Halbach Array Lens

Finally we consider a Halbach array lens with the same parameters as our existing lens. Similar to the previous section we will analyze the lens by characterizing the baseline and then a truncated model. The field region length was reduced to 15.105 cm. Results are shown in Figure 3.14. Next we will consider the effects of material imperfections, deviations from the hard material model, and assembly tolerances (6,3,7). As mentioned, the soft material nature of permanent magnets has already been included in the Halbach lens model data in Figure 3.14. Using the hard material model shifts the focus about 5 cm closer to the magnet and size

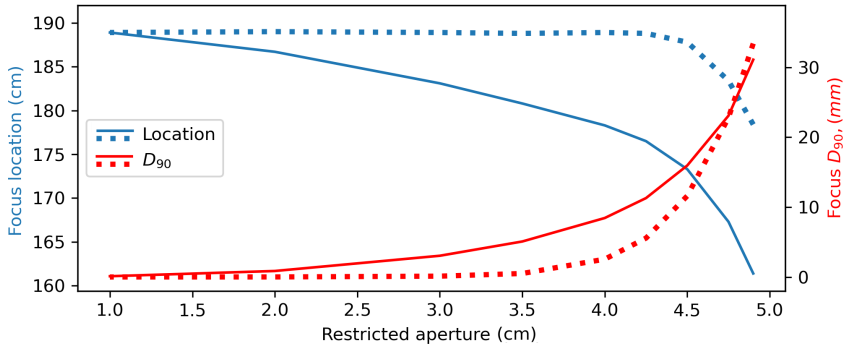

Figure 3.14: Performance of a Halbach array lens vs restricted aperture. The solid line indicates the results from the full lens. The dashed lines indicate the results when using only an inner plane free of fringe field effects.

of the focus changes about in proportion to the magnification. The majority of the difference comes from the bulk self-demagnetization effect, explained in Section 5.3. Meshing each cuboid into smaller domains and accounting for magnet-magnet interactions is a small correction that shifts the focus by a few additional mm.

Results from material imperfections using the tolerances specified in Section 2.1 are shown in Table 3.5. Imperfections are randomly applied to each cuboid. Results from assembly misalignment are shown in Table 3.6 for several aperture sizes. Translational assembly tolerance is $\pm$ 250 μm and rotational assembly tolerance is $\pm$ 0.5°. The model is not entirely physical because magnets can overlap, but the overlap is very small.

3.4.4 Analysis

Without fringe fields the Halbach array and ideally magnetized lens have very similar performance up to around 3-4 cm. After that the aberrations introduced by

Aperture (cm)	Location, error (cm)	D_{90}, error (mm)
4	178.1 ± 0.13 — [178.3]	8.33 ± 0.015 — [8.44]
3	182.6 ± 0.20 — [183.1]	3.12 ± 0.083 — [3.04]
2	186.7 ± 0.26 — [186.7]	1.49 ± 0.30 — [0.86]
1	187.5 ± 8.6 — [188.9]	2.93 ± 0.96 — [0.11]

Table 3.5: Mean and RMS values for focus location and size for standard neodymium permanent magnet imperfections. Quantities in square brackets ([]) are the values without imperfections.

Aperture (cm)	Location, error (cm)	D_{90}, error (mm)
4	178.1 ± 0.21 — [178.3]	8.35 ± 0.058 — [8.44]
3	182.8 ± 0.18 — [183.1]	3.03 ± 0.017 — [3.04]
2	186.5 ± 0.27 — [186.7]	0.934 ± 0.086 — [0.86]
1	188.6 ± 0.90 — [188.9]	0.842 ± 0.48 — [0.11]

Table 3.6: Mean and RMS values for focus location and size for alignment errors. Translational tolerance is ± 250 μm and rotational is ± 0.5°. Quantities in square brackets ([]) are the values without imperfections.

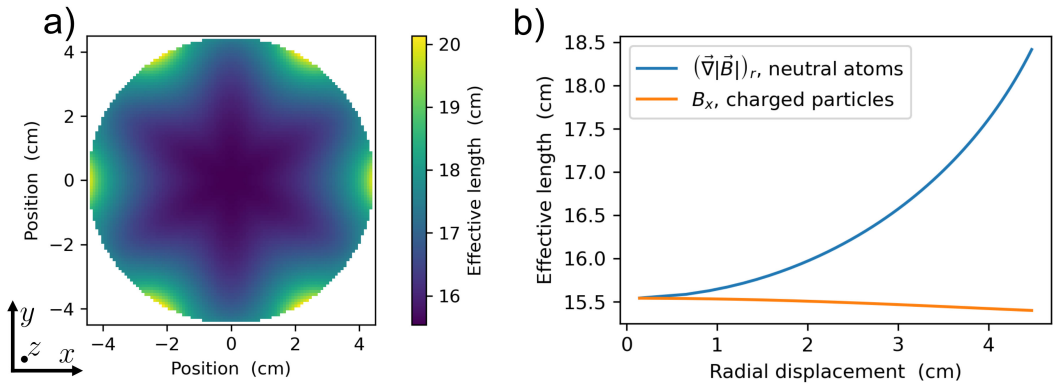

Figure 3.15: The effective length of our Halbach array lens. (a) A heatmap of the effective length across the bore radius. (b) The azimuthally averaged effective length as a function of bore radius. The same result is given for charged particles.

discrete magnets begin to become substantial. Not surprisingly, with fringe fields removed the ideally magnetized lens is effectively an ideal lens. Artificially disabling longitudinal forces has little effect on results so the fringe field's effect on focusing is primarily by reducing transverse field purity.

The effects of fringing fields are analogous to spherical aberrations in light optics, which result in greater focusing power for larger apertures. In a magnetic lens this is explained by a concept in accelerator physics known as *effective length* [91,92]. As its name implies, it is the length of the component that the particle effectively experiences from the field forces. Fields that extend past the input/output of a component will increase the effective length as long as the fields inside aren't reduced by a compensating amount.

The effective length is given by

$$L_{\text{eff}}(x,y) = \frac{\int g(x,y,z)dz}{g_0(x,y,z_{\text{center}})}, \tag{3.6}$$

where g is a transverse field quantity, g_0 is the value of g in the center of the component, x and y are transverse coordinates inside the component's bore, and z is the distance longitudinally through the component. The effective length is measured at some transverse x and y with the integral performed along z from $-\infty$ to $+\infty$ (in reality a sufficiently long distance). For a neutral atom lens $g(x,y,z)$ equals the radial component of $\vec{\nabla}B_0(x,y,z)$. For a charged particle g is proportion to the relevant components of $\vec{B}$.

For an ideal lens the effective length is constant for all valid x and y and is equal to the length of the lens. An effective length that increases with displacement from the bore centerline would signify spherical aberrations because particles experience focusing forces over a greater length.

The effective length is shown for the Halbach array lens in Figure 3.15. The effective length's strong radial dependence is the explanation for the lens' spherical aberrations. The field value B_x does not have the same dependence however. It can be shown that quantities proportional to $\vec{B}$ have this feature in multipole fields produced by magnetized material [82]. For a Halbach array the relation does not exactly hold and thus the slight tilt. Unfortunately quantities such as $\vec{\nabla}B_0$ do not have the same property.

For smaller apertures material imperfections and misalignments are domi-

nant. However, in Section 6.7.2 we will see that this is likely a significant overestimate of the impact of imperfections and is the worst case scenario assuming the proper assembly technique and vendor is chosen. Unfortunately, for our existing lens this was not done. The assembly tolerances are also rather loose and likely an overestimate. Fortunately, these effects largely wash out when using the entire lens because the poorly focused atoms originating from near the lens center is a small fraction of the total flux. For example, the lens region of radius 1-4 cm contains $(4^2 - 1^2)/1^2 = 15$ times as many particles as the region of radius 0-1 cm.

The impact of imperfections can be understood by an analysis of the multipole spectrum. Imperfections introduce unwanted multipole terms which have a larger impact near the lens center where the field values are smaller. For example, consider the trivial case of simply shifting the lens vertically by δy. The new radial force for $r \gg \Delta y$ is

$$F_r \approx -Kr(1 + \sin(\theta)\,\Delta y/r) \approx -Kr, \tag{3.7}$$

while for $r \ll \Delta z$ it is

$$F_r \approx -K\Delta y(1 + r/\Delta y) \approx -K\Delta y, \tag{3.8}$$

where $K = 2\mu_B B_p / r_p^2$.

We can analyze the multipole spectrum of the magnetic field to determine the impact of material imperfections. From Equation 2.17 using the orthogonality of multipoles and considering only skew terms and radial components we get

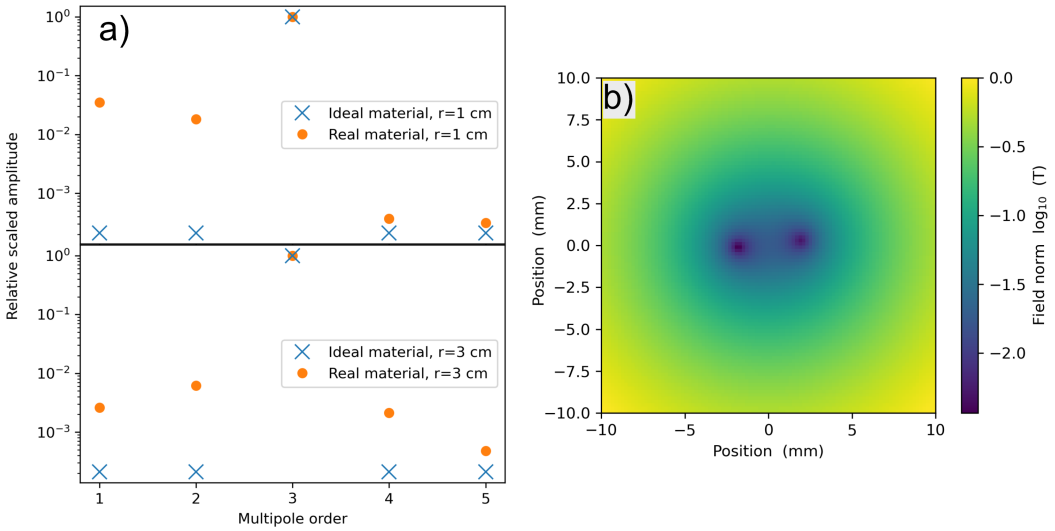

Figure 3.16: (a) Multipole spectrum of the atom lens with material imperfections at a radius of 1 and 3 cm. (b) $\log_{10}$ heatmap of an example of the difference between using perfect and imperfect magnets. Different random errors result in qualitatively similar plots.

$$
\begin{aligned}
B_r &= -\sum_n a_n r^{n-1} \cos(n\theta) \\
-\oint B_r a_k r^{k-1} \cos(n\theta) d\theta &= \oint a_k^2 r^{2(k-1)} \cos(k\theta)^2 d\theta \\
-\oint B_r \cos(k\theta) d\theta &= a_k r^{k-1} \pi \\
a_k &= \frac{-\oint B_r \cos(k\theta) d\theta}{r^{k-1}\pi},
\end{aligned}
\tag{3.9}
$$

where the integral is taken around a circle of radius r. Upright terms must be considered as well. Other integration schemes can be used such as an area integral. This integral gives the average amplitudes of the multipoles at radius r. For comparison it can be helpful to scale a_k by r^{k-1} to get a value that is proportional to the multipole's contribution to the total field at r and has units of tesla.

Figure 3.16(a) shows the scaled amplitudes relative to the hexapole amplitude for a range of multipole terms for circles at a radius of 1 cm and 3 cm for the atom lens with material imperfections. Amplitudes represent the norm of skew and upright amplitudes. Both circles were located at the lens center. The multipole spectrum will in general be different at different planes in the lens. Notice that the greater relative strength of additional poles for the 1 cm circle, particularly the dipole term. Figure 3.16(b) shows a $\log_{10}$ heatmap of the difference between the magnetic field for perfect and imperfect materials.

The two minima occur because of the addition of a dipole term to a hexapole term. This gives for the total field

$$\vec{B} = \left(a_3 r^2 \cos(3\theta) + a_1 \cos(\theta + \phi)\right) \hat{x} + \left(a_3 r^2 \sin(3\theta) + a_1 \sin(\theta + \phi)\right) \hat{y}, \tag{3.10}$$

where ϕ is the angle of the dipole field. For $a_3 r^2 \ll a_1$, valid near the lens center, it can be shown that

$$|\vec{B}|_{\text{near}} \approx a_1 + a_3 r^2 \cos(2\theta - \phi), \tag{3.11}$$

which decreases along the lines $\theta_1 = (\phi - \pi)/2$ and $\theta_1 = (\phi - 3\pi)/2$. Far from the lens origin we have $|\vec{B}|_{\text{far}} = a_3 r^2$. Smoothly joining the near and far solutions results in two minima separated by $\Delta\theta = \pi$ which is exactly what appears in Figure 3.16(b). Thus the central lens region appears roughly as two lensing regions. This can be clearly seen in images of the focus shown in Figure 3.17(b) for an aperture of 1 cm and 6 configurations of random material imperfections. A double focus is apparent in each image.

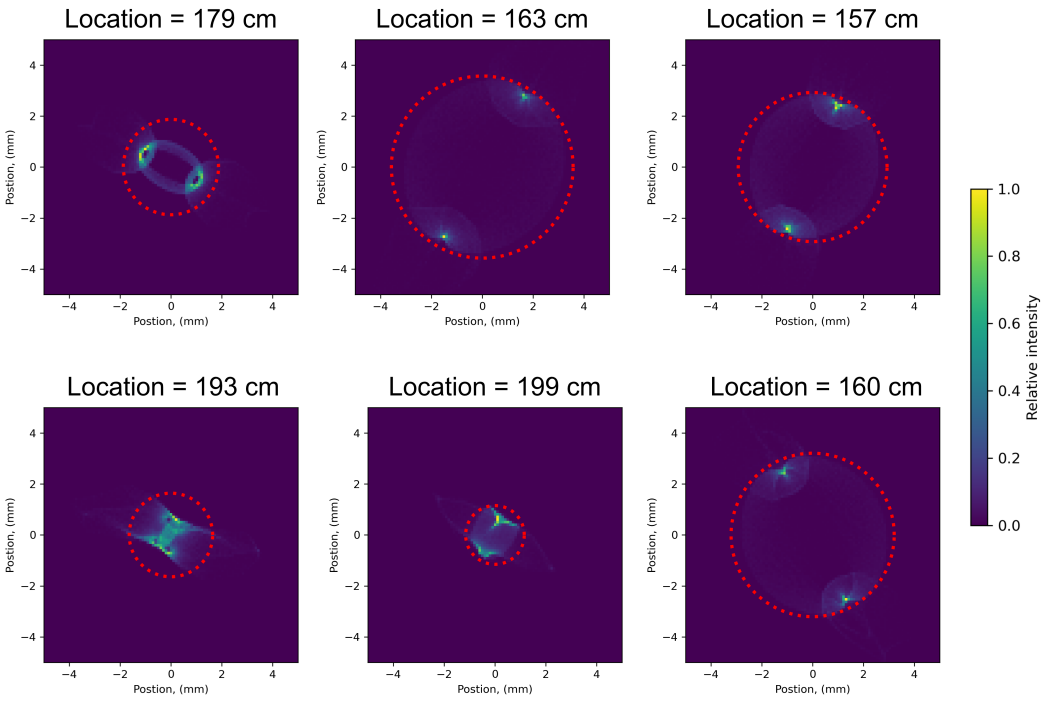

Figure 3.17: Heatmaps of relative intensity at the focus of 6 different configurations of random magnet errors in the atom lens with a 1 cm aperture. The red dotted circle depicts the D_{90}. The relative intensity is scaled in each figure independently.

Chapter 4

Neutral Atom Storage Ring Dynamics

4.1 Introduction

Given the success of particle accelerators, it is unsurprising that accelerator theory has reached an extremely high level of development. However, the term "accelerator physics" is somewhat of a misnomer because large portions of the theory deal with particles not under the influence of longitudinal acceleration. Much of this theory is readily extended to the study of a neutral atom storage ring. In the following chapter the most relevant aspects of accelerator theory [81, 91–93] are reviewed and extended to the case of neutral atoms in magnetic fields.

Before moving onto the accelerator formalism we will introduce bending and beam combination from an atomic physics perspective. We discussed focusing in Section 2.6. We will also derive a system of differential equations that result in linear solutions mapping initial to final phase space coordinates for a particle's trajectory through magnet elements. The mapping can be framed in terms of matrices analogous to the ABCD matrices in optics. Equations can also be derived that conveniently describe the behavior of ensembles of particles. Next we will compare this theory to results from a time stepping simulation and will find good agreement. We end by considering the use of different species and some numerical methods.

4.2 Neutral Atom Beam Manipulation

4.2.1 Bending

Let us begin by reviewing bending of charged particles in magnetic fields. Charged particles interact with the magnetic field via the familiar Lorentz force

$$\vec{F} = q\vec{v} \times \vec{B}, \tag{4.1}$$

where q and $\vec{v}$ are the particle's charge and velocity respectively. Dipole magnetic fields are conventionally used to bend charged particles, though additional higher order terms may be added to produce focusing or effects. The bending radius in a dipole field is found by balancing the centrifugal pseudo force and the magnetic force which gives

$$\begin{aligned} \frac{mv^2}{r_b} &= qvB_z \\ r_b &= \frac{mv}{qB_z}, \end{aligned} \tag{4.2}$$

where the particle's velocity is perpendicular to $\hat{z}$. For reference a singly ionized lithium atom with a speed 210 m/s in a 1 tesla field will have a bending radius of

$$\text{Li}^7 \text{ II bending radius} = \frac{1.165 \times 10^{-26}\,\text{Kg} \cdot 210\,\text{m/s}}{1.602 \times 10^{-19}\,\text{C} \cdot 1\,\text{T}} = 15\,\mu\text{m}. \tag{4.3}$$

Now consider neutral Li^7. Here the bending radius is given by

$$r_b = \frac{mv^2}{\mu_B |\vec{\nabla} B_0|}, \tag{4.4}$$

where the particle's velocity is perpendicular to the gradient $\vec{\nabla} B_0$ and B_0 is the

magnetic field norm $|\vec{B}|$. A reasonable region of magnetic field is perhaps a centimeter which gives an average gradient of ~ 100 T/m. The approximate bending radius for neutral Li^7 is then

$$\text{Li}^7 \text{ I bending radius} = \frac{1.165 \times 10^{-26}\,\text{Kg} \cdot (210\,\text{m/s})^2}{9.274 \times 10^{-24}\,\text{J/T} \cdot 100.0\,\text{T/m}} = 55\,\text{cm}, \tag{4.5}$$

which is more than 4 orders of magnitude larger than that for ionized lithium. Charged particle storage rings are orders of magnitude more compact than neutral particle storage rings.

But how do we produce this bending force? In Chapter 2 we used the multipole expansion to find a field configuration that yielded lensing. We can apply a similar procedure here. Any multipole higher than dipole would work because there is a non zero gradient. Applying the radial gradient operator ∇_r to Equation 2.19 gives

$$\nabla_r B_0 = B_p(n-1)\left(\frac{r^{n-2}}{r_p^{n-1}}\right), \tag{4.6}$$

which suggests the use of large values of n for bending. In principle this is true for a single particle, however we are interested in a beam of particles that may be spread out in the bore of the element. Using some fraction of the bore radius $\delta = r/r_p$ gives

$$\nabla_r B_0 = B_p(n-1)\left(\frac{\delta^{n-2}}{r_p}\right), \tag{4.7}$$

where it can be shown that the maximal value now depends on δ. For δ closer to 1 the optimal value of n is larger, while for δ closer to zero a smaller value is optimal. In a beam of particles we want δ to be small enough for the magnet bore

to accommodate the width of the beam. For a single particle, δ could be very close to 1. Using a value of $\delta = 0.95$ would give a bending radius of about 7.5 cm for ^{7}Li. One can continue down this path and find the optimal multipole for the given conditions.

An obvious alternative is to simply use a revolved hexapole lens, or a hexapole *waveguide*. This has the property of being an approximate linear transformation and we know it possesses reasonably well behaved focusing in practice. We will see that this linear behavior of hexapole lenses and waveguides is essential to accelerator theory. It is also a low order multipole which qualitatively agrees with the principle of smaller values of δ for a particle beam. It is possible that some higher order term, or even combination of terms is better, but in general the behavior of the system then has to be analyzed entirely numerically and it is not clear which choice is superior. For these reasons I chose not to pursue other multipole benders, though they may hold promise.

A magnetic hexapole bender, from this point forward referred to as a bender, can be considered a trap. The trap depth is $V_{\text{trap}} = \mu_B B_p$ in the lab frame. In the particle frame the effects of bending can be considered an asymmetric modulation of the trap depth, decreasing the depth outward radially and increasing the depth inward radially as shown in Figure 4.1. To maximize trapping particles must be mode matched into the bender, otherwise they will oscillate about the minimum. Another way to say this is that there is an ideal orbit for the particle, and if it enters off of that orbit it will oscillate about the orbit. We will discuss this concept extensively in the following sections. The optimal mode matching radius for loading

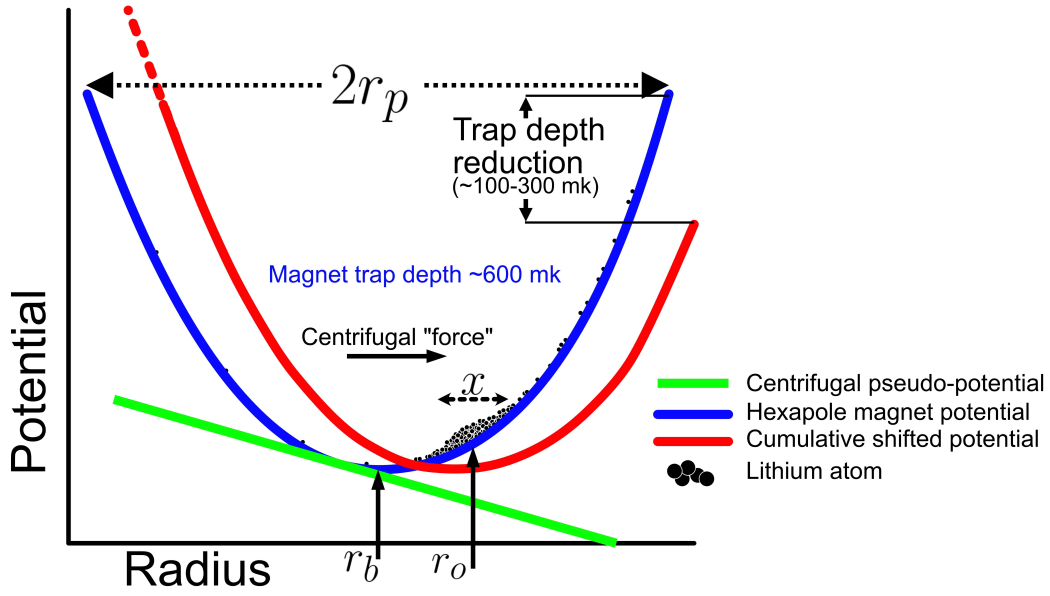

Figure 4.1: The trap model of bending. Particles are displaced from the center of the magnetic potential because of the centrifugal pseudo force. From the particle's perspective the trap height is reduced on one side and increased on the other.

is the radius which results in zero net force in the particle's frame. In the particle's frame the force is given by

$$F = F_c + F_m = \frac{mv^2}{r} - 2\mu_B B_p \frac{r - r_b}{r_p^2}, \tag{4.8}$$

where r_b is the radial center of the lens bore. For $F = 0$ we solve for r to get the mode matching radius, r_o,

$$r_o = \frac{r_b}{2} + \frac{1}{2}\sqrt{r_b^2 - \frac{2mv^2 r_p^2}{\mu_B B_p}}. \tag{4.9}$$

A particle loaded with $r \neq r_o$ will experience restoring forces. To determine these forces we must first account for the conservation of angular momentum by plugging $L = mvr$ into Equation 4.8 to get

$$F = F_c + F_m = \frac{L^2}{mr^3} - 2\mu_B B_p \frac{r - r_b}{r_p^2}. \tag{4.10}$$

Making the approximation that deviations from r_o are small and expressing these deviations as $x = r - r_o$ we can expand Equation 4.10 to get

$$F = F_c + F_m \approx \left(\frac{L^2}{mr_o^3} - 2\mu_B B_p \frac{r_o - r_b}{r_p^2}\right) - \left(\frac{3L^2}{mr_o^4} + \frac{2\mu_B B_p}{r_p^2}\right) x, \tag{4.11}$$

where the first term in parenthesis is zero from Equation 4.8 with $r = r_o$ (Figure 4.1). In the particle's frame it experiences a restoring force with an effective spring constant

$$K = \frac{3L^2}{mr_o^4} + \frac{2\mu_B B_p}{r_p^2}. \tag{4.12}$$

The first term originates from conservation of angular momentum resulting in the particle slowing down for outward radial displacements and the centrifugal pseudo force then weakening, or speeding up for inward displacements and the centrifugal pseudo force then strengthening.

A detail we have overlooked is the slowing a particle experiences as it enters a hexapole lens as discussed in Section 2.6. This will affect the orbit of the particle by changing the initial angular momentum. By applying energy conservation it is easy to show that

$$\Delta v \approx \frac{\mu_B B_p (r - r_b)^2}{m v r_p^2}, \tag{4.13}$$

where Δv is the reduction in velocity. With typical values of $B_p = 1$ Tesla and $r_p = 1$ m and entering the bore halfway to the pole face we get $\Delta_v \approx 1$ m/s. This is a small but not totally negligible value. This effect can be included in Equation 4.9.

So far we have been considering a bender formed by a continuous revolution of a lens. This can of course be done by using short revolved permanent magnets, but the cost of producing such exotic permanent magnets is prohibitive. A more practical approach we have elected to use is to approximate the bender as a series of short lenses.

4.2.2 Injection

The storage ring can be thought of as a toroidal trap. Particles cannot be introduced into a trap that confines in all directions without an irreversible process

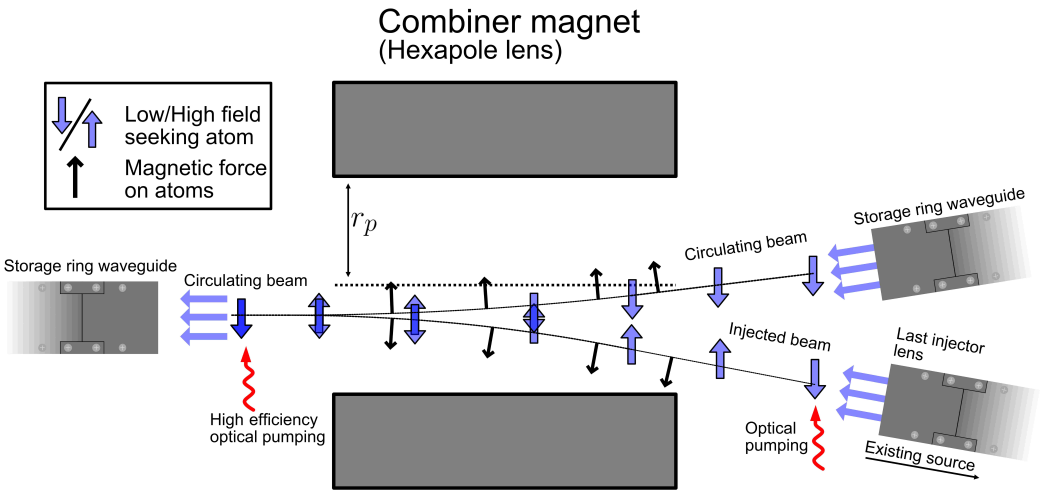

Figure 4.2: Diagram of continuously injecting the atom beam into the storage ring. The last stage of optical pumping must be highly efficient to avoid excessive losses of already injected atoms.

because to have entered the trap they must have had enough energy to climb the potential and will then be able to exit. With a pulsed source this can be done by rapidly turning the trap on. In accelerators this is done with a component known as a kicker magnet which abruptly shifts the particle bunch from one beam line to another. A continuous source requires a continuous irreversible process. Optical pumping is a natural choice for this. We refer to the continuous component as a combiner.

The problem is to guide an injected beam on top of an existing circulating beam. This is quite similar to the famous Stern-Gerlach experiment. One way to achieve this is to simply use a Stern-Gerlach magnet which applies a nearly constant force. The injected beam is optically pumped into a high field seeking

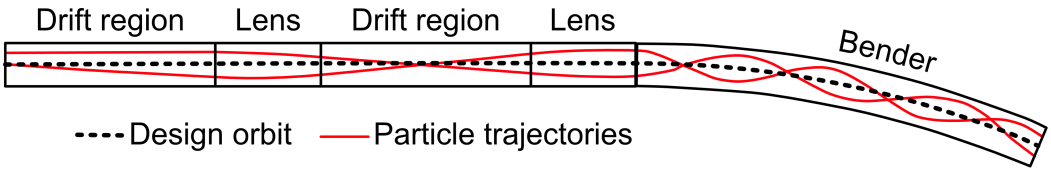

Figure 4.3: An example of an accelerator lattice. In this example the bending element is a curved lens. The drift regions are regions of the lattice approximately free of electromagnetic fields.

state (untrapped) and will bend in the opposite direction of the circulating low field seeking state (trapped) beam. The two beams combine at the output and both are pumped into the low field seeking state. The efficiency of this pumping will introduce a loss mechanism.

An alternative approach is to combine the beams using a hexapole lens as shown in Figure 4.2. The principle is identical to the Stern-Gerlach magnet, and has the added benefit that the fields are confining for the circulating atoms though they are diverging for the injected atoms. If particles that do enter the ring circulate for longer periods of time then this can be an advantageous trade off.

4.3 Particle Accelerators

An accelerator is composed of a series of components that guide and deflect a beam of charged particles. In analogy with light optics, the system of elements and phenomena of deflection is also called **beam optics**. Components include bending sections, focusing sections, accelerating sections, insertion devices, field free regions and RF beam cooling components. Field free regions are conventionally referred

to as a **drift** regions. The components are typically far enough apart such that the interactions between components are neglected to first order. Magnetic fields are the primary method of deflecting particles, particularly dipole magnetic fields for bending and quadrupole magnetic fields for focusing. Higher order fields such as hexapole or octupole can be used for correcting aberrations and adjusting other perturbations.

The series of magnetic deflecting components, or elements, is referred to as a lattice. The lattice defines the **design orbit**, shown in Figure 4.3, which is the orbit that a particle with the ideal initial phase space coordinates would travel along. This ideal particle travels along the design orbit at the **design speed** or **design momentum**. Particles that deviate from the ideal initial conditions will still travel in the vicinity of this orbit, though in general will undergo oscillations about it known as **betatron oscillation**. In a purely linear accelerator this orbit will be a straight line, whereas in a circular accelerator some lattice elements will bend the orbit and it will return to its starting point after traveling a total bending angle of 2π radians. In a lattice with bending, a particle with zero initial transverse displacement and momentum but with the incorrect longitudinal momentum, known as an **off momentum** particle, will deviate from the design orbit because of centrifugal pseudo forces. This is referred to as **chromatic** behavior, another analogy to light optics.

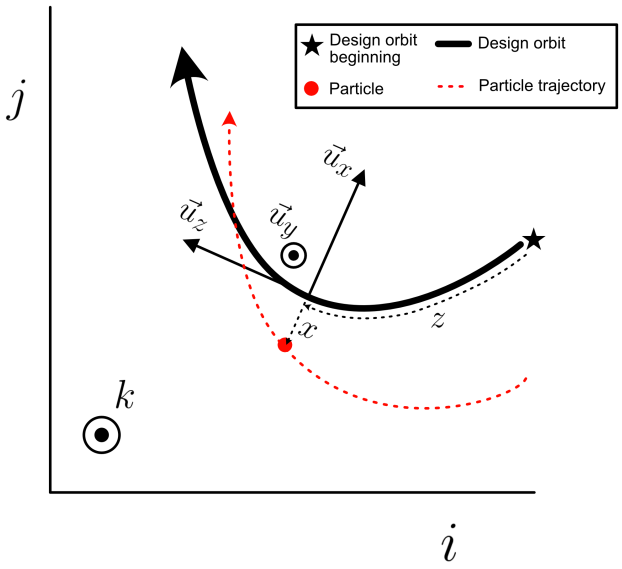

Figure 4.4: The Frenet-Serret coordinate system. The design orbit is the orbit of the particle with ideal initial conditions. Particles that deviate from these conditions travel in a coordinate system that moves along and is relative to the design orbit. This simplifies the equations of motion and analysis, but introduces pseudo forces that must be accounted for. The lab frame is represented by the ijk coordinate system.

4.4 Frenet-Serret Coordinate System

A purely linear lattice is adequately described by a Cartesian coordinate system and a purely circular lattice with constant bending radius is well described by cylindrical coordinates. However, a real storage ring is composed of discrete linear and bending sections which may have different bending radii. This motivates the adoption of a more general curvilinear coordinate system, the Frenet-Serret coordinate system depicted in Figure 4.4. In this coordinate system a particle's phase space coordinates are defined relative to the design orbit. A Cartesian coordinate system travels along the design orbit, with the unit vector $\hat{u}_z$ tangent to the orbit, $\hat{u}_x$ horizontally perpendicular to the orbit, and $\hat{u}_y$ vertically perpendicular to the orbit.

For later use we will derive the relationship between an infinitesimal change in the global coordinate system and the Frenet-Serret coordinate system. First, the particle's location in a global coordinate system can be described by

$$\vec{r} = \vec{r}_o(z) + x(z)\vec{u}_x + y(z)\vec{u}_y, \tag{4.14}$$

where $\vec{r}_o(z)$ is the coordinates of the orbit at distance z along the orbit, and $\vec{u}_x$ and $\vec{u}_y$ are the orbit unit vectors in the global coordinate system. Assuming no bending in the vertical direction we have

$$\vec{u}_z(z) = \frac{d\vec{r}_o(z)}{dz}, \quad \frac{d\vec{u}_x(z)}{dz} = k_0(z)\vec{u}_z(z), \text{ and } \frac{d\vec{u}_y(z)}{dz} = 0, \tag{4.15}$$

where the bending radius r_0 is expressed as $k_0 = 1/r_0$. In the following sections the

reason for this notation will become clear. With the above relations it can be shown that

$$d\vec{r} = dx\vec{u}_x + dy\vec{u}_y + hdz\vec{u}_z, \tag{4.16}$$

where

$$h = 1 + k_0 x. \tag{4.17}$$

4.5 Equations of Motion

The equations of motion in the Frenet-Serret coordinate system can be derived with Lagrangian mechanics. In the following derivation gravity is neglected. This is justified because the force of gravity is orders of mangitude weaker than the magnetic force a particle typically experiences. The kinetic energy in Frenet-Serret coordinates is given by

$$T = \frac{m}{2}\dot{\vec{r}} \cdot \dot{\vec{r}} = \frac{m}{2}\left(\dot{x}^2 + \dot{y}^2 + h^2\dot{z}^2\right), \tag{4.18}$$

where we used Equation 4.16.

Combining this with the potential for a paramagnetic atom in a magnetic field the Lagrangian is then

$$L = T - V = \frac{m}{2}\left(\dot{x}^2 + \dot{y}^2 + h^2\dot{z}^2\right) - \mu B_0(x, y, z), \tag{4.19}$$

where μ is positive for low field seeking and negative for high field seeking and is constant. Our first approximation to arrive at linear solutions is to replace the

continuous dependence of $B_0(x, y, z)$ on z with a rectangle function as

$$B_0(x, y, z) \approx \sum_i B_{0i}(x, y)\Pi(z, z_i, L_i), \tag{4.20}$$

where Π is a rectangle function with value 1 between z_i to $z_i + L_i$, the index i refers to the ith magnetic element, z_i the beginning of the ith magnetic element, and L_i the length of the ith element. This is known as the hard edge approximation which assumes that magnetic fields abruptly turn on and off when entering or exiting a magnetic element. Additionally, we assume that the degree of bending k_0 follows a similar rectangle function relation. In Section 4.14 we will discuss methods to improve this approximation. For convenience the term $B_0(x, y)$ will be assumed to refer to Equation 4.20 with the appropriate value of i.

The new Lagrangian is then

$$L = \frac{m}{2}(\dot{x}^2 + \dot{y}^2 + h^2\dot{z}^2) - \mu B_0(x, y). \tag{4.21}$$

Recall that Lagrange's equation without constraints is

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{q}_i} - \frac{\partial L}{\partial q_i} = 0. \tag{4.22}$$

We can immediately note the presence of a conserved quantity because

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{z}} = 0, \tag{4.23}$$

with value

$$\frac{\partial L}{\partial \dot{z}} = mh^2\dot{z} = m(1+k_0x)^2\dot{z} = m(1+k_0x_i)^2\dot{z}_i = l, \tag{4.24}$$

where l is constant and x_i and z_i are the initial values entering an element of constant k_0.

The conserved quantity is simply angular momentum expressed in the Frenet-Serret coordinate system. Keep in mind that this is for a specific element, and l will in general have different values in different elements. The full equations of motion are

$$\begin{aligned} &m(1+k_0x)^2\dot{z} = l, \\ &m\ddot{x} - k_0m(1+k_0x)\dot{z}^2 + \mu\frac{\partial}{\partial x}B_0(x,y) = 0, \\ &\text{and} \\ &m\ddot{y} + \mu\frac{\partial}{\partial y}B_0(x,y) = 0. \end{aligned} \tag{4.25}$$

Using the conserved quantity l, the equations of motion can be reduced to

$$\begin{aligned} &m\ddot{x} - \frac{k_0l^2}{m(1+k_0x)^3} + \mu\frac{\partial}{\partial x}B_0(x,y) = 0 \\ &\text{and} \\ &m\ddot{y} + \mu\frac{\partial}{\partial y}B_0(x,y) = 0. \end{aligned} \tag{4.26}$$

It is advantageous to change the parameterization from time to distance along the design orbit z. In this parameterization time derivatives are replaced by derivatives

with respect to z represented by primes. The new derivative operators are given by

$$\dot{q}_i = \frac{dq_i}{dt} = v\frac{dq_i}{ds} = v\frac{dz}{ds}\frac{dq_i}{dz} = \frac{v}{s'}\frac{dq_i}{dz} = \frac{v}{s'}q_i'$$

and (4.27)

$$\ddot{q}_i = \frac{d}{dt}\frac{dq_i}{dt} = \frac{v^2}{s'^2}(q_i'' - s''\frac{q_i'}{s'}),$$

where

$$s' = |\frac{d\vec{r}}{dz}| = \sqrt{x'^2 + y'^2 + h^2}, \tag{4.28}$$

v is the speed of a particle along its own trajectory, s is the distance along that trajectory, and Equation 4.16 was used. Our new equations of motion are

$$m\frac{v^2}{s'^2}(x'' - s''\frac{x'}{s'}) - \frac{k_0 l^2}{m(1+k_0 x)^3} + \mu\frac{\partial}{\partial x}B_0(x,y) = 0$$

and (4.29)

$$m\frac{v^2}{s'^2}(y'' - s''\frac{y'}{s'}) + \mu\frac{\partial}{\partial y}B_0(x,y) = 0.$$

To arrive at linear solutions, several approximations must be made. The first approximation is that transverse particle displacements are small compared to bending radii which gives $k_0 x \ll 1$ and $h \approx 1$. The next approximation is the familiar paraxial approximation from optics, which assumes that rays and the optical axis form small angles. Here this corresponds to $x' = \dot{x}/\dot{z} \ll 1$ and $y' = \dot{y}/\dot{z} \ll 1$, and thus $s' \approx 1$ and $s'' \approx 0$. To motivate this approximation, consider that the velocity for a lithium atom to just overcome a trap depth of 1 tesla is about 40 m/s which corresponds to an angle of about 0.19 radians. This is approximately the maximum transverse velocity one could expect from a trapped lithium atom, while the typical

value is half or less. Finally, combining the above assumptions we can use $l \approx mv$.

Applying these approximations yield

$$\begin{aligned} &x'' - k_0(1 - k_0 x) + \frac{(1 + 2k_0 x)}{mv^2}\mu_0 \frac{\partial}{\partial x} B_0(x, y) = 0 \\ &\text{and} \\ &y'' + \frac{(1 + 2k_0 x)}{mv^2}\mu_0 \frac{\partial}{\partial y} B_0(x, y) = 0. \end{aligned} \tag{4.30}$$

In our new z parameterization, coupling between x and y has become explicit. Notice that this is only present in regions with bending, i.e. $k_0 = (1/r_0) \neq 0$. This coupling occurs because as a particle travels through a bend, horizontal deviations result in changes to path length and thus transit time. In the time parameterization this would appear when integrating the equations of motion.

We can rewrite the bending force from Equation 4.10 to get an explicit relation for the magnetic force. Doing so yields

$$\begin{aligned} &\mu_0 \frac{\partial}{\partial x} B_0(x, y) = mv_o^2 k_0 + k_m x \\ &\text{and} \\ &\mu_0 \frac{\partial}{\partial y} B_0(x, y) = k_m y, \end{aligned} \tag{4.31}$$

where k_m is the magnetic force spring constant $2\mu B_p / r_p^2$, v_o is the design speed, and k_m is positive(negative) for low(high) field. We have used the design speed in Equation 4.11 because that yields the correct bending radius of r_o. Substituting

this into Equation 4.30 and expanding to first order in x gives

$$\begin{aligned} & x'' + \left(k_0^2 + \frac{2k_0^2 v_o^2}{v^2} + \frac{k_m}{mv^2}\right) x + k_0 \frac{v_o^2}{v^2} - k_0 = 0 \\ & \text{and} \\ & y'' + \frac{k_m y}{mv^2} = 0. \end{aligned} \tag{4.32}$$

Recall that v is the speed of a particle along its own trajectory, and v_o is the speed of the ideal particle traveling along the design orbit. Let us replace v with the relation

$$v = v_o(1+\delta) \rightarrow \delta = \frac{v - v_o}{v_o}. \tag{4.33}$$

The particle's deviation from the design speed is now expressed with the variable δ. A particle with $\delta \neq 0$ is an off momentum particle. If we assume $\delta \ll 1$ then we can expand Equation 4.32 in δ. This is reasonable considering that even for a 100 mK beam of lithium atoms the FWHM spread in velocity is ~ 26 m/s, yielding $\delta \approx 0.1$. Inserting Equation 4.33 into Equation 4.30 and expanding to first order in $\delta \cdot x$ yields

$$\begin{aligned} & x'' + \left(\frac{k_m}{mv_o^2} + 3k_0^2\right) x - 2\delta k_0 = 0 \\ & \text{and} \\ & y'' + \frac{k_m y}{mv_o^2} = 0. \end{aligned} \tag{4.34}$$

We have finally arrived at the desired result: a system of equations of motion that yield linear solutions. Along the way we made several approximations, though it will be shown in Section 4.12 that our equations of motion work reasonably well. Once again, note that for an element with no nominal bending, such as a linear

waveguide, k_0 is zero and the equations of motion for x and y are identical.

Let us inspect each term in our equations of motion. The term $\pm k/mv_o^2$ is the effective spring constant of the magnetic element. For larger values of mass the effective spring constant is reduced because of the larger inertia and therefore resistance to deflection. Larger values of v_o reduce the effective spring constant for more subtle reasons. The term is squared because of two phenomena. First, faster particles spend less time in an element and therefore experience proportionally less bending. Second, recall that the z parameterization resulted in slopes (x') and rates of change of slopes (x'') rather than velocity ($\dot{x}$) or acceleration ($\ddot{x}$). A faster particle experiencing the same total transverse momentum kick would result in proportionally smaller changes in x' and x''. Thus the combined effect is the presence of a v_o^2 term. The term $3k_0^2$ originates from the conserved quantity l which is just conservation of angular momentum introducing a restoring pseudo force. Finally, the term $2\delta k_0$ results from particles with non design speeds experiencing a magnetic force which does not exactly balance their centrifugal pseudo-force. The centrifugal pseudo-force depends on speed to the second power, and so the expansion in δ results in a factor of 2.

4.6 Solution and Matrix Formalism

Before solving Equations 4.34 we will rephrase them in a more general form

$$u''(z) + Ku(z) - P = 0, \tag{4.35}$$

where $P = 2\delta k_0$ and $K = \frac{k_m}{mv_o^2} + 3k_0^2$. Because of the relative strength of the magnetic forces, we will assume that the sign of K is dictated by the sign of k_m.

This is an inhomogenous second order ordinary differential equation to which the solution is well known. The total solution is the sum of two solutions,

$$u(z) = u_c(z) + U_p(z), \tag{4.36}$$

where u_c is the *complementary* solution when $P = 0$ and U_p is the *particular* solution for the system when $P \neq 0$. We will solve this inside a single element of constant P and K with z replaced by $\Delta z = z - z_i$ where z_i is the beginning of the element. The solutions are

$$\begin{aligned}
u_c(\Delta z) &= u_{ci}\cosh\Omega + \frac{u'_{ci}}{\sqrt{-K}}\sinh\Omega, \\
u'_c(\Delta z) &= u_{ci}\sqrt{-K}\sinh\Omega + u'_{ci}\cosh\Omega, \\
U_p(\Delta z) &= U_{pi}\cosh\Omega + \frac{U'_{pi}}{\sqrt{-K}}\sinh\Omega + \frac{P}{K}(1-\cosh\Omega), \\
&\text{and} \\
U'_p(\Delta z) &= U_{pi}\sqrt{-K}\sinh\Omega + U'_{pi}\cosh\Omega - \frac{P\sqrt{-K}}{K}\sinh\Omega,
\end{aligned} \tag{4.37}$$

where the i subscript indicates the value of u or U at $\Delta z = 0$, and $\Omega = \Delta z\sqrt{-K}$. For low field seeking, $K > 0$ and the hyperbolic sines and cosines become ordinary sines and cosines which yields oscillatory behavior. For high field seeking, $K < 0$ which results in $U_p(\Delta z) \to \infty$ and $U'_p(\Delta z) \to \infty$ for $\Delta z \to \infty$.

The two solutions have distinct physical significance. For particles with design speed, $\delta = 0$, $u_c(z)$ is the only solution and describes the behavior of the

particles relative to the design orbit. For off momentum particles ($\delta \neq 0$) there can be an additional offset from the design orbit of value $U_p(z)$. This is interpreted as a shifted orbit about which the particle oscillates rather than the design orbit. This is a first order chromatic correction.

Equations 4.37 are most helpfully represented as three 3x3 matrices for each of the conditions $K = 0$, $K < 0$, and $K > 0$. Additionally, we will replace P with the value $2k_0$. The full matrices for an element are

$$
\begin{aligned}
K = 0 : \mathbf{M}(\Delta z) &= \begin{bmatrix} 1 & \Delta z & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \\
K > 0 : \mathbf{M}(\Delta z) &= \begin{bmatrix} \cos\Omega & \sin\Omega/\sqrt{K} & 2k_0\left(1-\cos\Omega\right)/K \\ -\sqrt{K}\sin\Omega & \cos\Omega & 2k_0\sin\Omega/\sqrt{K} \\ 0 & 0 & 1 \end{bmatrix},
\end{aligned}
\tag{4.38}
$$

and

$$
K < 0 : \mathbf{M}(\Delta z) = \begin{bmatrix} \cosh\Omega & \sinh\Omega/\sqrt{|K|} & 2k_0\left(1-\cosh\Omega\right)/K \\ \sqrt{|K|}\sinh\Omega & \cosh\Omega & 2k_0\sinh\Omega/\sqrt{|K|} \\ 0 & 0 & 1 \end{bmatrix},
$$

where $\Omega = \sqrt{|K|}\Delta z$. $K = 0$ represents a field free region, $K > 0$ a focusing region, and $K < 0$ a de-focusing region. The total matrix for an element of length L is $\mathbf{M}(L)$. The matrices can also be used to map the trajectory through an element by replacing Δz with the depth into the element. We have only considered elements in which the bending radius is constant, but this is not the case for the combiner. We will address this numerically in Section 4.14.

The above matrices are commonly known as transfer matrices, and map initial conditions to final conditions through an element such as a lens, drift region, or bender. They are analogous to the well-known ABCD matrices in ray optics. In the x dimension chromatic effects can be accounted for by the use of a length three vector by

$$\mathbf{X} = \begin{bmatrix} x \\ x' \\ \delta \end{bmatrix} \qquad \mathbf{X}_f = \mathbf{M}\mathbf{X}_i, \tag{4.39}$$

where $\mathbf{X}_i$ is the initial value of $\mathbf{X}$ and $\mathbf{X}_f$ is final value after traveling through the element(s). In the y dimension the matrices can be reduced to their upper left 2x2 sub-matrices and a length two vector is used. One can also use a length five vector, $[x, x', y, y', \delta]$ with the appropriate 5x5 matrix.

The transfer matrices have the interesting and useful property that their determinant is always 1. In the case of $K = 0$ this is readily apparent. For $K \neq 0$ note that the matrices have the form

$$\begin{bmatrix} C & S & E \\ C' & S' & F \\ 0 & 0 & 1 \end{bmatrix}, \tag{4.40}$$

where C and S are ordinary or hyperbolic sine and cosine functions, and C' and S' are the respective z derivatives. The determinant is

$$\det \mathbf{M} = CS' - SC' = 1 \tag{4.41}$$

by basic ordinary and hyperbolic trig identities. We will make use of this property later.

Because of the linearity of the matrices, a series of elements can be represented by a single transfer matrix by multiplying the sequence of transfer matrices together in order, such as

$$\mathbf{M}_{\text{total}} = \prod_{i}^{N} \mathbf{M}_i, \tag{4.42}$$

where N is the number of elements, and M_i is the ith transfer matrix in the x or y dimension (or both if using the full 5x5 transfer matrix). For a periodic lattice, i can be extended to values greater than the total number of elements for particles making multiple revolutions.

We may also be interested in the transfer matrix some total depth z into the lattice. This is found by

$$\mathbf{M}(z) = \begin{cases} \mathbf{M}_0(z), & \text{if } z \leq L_0 \\ \mathbf{M}_{j+1}(z - z_{j+1}) \prod_i^j \mathbf{M}_i, & \text{otherwise.} \end{cases} \tag{4.43}$$

4.6.1 Periodic Stability Criteria

Particles will ideally experience a large number of revolutions in a circular lattice. For example, at the LHC a beam of protons will circulate several hundred million times. One factor in achieving this is the *stability* of the lattice, which loosely speaking is the degree to which particles do not wander away from the design orbit. There are many considerations to take into account when designing a stable periodic lattice, and the simplest is the stability of the transfer matrices.

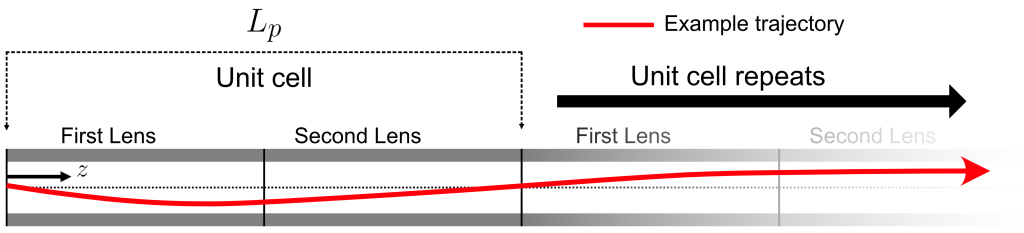

Figure 4.5: A simple periodic linear lattice. The unit cell, the smallest unit of symmetry, is two lenses. A particle that reaches the end of the second lens is effectively mapped back to the input of the first lens.

In a storage ring the lattice is periodic, that is

$$\mathbf{M}(z) = \mathbf{M}(z + L_p), \tag{4.44}$$

where L_p is the path length of the design orbit for one revolution. For N circulations we have the transfer matrix

$$\mathbf{M}_{\mathrm{N}} = \mathbf{M}_{\mathrm{period}}^{N}. \tag{4.45}$$

It is not guaranteed that as $N \to \infty$ the components of M_∞ do not grow unbounded. Consider the trivial example of a drift region of length L_d with transfer matrix

$$\mathbf{M}_{\mathrm{drift}} = \begin{bmatrix} 1 & L_d \\ 0 & 1 \end{bmatrix}, \tag{4.46}$$

which raised to the power N can be shown to give the transfer matrix

$$\mathbf{M}_{\mathrm{drift}}^{N} = \begin{bmatrix} 1 & NL_d \\ 0 & 1 \end{bmatrix}, \tag{4.47}$$

which grows unbounded as $N \to \infty$.

Let us consider a simple periodic lattice of two ideal lenses, shown in Figure 4.5. One can think of a periodic lattice as a crystal, with the smallest repeating unit known as a *unit cell*. Both lenses have a length of 10 cm and bore radius of 1 cm. The first lens has $B_p = 0.9$ tesla and the second has $B_p = 0.1$ tesla. The total transfer matrix in x or y, neglecting chromaticity and assuming a design speed of 210 m/s, is

$$\mathbf{M}_{\text{total}} = \mathbf{M}_{\text{lens 2}}\mathbf{M}_{\text{lens 1}} = \begin{bmatrix} -0.373 & 0.023 \\ -13.7 & -1.84 \end{bmatrix},$$

where

$$\mathbf{M}_{\text{lens 1}} = \begin{bmatrix} 0.825 & 0.094 \\ -3.397 & 0.825 \end{bmatrix} \tag{4.48}$$

and

$$\mathbf{M}_{\text{lens 2}} = \begin{bmatrix} -0.230 & 0.054 \\ -17.5 & -0.230 \end{bmatrix}.$$

Raising $\mathbf{M}_{\text{total}}$ to the 100th power gives

$$\mathbf{M}_{\text{total}}^{100} = \begin{bmatrix} -2.03E+19 & -1.79E+18 \\ 1.07E+21 & 9.41E+19 \end{bmatrix}, \tag{4.49}$$

which would map even nanometer or nanoradian initial values to unrealistically large final values. The results of a time-stepping simulation of the trajectory of particle in this lattice are shown in Figure 4.6. Even in a system with no drift regions and only elements with a restoring force, we still encounter unconditionally unstable behavior.

This behavior can be predicted using the eigenvalues of the transfer matrix.

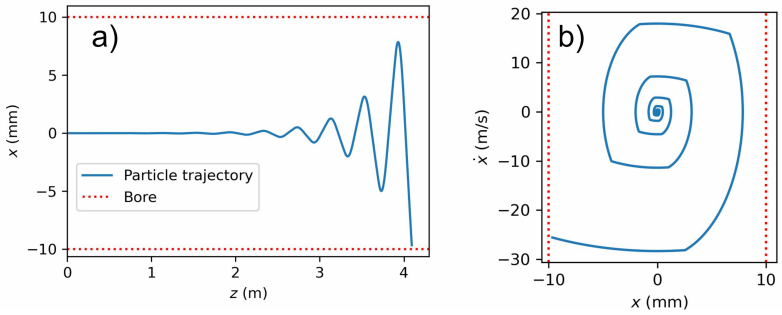

Figure 4.6: Results of numerical time-stepping through the lattice shown in Figure 4.10. (a) The particle's trajectory along the length of the lattice. The particle collides with the lens bore after traveling through 20 unit cells. (b) The particle's trajectory through phase space. The particle's initial position in phase space is amplified through every lens.

The eigenvalues, λ, of a matrix are a measure of the magnitude of a matrix, and so must remain bounded if raised to the power N for $N \to \infty$. The eigenvalues of a 2x2 transfer matrix are found by solving

$$\det\left(\mathbf{M} - \lambda \mathbf{I}\right) = \begin{bmatrix} A - \lambda & B \\ C & D - \lambda \end{bmatrix} = \lambda^2 - \lambda\left(A + D\right) + 1 = 0, \tag{4.50}$$

where we made use of Equation 4.41. This yields the two eigenvalues

$$\lambda_{\pm} = \frac{1}{2}(A + D) \pm \sqrt{\left(\frac{1}{2}(A + D)\right)^2 - 1} \tag{4.51}$$

and the diagonalized matrix

$$\begin{bmatrix} \lambda_{\pm} & 0 \\ 0 & \lambda_{\pm} \end{bmatrix}. \tag{4.52}$$

This matrix must have eigenvalues $|\lambda_{\pm}| <= 1$ otherwise it will grow unbounded as

$N \to \infty$. Solving $|\lambda_{\pm}| <= 1$ yields the stability condition

$$|A + D| = |\mathrm{Tr}\,(\mathbf{M})\,| \leq 2. \tag{4.53}$$

The trace of the total transfer matrix in Equation 4.48 is $\sim$ -2.2, which is an unstable transfer matrix. Changing the value of the pole face magnetic field of the second lens from 0.9 tesla to 0.5 tesla yields $A + D \approx -1.1$, a stable solution.

The eigenvalue equation of a 3x3 transfer matrix is

$$\det\left(\mathbf{M}_{3\mathrm{x}3} - \lambda\mathbf{I}\right) = \begin{bmatrix} A-\lambda & B & E \\ C & D-\lambda & F \\ 0 & 0 & 1-\lambda \end{bmatrix} = \lambda^2 - \lambda\,(A + D) + 1, \tag{4.54}$$

which yields the same eigenvalue problem. Thus the chromatic behavior of the lattice does not affect the stability.

It is important to keep in mind that the stability condition we derived is in the context of the approximations and assumptions that led us to linear equations of motions. Substantial deviations from these approximations and assumptions could yield unstable lattice configurations, even if Equation 4.53 is satisfied. Nonlinear forces can introduce instability into specific regions of phase. There are also other stability conditions such as the avoidance of resonances, discussed in Section 4.8, that are not accounted for in this analysis and impose additional stability criteria.

4.6.2 Stability Regions

In the previous section we observed that a periodic lattice of two lenses could be tuned into or out of a region of instability by changing the strength of the magnetic fields. This implies that one cannot simply choose what seem like reasonable values for key parameters and achieve a viable lattice. In fact, even for simple systems the dependence of the stability on element parameters can be quite complex. For example, consider Figure 4.7 which plots stability against the field strength in the lenses of the lattice in Figure 4.5. More complicated systems with more variable parameters will in general have more complicated higher dimensional regions of stability.

The size of stability regions is negatively affected by the addition of non-confining elements. We can study this by adding a drift region between each lens in our two lens lattice. Figure 4.8 shows the reduction in the size of stability regions as the length of the drift element increases.

4.6.3 Dispersion

In Section 4.6 we derived a set of matrices (Equation 4.38) that map a particle's initial coordinates to final coordinates while including chromatic effects. An alternative formulation that is particularly helpful for periodic lattices is possible. To arrive at this, we introduce a different form of Equation 4.36,

$$u(z) = u_\beta(z) + D(z)\delta, \tag{4.55}$$

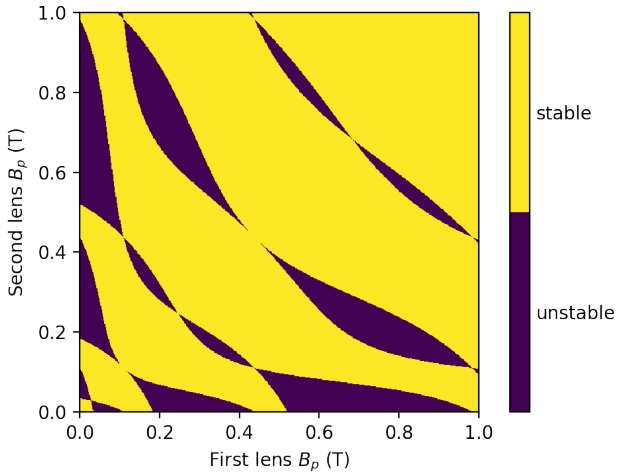

Figure 4.7: Stability plot of the two lens system shown in Figure 4.5 as a function of field strength of each lens.

where $D(z)$ is known as dispersion, and $u_\beta(z)$ is the trajectory when $\delta = 0$. $D(z)\delta$ can be thought of as a new orbit offset from the design orbit that the particle oscillates about with displacement $u_\beta(z)$. The equation of motion for $D(z)$ is Equation 4.35 with $\delta = 1$,

$$D''(z) + KD(z) - 2k_0 = 0. \tag{4.56}$$

This has the same 3x3 matrix solutions we already found. The vector we now propagate is

$$\mathbf{D}(z) = \begin{bmatrix} D(z) \\ D'(z) \\ 1 \end{bmatrix}, \tag{4.57}$$

according to

$$\mathbf{D}_f = \mathbf{M}\mathbf{D}_i. \tag{4.58}$$

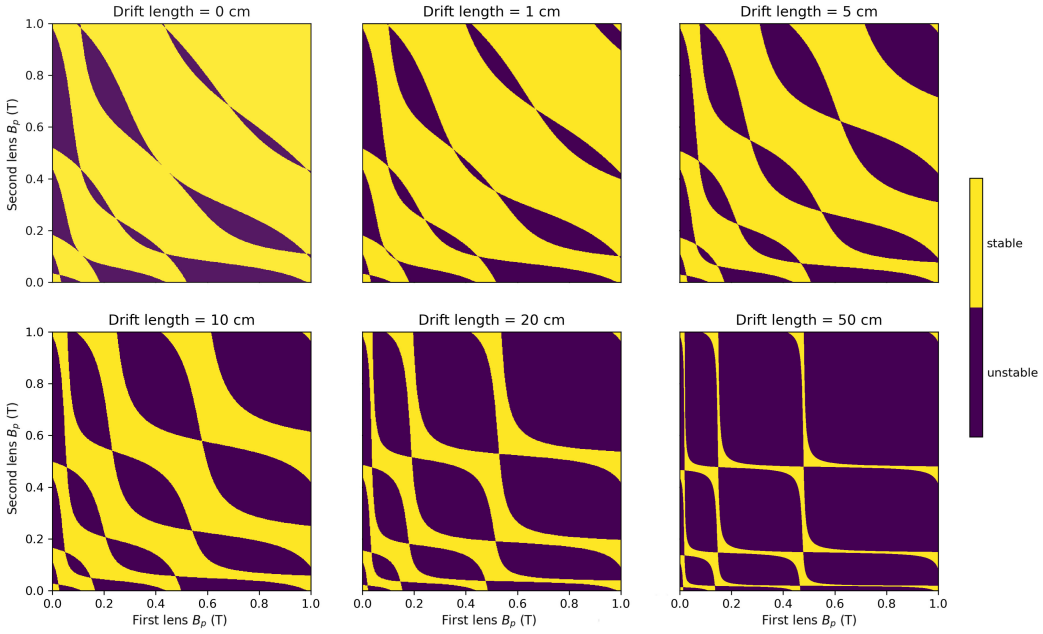

Figure 4.8: Stability plots of the two lens system shown in Figure 4.5 with drift regions between each lens. As the length of the drift elements increases, the size of the stability regions as a function of lens strength decreases.

An important detail is that dispersion is only generated when the third column of the first and second row of the 3x3 transfer matrix is nonzero, which only occurs for elements with bending.

To summarize, we set $\delta = 1$ so that the complementary solution to Equation 4.56 can be simply scaled by δ to find the resulting trajectory offset. It is a feature of the complementary solution that it is proportional to the source term, which is $2k_0\delta$. We are only interested in the complementary solution to Equation 4.56 because it is still just Equation 4.35, to which we already found the particular solution, labeled $u_\beta(z)$ in this case.

To find $D(z)$ in a linear lattice we propagate the vector 4.57 through the lattice with initial value $[0, 0, 1]$. For periodic lattices we have $D(z) = D(z + L_p)$, giving

$$\mathbf{D}(z + L_p) = \mathbf{D}(z) = \mathbf{M}_{\text{period}}\mathbf{D}(z). \tag{4.59}$$

With some algebra this yields

$$\begin{aligned} D(z)' &= \frac{m_{21}m_{13} + m_{23}(1 - m_{11})}{2 - m_{11} - m_{22}} \\ &\text{and} \\ D(z) &= \frac{m_{12}D' + m_{13}}{1 - m_{11}}, \end{aligned} \tag{4.60}$$

where m_{ij} is the periodic transfer matrix element in the ith column and jth row at location z.

4.7 The Hill Equation

So far we have developed a formalism that allows us to propagate a particle's coordinates through a series of elements using a matrix formalism. With this matrix formalism we can also predict whether a periodic lattice will be stable and able to support many particle revolutions. Now we will derive a method suited to determining the behavior of an ensemble of particles in a periodic lattice. To arrive at this method, we start by solving the Hill differential equation

$$u'' + K(z)u = 0, \tag{4.61}$$

where $K(z)$ is a piecewise periodic function and u is x or y. This differential equation appears in many problems in physics. We can solve this by substituting the ansatz

$$u(z) = Aw(z)\cos\left(\Psi(z) + \phi\right) \tag{4.62}$$

into Equation 4.61, which yields

$$A\left[w'' - w\Psi'^2 - K(z)w\right]\cos\left(\Psi + \phi\right) - \left[2w'\Psi' + w\Psi''\right]\sin\left(\Psi + \phi\right) = 0, \tag{4.63}$$

where $\Psi = \Psi(z)$ and $w = w(z)$. For a solution where $A \neq 0$ and $\Psi + \phi$ is not a constant, the two terms must equal zero independently everywhere in the lattice. This requires

$$w'' - w\Psi'^2 - K(z)w = 0 \tag{4.64}$$

and

$$2w'\Psi' + w\Psi'' = 0, \tag{4.65}$$

with Equation 4.65 yielding for Ψ

$$\Psi(z) = \int_0^z \frac{ds}{w^2(s)}. \tag{4.66}$$

Inserting this into Equation 4.64 gives

$$w'' - \frac{1}{w^3} - K(z)w = 0. \tag{4.67}$$

Unfortunately this equation has no known analytic solutions, but in the

following sections we will work out methods to circumvent this issue. To facilitate this, we will replace the term $w(z)$ with $\sqrt{\beta}$. $\beta(z)$ is known as the beta function. For similiar reasons, we will also replace A with $\sqrt{\epsilon}$, where ϵ is known as the emittance. Both β and ϵ have units of distance. The final solution and its derivative are

$$u(z) = \sqrt{\epsilon\beta(z)}\cos\left(\Psi(z) + \phi\right) \tag{4.68}$$

and

$$u'(z) = -\sqrt{\frac{\epsilon}{\beta(z)}}\left[\alpha(z)\cos\left(\Psi(z) + \phi\right) + \sin\left(\Psi(z) + \phi\right)\right], \tag{4.69}$$

where $\alpha(z) = \frac{-\beta'(z)}{2}$. It can be shown with Floquet's Theorom that

$$\beta(z) = \beta(z + L_p). \tag{4.70}$$

Notice that there are 4 variables in Equation 4.68: ϵ, $\beta(z)$, $\Psi(z)$, and ϕ. The values $\beta(z)$ and $\Psi(z)$ are given by solutions to Equations 4.64 and 4.65. Thus, these terms are independent of the initial particle conditions u_i and u_i'. The terms ϵ and ϕ must then depend on the initial particle conditions. Notice that these value have no z dependence. Notice as well that because of the squared term in Equation 4.66, Ψ must be a monotonically increasing function. A particle with a given ϵ and ϕ will thus travel through the lattice with its phase advancing according to $\Psi(z)$ performing oscillations with maximum amplitude $\sqrt{\epsilon\beta(z)}$. The envelope that the particle traces out is then

$$\text{Amplitude envelope} = \sqrt{\epsilon\beta(z)}. \tag{4.71}$$

For an ensemble of particles where the largest value of emittance is $\epsilon_{\max}$, we have

$$\text{Ensemble envelope} = \sqrt{\epsilon_{\max}\beta(z)}. \tag{4.72}$$

Previously if we wanted to determine the envelope of an ensemble of particles in a periodic lattice, we would have needed to apply matrix multiplication to each particle's initial coordinates to propagate it through the lattice. Now, we only need to determine $\beta(z)$ and $\epsilon_{\max}$.

Notice that we did not include dispersion here. This means our solution is the particular solution to Hill's equation. The complementary solution is included by simply adding $\delta D(z)$ and $\delta D'(z)$ to Equation 4.68 and 4.69 respectively.

4.7.1 Acceptance

Particles with larger values of ϵ will oscillate with larger amplitudes. Eventually a value of ϵ is reached that causes oscillation amplitudes that cannot be supported by the lattice. The value that this occurs at is known as the **acceptance** and from Equation 4.69 is given by

$$\text{Acceptance} = \epsilon_{\max} = \left(\frac{d(z)^2}{\beta(z)}\right)_{\min}, \tag{4.73}$$

where $d(z)$ is the aperture of the lattice at z.

An off momentum particle may have the same emittance as an on momentum particle, but its shifted orbit will change its position relative to the aperture which

can be viewed as a different sized aperture. We can replace Equation 4.73 with

$$\text{Acceptance} = \epsilon_{\max} = \left(\frac{(d(z) - \delta D(z))^2}{\beta(z)} \right)_{\min} \tag{4.74}$$

to account for chromatic effects. Asymmetric lattices can be accounted for by using two values for $d(z)$.

4.7.2 Tune

The phase of particle oscillations, Equation 4.68, increases monotonically with z according to Equation 4.66. In a periodic lattice the phase will grow by the same amount every revolution because of relation 4.70. As we will see, this periodic increase in phase is extremely important in the study of resonances in periodic lattices. To facilitate this a new variable, **tune**, is introduced and defined as

$$\mathrm{Q} = \frac{\Psi(z + L_p) - \Psi(z)}{2\pi} = \frac{1}{2\pi} \int_0^{L_p} \frac{1}{\beta(z)}, \tag{4.75}$$

where Equation 4.66 was used. The tune is simply the number of oscillations a particle undergoes within its envelope, $\sqrt{\epsilon\beta(z)}$, per unit cell. Because the tune is proportional to the total phase advance per element we have

$$\mathrm{Q} \propto \sum_i \Omega_i \propto \sqrt{\frac{\mu \overline{B_p}}{m}} \frac{L_p}{v}, \tag{4.76}$$

where $\overline{B_p}$ is the average magnetic pole face strength.

Particles with different longitudinal velocities will experience different tune

shifts per revolution. The absolute shift is

$$\Delta \mathrm{Q} = \mathrm{Q}_{\delta \neq 0} - \mathrm{Q}_{\delta = 0} = \mathrm{Q}\frac{v}{v + \Delta v} - \mathrm{Q} \approx |\delta|\mathrm{Q}. \tag{4.77}$$

This can be used to define a coherence length for a swarm of particles. A reasonable definition is the number of revolutions, N_c, for the tune to drift by a value of 1 which means the phase has drifted by 2π. For a swarm of particles with velocity width of about $2\Delta v$ the coherence length in terms of revolutions is then

$$N_{\mathrm{c}} = \frac{v}{|\Delta v|}\frac{1}{\mathrm{Q}} = \frac{1}{\mathrm{Q}|\delta|}, \tag{4.78}$$

where v is the design speed. As expected for $\Delta v = 0$ the coherence length is infinite. In reality the tune shift will depend on the initial position of the particles as well because of nonlinear forces.

4.8 Resonances

So far we have developed a linear theory of a storage ring. There are numerous considerations and methods to improve the realism of our model [91]. We will only consider perhaps the most important: resonances.

An otherwise ideal lattice can possess undesired forces which in a periodic system may be disastrous. Consider a short ideal lens installed with a small shift in the horizontal direction of Δx. This results in a new force

$$F_x = -K_m\left(x + \Delta x\right) = -K_m x - K_m \Delta x = -K_m x + F_{x0}, \tag{4.79}$$

where we have effectively introduced a small constant force of F_{x0}. This is known as a dipole error because it results from an unwanted dipole contribution in the case of charged particles. Because the lens is short this new force can be considered a small kick to the particle's angle resulting in $x' \to x' + \Delta\theta$. In general the effect of this kick will average to zero because the particle will arrive with a different phase each revolution. However, in the special case of an integer lattice tune, the particle will experience the force in a coherent manner. It will then continually receive a small increase of $\Delta\theta$ and its emittance may grow large enough that it is lost. Tune must be avoided when designing a lattice integer. Similar reasoning can be applied to higher order errors, with the next term being the quadrupole error from an unwanted lensing force. The resonance condition is then a half integer tune. Higher order terms have diminishing impact.

The derivation of the impact of resonances is rather tedious and involved. The approach is to determine the new equilibrium orbit after many revolutions under the influence of errors distributed around the lattice. We will simply write down the results for the leading dipole term,

$$u(z) = \frac{\sqrt{\beta(z)}}{2\sin(\pi Q)} \sum_i \sqrt{\beta_i}\theta_i \cos\left((\Psi(z) - \phi + \pi)Q\right), \tag{4.80}$$

where dipole errors are located at z_i, and θ_i and β_i are the dipole kick and beta function at z_i respectively. This result is only valid for dipole errors that can be approximate as discrete kicks. It also assumes the presence of damping and is the equilibrium orbit after several damping time constants. The relevant detail however is that the effects of dipole errors are amplified by the term $1/\sin(\pi Q)$, where it is

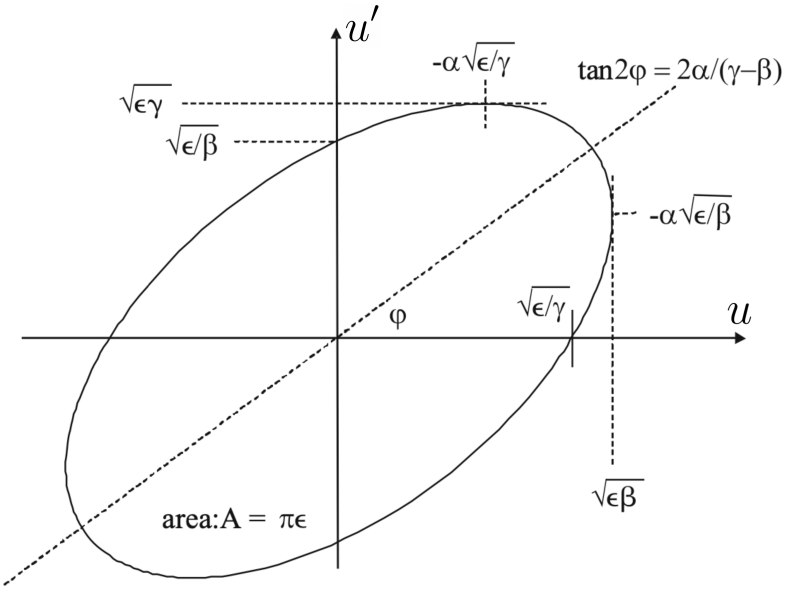

Figure 4.9: Constant emittance ellipse in phase space. In general the ellipse will have a different shape at each location z in the lattice.

clear that integer tunes and near integer tunes are to be avoided.

4.9 Phase Space

In the previous section we derived a relationship for the envelope a particle traces out in position space. We can extend this idea to a 2 dimensional phase space of x, x' and y, y'. Rather than an envelope, a particle will be confined to exist on the perimeter of an ellipse that periodically evolves through the lattice. Similiar to the amplitude envelope $\sqrt{\epsilon_{\max}\beta(z)}$, we require a relation that does not depend on the phase $(\Psi(z) + \phi)$ of the particle. Eliminating the phase in Equation 4.68 and

4.69, with some algebra we get

$$\gamma(z)u(z)^2 + 2\alpha(z)u(z)u'(z) + \beta(z)u'^2(z) = \epsilon, \tag{4.81a}$$

$$\alpha(z) = \frac{-\beta'(z)}{2}, \tag{4.81b}$$

and

$$\gamma(z) = (1 + \alpha(z)^2)/\beta(z), \tag{4.81c}$$

where the definition of $\alpha(z)$ has been intentionally repeated for reference. In this context the parameters $\beta(z)$, $\alpha(z)$, and $\gamma(z)$ are commonly referred to as **Twiss parameters**.

Equation 4.81 is the equation of an ellipse as plotted in Figure 4.9. A particle with emittance ϵ at location z will be located somewhere along the perimeter of the ellipse. In a periodic lattice a particle making multiple revolutions will in general arrive at location $z + NL_p$ at a different point on the perimeter of the ellipse than previous arrivals. If one plots the u, u' values of the particle at location $z + NL_p$, the ellipse in Figure 4.9 would be produced for sufficient N. This is also known as a Poincaré map and appears frequently in the study of dynamical systems. Note that a particle will only fully trace out the ellipse as $N \to \infty$ if the phase advances by an irrational number every L_p.

To account for off momentum particles we must include the effects of dispersion. An off momentum particle's emittance ellipse is relative to the its dispersive

orbit. A more general relation for emittance in the x dimension is then

$$\gamma(z)x_\epsilon(z)^2 + 2\alpha(z)x_\epsilon(z)x'_\epsilon(z) + \beta(z)x'^2_\epsilon(z) = \epsilon,$$

where

$$x_\epsilon(z) = x(z) - \delta D(z) \tag{4.82}$$

and

$$x'_\epsilon(z) = x'(z) - \delta D'(z).$$

We are often more interested in the emittance of a distribution of particles than of a single particle. There are several ways to quantify this. A common approach is to define the emittance statistically as

$$\epsilon_{\text{RMS}} = \sqrt{\langle u^2\rangle\langle u'^2\rangle - \langle uu'\rangle^2}, \tag{4.83}$$

where the total emittance is $\epsilon = 4\epsilon_{\text{RMS}}$ [94]. Another common approach is to use the emittance of the 90th percentile particle.

The relative density and temperature of a circulating beam can be related to the Twiss parameters. The transverse maximum position is given by $\sqrt{\epsilon\beta}$ and the maximum transverse velocity is $\epsilon\gamma$. Thus

$$n \propto \sqrt{\beta}$$

and (4.84)

$$T \propto \sqrt{\gamma},$$

where n is the particle density and T is the particle temperature.

4.10 Determining the Beta Function

Section 4.7 introduced a powerful technique for understanding the behavior of ensembles of particles with the derivation of Equation 4.68. However, Equation 4.67 could not be solved analytically and the value of $\beta(z)$ was left unspecified. Here we derive a method to determine the value of $\beta(z)$.

First, note the relationship

$$\epsilon = \gamma u^2 + 2\alpha u u' + \beta u'^2 = \begin{bmatrix} u & u' \end{bmatrix} \begin{bmatrix} \gamma & \alpha \\ \alpha & \beta \end{bmatrix} \begin{bmatrix} u \\ u' \end{bmatrix} = \mathbf{U}^T \mathbf{B}^{-1} \mathbf{U}, \tag{4.85}$$

where

$$\mathbf{B} = \begin{bmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{bmatrix} \tag{4.86}$$

is the beta matrix and we have dropped the explicit z dependence for notational simplicity.

The vector $\mathbf{U}(z)$ is transformed according to

$$\mathbf{U}_2 = \mathbf{M}_{12} \mathbf{U}_1 \tag{4.87}$$

where the subscripts $_1$ and $_2$ correspond to positions z and $z + \Delta z$ respectively. We can determine the relationship between $\mathbf{B}_1$ and $\mathbf{B}_2$ in terms of transfer matrices. Starting with

$$\epsilon = \mathbf{U}_1^T \mathbf{B}_1^{-1} \mathbf{U}_1, \tag{4.88}$$

applying some elementary matrix identities, and rearranging terms we can arrive at

$$\epsilon = \mathbf{U}_2^T \left(\mathbf{M}_{12}\mathbf{B}_1^{-1}\mathbf{M}_{12}^T\right)^{-1} \mathbf{U}_2, \tag{4.89}$$

which finally leads to

$$\mathbf{B}_2 = \mathbf{M}_{12}\mathbf{B}_1\mathbf{M}_{12}^T. \tag{4.90}$$

In the case of a linear lattice, if we can determine the value of the beta matrix at some point in the lattice we can then use Equation 4.90 to produce $\beta(z)$ at any other point. For an injection system this is easily done by backwards propagating $\mathbf{B}(z_{\text{inj}})$ from the ring through the injector, where z_{inj} is the injection location.

For a periodic lattice we have the relation 4.70, which leads to

$$\begin{gathered}\mathbf{B}(z + L_p) = \mathbf{M}_{\text{total}}\mathbf{B}(z)\mathbf{M}_{\text{total}}^T = \mathbf{B}(z) \\ \begin{bmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{bmatrix} = \begin{bmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{bmatrix} \begin{bmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{bmatrix} \begin{bmatrix} m_{11} & m_{21} \\ m_{12} & m_{22} \end{bmatrix}\end{gathered} \tag{4.91}$$

since the components of $\mathbf{B}$ are functions of only β. Some algebra then yields the desired relations

$$\beta = \frac{2m_{12}}{\pm\sqrt{4 - (m_{11} + m_{22})^2}} = \frac{2m_{12}}{\pm\sqrt{4 - \text{Tr}\,(\mathbf{M})^2}}, \tag{4.92a}$$

$$\alpha = \frac{m_{11} - m_{22}}{2m_{12}}\beta, \tag{4.92b}$$

and

$$\gamma = \frac{1+\alpha^2}{\beta}, \tag{4.92c}$$

where the stability relation Equation 4.53 that we derived by general arguments about eigenvalues has appeared in Equation 4.92a.

It is worth pointing out a minor subtlety in Equations 4.92. The matrix term m_{12} is not guaranteed to be positive, which could result in $\beta < 0$. This is non-physical because the particle's envelope is proportional to $\sqrt{\beta}$. Thus the sign of the term $\pm\sqrt{4 - \mathrm{Tr}\left(\mathbf{M}\right)^2}$ must be chosen correctly based on the sign of m_{12}.

4.11 Differences With Charged Particles

So far we have largely neglected discussion of charged particles for which the theory was originally developed. The major differences are

1. In Equation 4.34 the term $2\delta k_0$ is replaced by δk_0.
2. In Equation 4.34 v_o^2 is replaced by v_o.
3. Multipoles yield different forces for charged particles.
4. Charged particles produce radiation when accelerated which introduces intrinsic damping into the system.

The first 2 differences arise from the fact that a charged particle's velocity appears in the Lorentz force. For neutral atoms this would be similar to a situation where the magnetic field strength is multiplied by a factor of velocity. This would

result in every v^2 term being reduced to v. This implies that the behavior of a lattice for neutral atoms is more sensitive to changes in particle velocity.

The third difference results in different roles for multipoles than we saw for neutral atoms. Dipoles produce bending for charged particles and quadrupoles result in lensing. However, there is a difficulty in charged particle lensing that does not occur in neutral atom lensing. The magnetic field from an upright quadrupole (Equation 2.17) is

$$\vec{B} = -gy\hat{x} - gx\hat{y}, \tag{4.93}$$

which gives a Lorentz force of

$$\vec{F} = gv_z x\vec{x} - gv_z y\vec{y}, \tag{4.94}$$

where g is the field gradient with units T/m and a positive charge is assumed. Notice that the force is focusing for y and defocusing for x. It can be shown that a series of quadrupoles of alternating polarity will result in overall focusing of the beam. This approach is known as alternating-gradient focusing or strong focusing for historical reasons. Higher order multipoles and combinations of multipoles also find roles in accelerators. Sextupole and octupole magnets can be judiciously used to reduce chromatic aberrations and can correct for some perturbations. See [91] for more details.

The fact that charged particles emit radiation when accelerating is of great advantage to charged particle storage rings. This results in particle dampening that counteracts unwanted particle deviations to some extent. However, a downside is

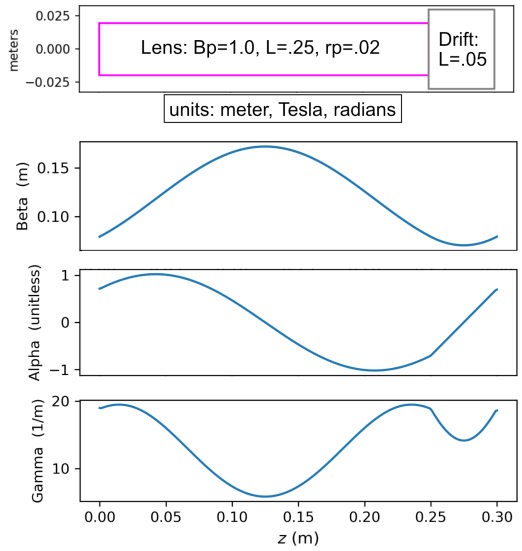

Figure 4.10: Graphical depiction of a periodic linear lattice and the profiles of Twiss parameters through the unit cell. The parameters repeat periodically.

that particles cannot circulate without being continually re-accelerated because they lose energy at every turn.

4.12 Comparison With Time Stepping

4.12.1 Linear Lattice

We can compare the theory we have developed to simulated results of ideal elements. First, consider a simple periodic lattice composed of a lens followed by a drift element. Figure 4.10 shows a graphical representation of the lattice along with its Twiss parameters.

Next, a swarm of particles is initialized on a grid in phase space in the x dimension and numerically time-stepped through the lattice. Results are shown in

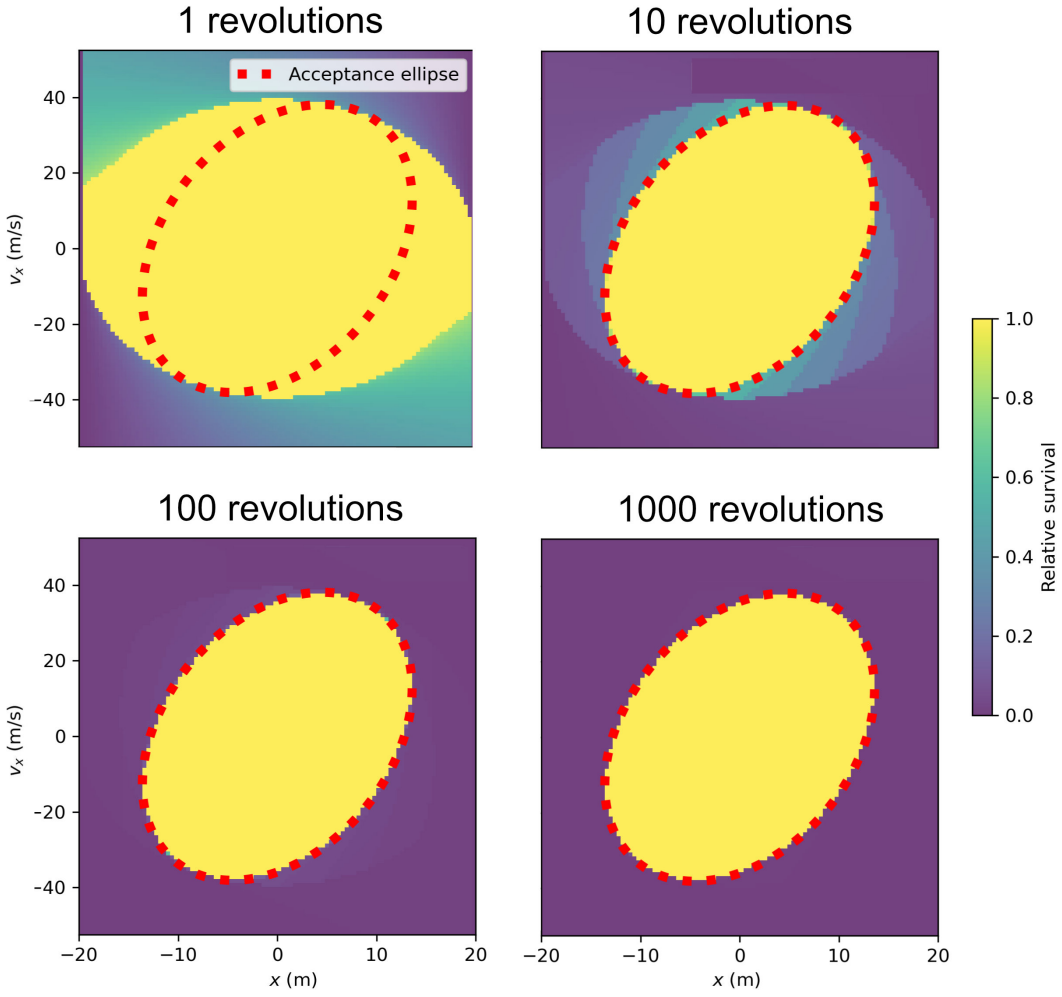

Figure 4.11: Plots of relative survival through the periodic lattice in Figure 4.10 versus initial phase coordinates. Revolutions are passes through the unit cell. Relative survival is the fraction of revolutions a particle makes of the maximum possible. For high numbers of revolutions, there is good agreement between theory and prediction. For low numbers more revolutions are required for high emittance particles to collide with an aperture.

Figure 4.11. After a sufficient number of revolutions, the simulated survival agrees with the predicted survival. Multiple revolutions are required because though many particles may exceed the acceptance, it will take in general multiple revolutions for the particle to actually collide with the limiting aperture. Results for x and y are visually indistinguishable because the Twiss parameters are identical in both dimensions and the effects of gravity are negligible.

It is not surprising that there is such excellent agreement between theory and simulation for a linear lattice. When deriving the equations of motion (Section 4.5) we made numerous approximations. However, in the case of a linear lattice where $k_0 = 1/r_0 = 1/\infty = 0$, most approximations are not required.

4.12.2 Circular Lattice

We saw excellent agreement between predicted and simulated results for a linear lattice in the previous section. Now we turn our attention to the more challenging model of a circular lattice. A graphical representation of the lattice along with its Twiss parameters is shown in Figure 4.12. Once again, a swarm of particles is generated and time stepped through the lattice. Results are shown in Figure 4.13.

In the y direction there is still excellent agreement between theory and simulation, however in the x direction this is not the case. A region of particle loss has appeared within what is predicted to be a region of survival. It also turns out that the shape and location of this region is dependent on the specifics of the lattice.

Consider Figure 4.14, which depicts the survival as a function of the length

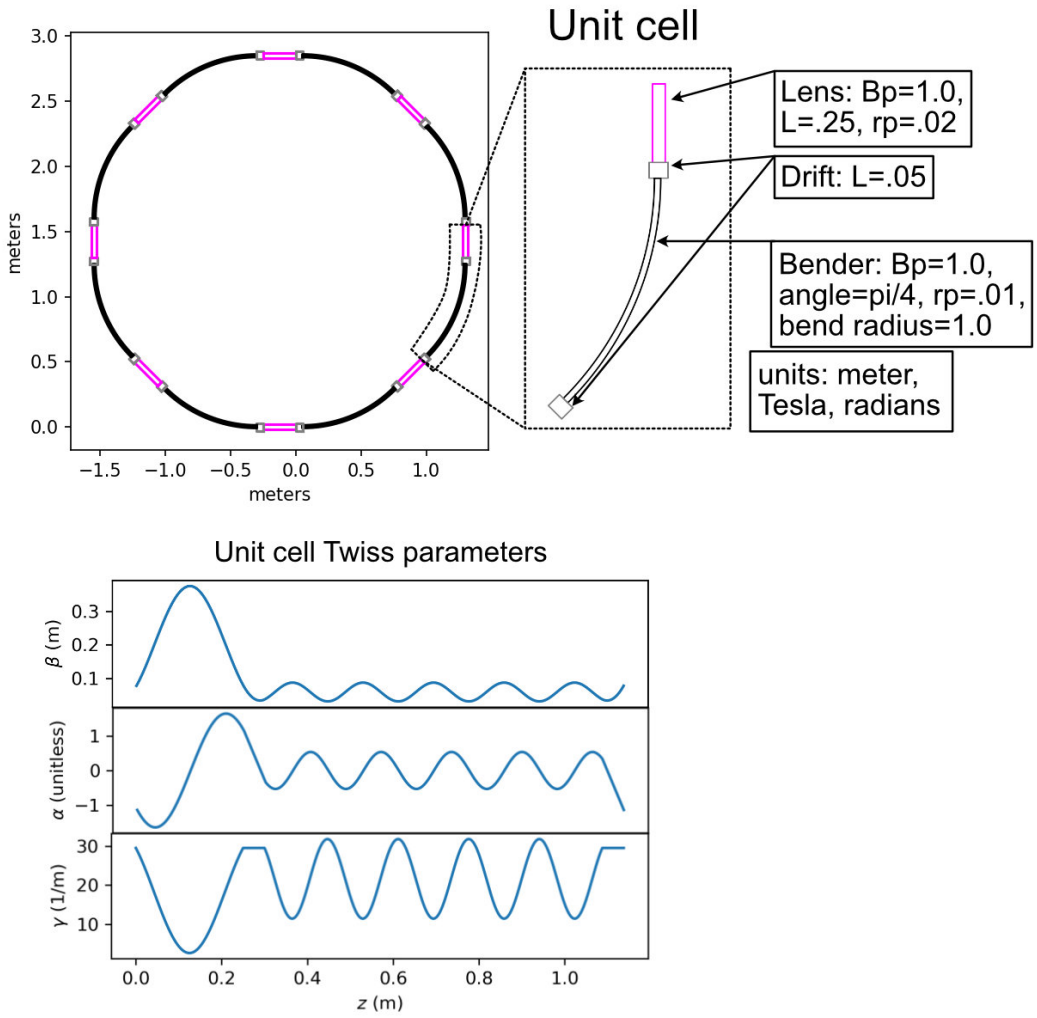

Figure 4.12: Graphical depiction of periodic linear lattice and profiles of the Twiss parameters through the unit cell. The lattice is composed of 8 unit cells, each deflecting a particle by 45 degrees.

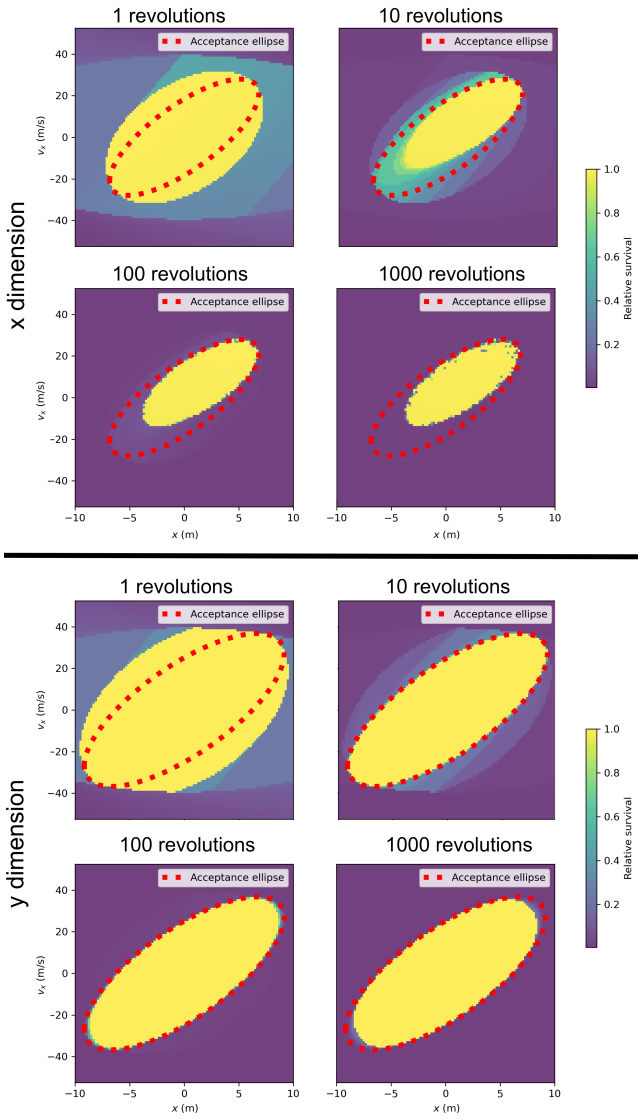

Figure 4.13: Images of relative survival through the periodic lattice in Figure 4.10 versus initial phase coordinates in the x and y dimension. Revolutions refers to passes through the unit cell. Relative survival is the fraction of revolutions a particle makes of the maximum possible.

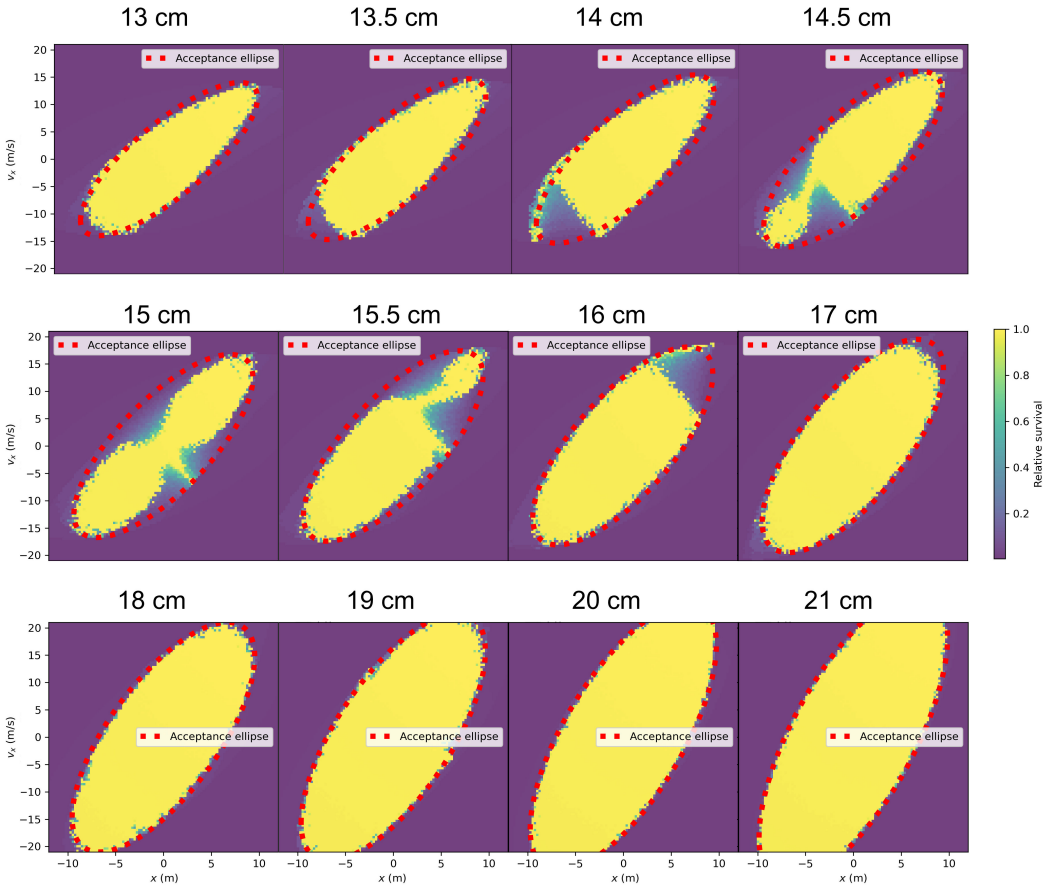

Figure 4.14: Plot of particle survival as a function of lens length in the lattice shown in Figure 4.12. Notice the appearance of non-linear effects for a range of lens lengths. Also note that the acceptable particle emittance increases. Though not shown, it eventually decreases.

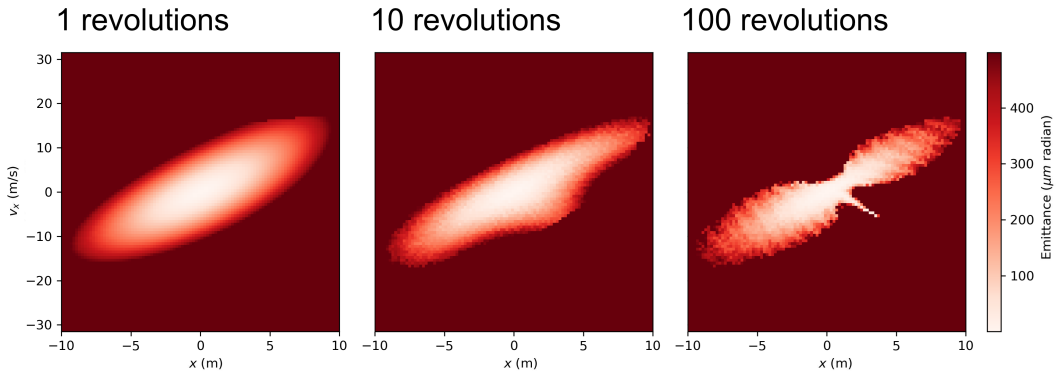

Figure 4.15: Particle emittance versus initial phase space coordinates for the lattice in Figure 4.12. Revolutions refers to passes through the unit cell. Nonlinear forces result in a change in emittance not explained by linear theory.

of the lens. Only within a range of lengths is the disagreement between predicted and simulated results present. The disagreement moves through the phase space ellipse as a function of lens length. Notice that longer lens lengths allow for the acceptance of particles with larger emittance values. The trend eventually reverses. This hints at an optimization problem if one is interested in maximizing the volume of phase space that can be accepted. This is a form of a mode matching problem.

Next, consider the behavior of the emittance of each particle as shown in Figure 4.15 as a function of revolutions for a lens length of 15 cm. Recall that we defined the emittance to be a constant, yet in Figure 4.15 it is clearly increasing for many particles. This is a nonlinear effect from the presence of bending that is not accounted for in our linear theory.

A comparison of the stability behavior of the circular lattice is shown in Figure 4.16. Particles are initialized with zero transverse phase space coordinates

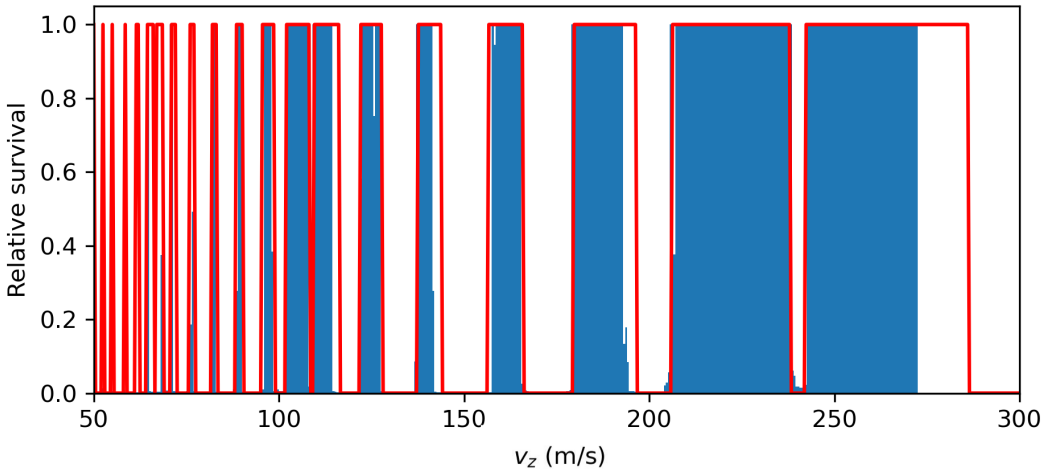

Figure 4.16: Particle survival versus longitudinal velocity. A spectrum of deadbands, regions with no survival, is present and is a general feature of lattices.

and traced for 10 seconds. Over a wide range of longitudinal velocities there is good agreement between simulation and theory. For higher velocities a cutoff occurs around 270 m/s that is not explained by stability criteria. However, accounting for dispersion yields a cutoff of about 270 m/s. Notice that the regions of stability are narrower and closer together at lower velocities. This occurs because unstable regions occur periodically with respect to tune and tune is inversely proportion to longitudinal velocity (Equation 4.76). The unstable regions are analogous to deadbands in a transfer function.

Finally we can compare the predicted survival to the simulated value for a swarm of particles in phase space. Here we use a swarm of 1000 particles pseudo-randomly distributed in phase space. Transverse particle positions are distributed in a circle of 10 mm diameter, transverse velocities are distributed in a circle of 20

m/s diameter, and longitudinal velocities are distributed between 200 and 220 m/s. Particles are traced for 10 seconds. Theory predicts 56% of particles survive while simulation indicates 49% survive.

4.13 Different Species

Assuming our approximations hold, we can use the preceding theory to design a neutral atom storage ring for an atomic or molecular species. An obvious question is whether other species can be used. As far as the physics of the storage ring is concerned, only a particle's mass, velocity, and magnetic moment are relevant. If these parameters can be varied in a manner that does not change the equations of motion, then the answer is yes.

From Equation 4.34 the general equation of motion is

$$u''(z) + \left(\pm \frac{2\mu B_p(z)}{m v_o^2 r_p(z)^2} + 3k_0(z)^2 \right) u(z) - 2\delta k_0(z) = 0, \tag{4.95}$$

where it is assumed the magnetic moment μ is a constant. m and μ are characteristics of the species, while v_o is characteristic of the source of the species. Our previous discussion of the design orbit implied that the speed v_o is a feature of the lattice, but this is actually not the case. Recall that the design orbit in a bending element is determined by balancing the centrifugal pseudo force and magnetic force. From Equation 4.9 the particle parameters that defined this radius all appear in the term $2mv_o^2/\mu$, which also appears in our equation of motion.

Therefore we can write down a species dependent factor for which a specific

lattice is designed to operate with. This factor is

$$\xi = \frac{2\mu}{mv^2} = \frac{\mu}{E_k}, \tag{4.96}$$

where μ is the magnetic moment of the species and is assumed to be constant, m is the mass of the species, v is the longitudinal velocity of the species, and E_k is the particle's kinetic energy. The subscript was dropped from v_o. We will refer to this factor as the **species factor**. A given lattice designed with a specific ξ can be used with any species/source combination that results in the same ξ.

A possible source design is one which generates a thermal particle distribution to load a storage ring. The storage ring samples some half width Δv about some point v_k which is typically the peak or mean of the distribution. These can be related to temperature and mass by

$$\begin{aligned} v_k &= C_k\sqrt{T/m} \\ &\text{and} \\ \Delta v &= \beta\sigma = \beta C_\sigma\sqrt{T/m}, \end{aligned} \tag{4.97}$$

where C_k and C_σ are constants that depend on the source characteristics (such as effusive, supersonic, etc), σ is the variance squared, and β is a unitless number. Here thermal distribution is being used in a general sense to mean any distribution were some central value of the longitudinal velocity is proportional to $\sqrt{T}$. Plugging this

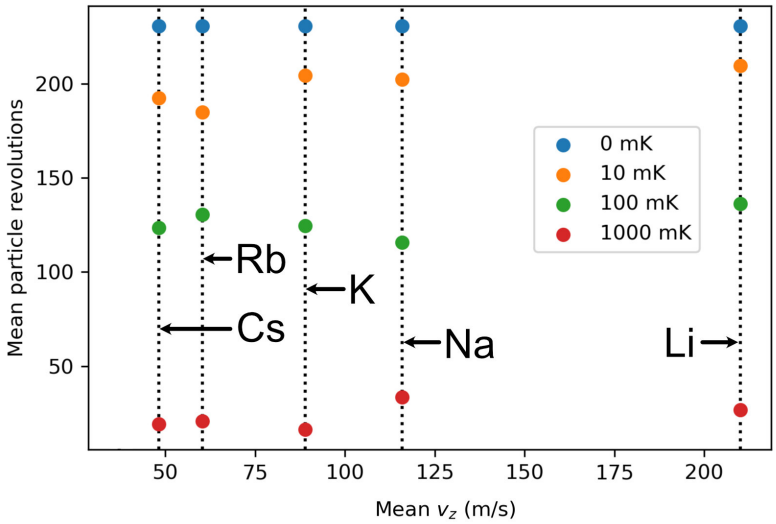

Figure 4.17: Mean revolutions of a swarm of particles through a periodic lattice versus initial longitudinal velocity. The stable alkali atoms are shown. Several initial temperatures are shown. Revolutions are scaled by $1/v_z$ relative to Li for comparison. Maximum revolutions possible is 230.8 for Li.

into δ and ξ gives

$$\xi = \frac{2u}{C_k^2 T}$$

and (4.98)

$$\delta = \frac{\Delta v}{v_k} = \beta \frac{C_\sigma}{C_k},$$

which now only depend on μ and T.

The equations of motion for a thermal source then only depend on the temperature of the source and the magnetic moment. This is a very convenient property. For example, a storage ring loaded with a source of neutral alkali atoms at temperature T could be as efficiently loaded with another source of a different species of alkali atoms at the same temperature.

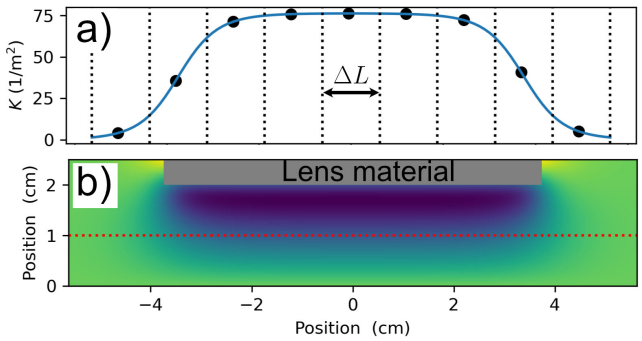

Figure 4.18: Modeling a lens as short slices to account for fringe fields.(a) The spring constant $K = 2\mu B_p/mv^2r_p^2$ versus position in the lens along the red dashed line in image b. The black circles indicate the location of the transfer matrix produced from ΔL and each averaged value of K. (b) Heatmap of the transverse force applied by the lens.

Let us test this model in the periodic lattice shown in Figure 4.12 for the ground state alkali atoms. Figure 4.17 shows the mean number of particle revolutions for several temperatures of swarms launched from a point source within the lattice. 100 particles are sampled from a thermal distribution and time stepped for 10 seconds. There is good agreement over a wide range of masses. In Section 6.6 we will see that this phenomena holds for a more realistic lattice.

4.14 Numerical Methods

Real magnetic elements do not obey the hard edge approximation as their fields extend past their input/output planes. This model can be improved by taking into account the effective length described in Section 3.4.4. The effective length is used in place of the material length for the transfer matrices (Equation 4.38). A further improvement is to model the element as small slices and construct a transfer

matrix for each slice as shown in Figure 4.18 for a lens of length 7.5 cm and bore radius of 2 cm. The total matrix is the product of each slice. In the case of a lens this is done by determining the average value of K in each segment of length ΔL and producing a transfer matrix from K and ΔL.

For the lens shown in Figure 4.18 the transfer matrices for the three methods are

$$\begin{aligned}
&\text{Naive model: } \mathbf{M} = \begin{bmatrix} 1.337 & 0.1285 \\ 6.138 & 1.337 \end{bmatrix}, \\
&\text{Effective length model: } \mathbf{M} = \begin{bmatrix} 1.337 & 0.1283 \\ 6.130 & 1.337 \end{bmatrix}, \\
&\text{and} \\
&\text{Sliced lens model: } \mathbf{M} = \begin{bmatrix} 1.318 & 0.1214 \\ 6.072 & 1.318 \end{bmatrix},
\end{aligned} \tag{4.99}$$

where the matrices were computed for field values on or along the red dashed line in the figure. All three matrices are fairly similar, but their differences could become significant in a periodic system. Note that for the effective length and naive model two drift regions needed to be included at the entrance and exit to match the region used for the sliced lens model. Also note that because of field aberrations the values will change depending on the location of the red dashed line.

We have neglected to consider elements in which the bending radius varies, which occurs in the combiner. If the combiner is viewed as many short bending sections with each neighbor's radius changing by a small amount, it can be easily modeled numerically with the above method.

An additional method is to expand the transfer matrices to include mappings for higher order terms of $\tau^{n_1+n_2+n_3} = x_1^n y_2^n \delta_3^n$. We discarded terms higher than τ^1. They can be included, but the linear nature of the transfer matrices is lost.

Chapter 5

Simulation Overview

5.1 Introduction

The construction of a permanent magnet meter-scale neutral atom storage ring is a significant experimental undertaking. To the best of our knowledge it has never been done before and the theory developed in Chapter 4 is somewhat novel. Additionally, it has been proposed that it may not be possible to find a stable configuration [9]. Therefore to reduce the possibility of failure, a full simulation[1] of the storage ring was developed in Python, taking advantage of the Numba library [95] to increase performance.

A lattice can be specified as a series of elements such as drift regions, lenses, benders, and combiners. Idealized models can be used, or in the case of permanent magnet arrays, analytically computed magnetic fields can be specified. Results from a Finite Element Method (FEM) simulation, such as COMSOL Multiphysics, can be provided for iron based elements. For faster yet less accurate evaluations, approximate symmetry can be exploited to reduce the computational overhead of generating magnetic fields. Specified assembly tolerances and magnetic material

[1] The codebase can be found on GitHub (link: "https://github.com/BillyTheKidPhysicist/storageRing.git"). At the time of writing the commit ID is "d53c3b3c83f1e16559880ec7a681fdfaffb3775e".

imperfections can be included. Particles are time-stepped through the lattice, and phase space coordinates and energies can be logged at each step. The effects of Li-Li collisions can be approximately accounted for as well.

A specific storage ring design is created by coupling a linear injector lattice and circular ring lattice. The design is parameterized by variables such as element length or radii, typically 10-20 in number, and then used to find an optimal solution. This procedure can be quite computationally intensive considering each solution takes 1-10 minutes when exploiting symmetry. When including the effects of misalignments and material imperfections, evaluation time is on the order of an hour. An allocation at the Texas Advanced Computing Center (TACC) for the Lonestar6 supercomputer was awarded which aided modeling efforts. Each node of this computer has 128 processors and only a single node was needed. In the following chapter technical details of the simulation are described, as well as some brief theory of numerical methods.

5.2 Particle Time Stepping

They are a wide variety of time stepping procedures available [96]. Ultimately the velocity Verlet method was chosen for its simple and efficient implementation and favorable performance compared to other methods. Let us briefly review some of the theory behind time-stepping.

We will only consider solving first order Ordinary Differential Equations

(ODE) of the form

$$\frac{du(t)}{dt} = \dot{u} = f(u(t)), \tag{5.1}$$

for which a wide variety of well-established methods exist. In general, the methods yield procedures for mapping $u(t) \rightarrow u(t+h)$ where h is a sufficiently small value. The term h is commonly constant which allows us to define

$$u_n := u(nh). \tag{5.2}$$

Newton's equation of motion is a second order ODE given by

$$\ddot{x} = \frac{F(x)}{m} \tag{5.3}$$

and can be transformed into the two first order ODEs

$$\dot{x} = p/m \tag{5.4a}$$

and

$$\dot{p} = F(x), \tag{5.4b}$$

where we have assumed $\dot{m} = 0$. This can now be solved by numerical methods for Equation 5.1.

5.2.1 Euler Method

The simplest integration scheme for Equation 5.1 is known as the Euler method, or more precisely, the explicit Euler method. Tayloring expanding $u(t)$ to

first order to gives

$$u_{n+1} = u_n + f(u_n)h + O(h^2). \tag{5.5}$$

To solve Equation 5.1 up to time T one starts with an initial value of u_0 and applies the above procedure for T/h iterations. For Newton's equation we have

$$x_{n+1} = x_n + \frac{p_n}{m}h \tag{5.6a}$$

and

$$p_{n+1} = p_n + F(x_n)h. \tag{5.6b}$$

Ignoring any computer precision concerns, our solution is only an approximate one because of the discretization of time and the loss of higher order terms in Equation 5.5. In general it is impossible to know how much our approximate solution deviates from the exact solution because we do not have access to the exact solution. Instead, it is common to characterize error by the terms neglected in our solution method.

In the case of the Euler method it is clear from Equation 5.5 that the error per step, or local error, is

$$\epsilon_{\text{local}} = O(h^2). \tag{5.7}$$

In contrast to the local error, we are also concerned with the global error which is the order of the error accumulated over N steps. The global error is then

$$\epsilon_{\text{global}} = \epsilon_{\text{local}} N = \epsilon_{\text{local}} \frac{T}{h} = O(h). \tag{5.8}$$

The stability of a numerical method can be quantified either in a general

way, or for a specific problem. We are primarily concerned with harmonic behavior, expressed by

$$\ddot{x} = -\omega^2 m x, \tag{5.9}$$

for which the time stepping procedure according to the Euler method is

$$x_{n+1} = x_n + \frac{p_n}{m} h \tag{5.10a}$$

and

$$p_{n+1} = -\omega^2 m h x_n + p_n. \tag{5.10b}$$

This can be written in matrix form as

$$\mathbf{X}_{n+1} = \mathbf{M}\mathbf{X}_n. \tag{5.11}$$

If N iterations are carried out, then the solution is $\mathbf{X}_f = \mathbf{M}^N \mathbf{X}_i$. In analogy with Section 4.6.1, $\mathbf{X}_f$ will diverge if the magnitude of the eigenvalues of $\mathbf{M}$ are greater than 1. Some algebra yields for the eigenvalues

$$|\lambda| = \sqrt{1 + h^2\omega^2}, \tag{5.12}$$

which does not satisfy $|\lambda| \leq 1$ for non-trivial values of h and ω. Thus, the Euler method is unconditionally unstable for solving a harmonic oscillator.

It can be shown that in general the Euler method exhibits poor stability and energy conservation. For these reasons the method is rarely used outside of a pedagogical role.

5.2.2 Implicit Euler Method

A more practical method is the implicit Euler method, given by

$$u_{n+1} = u_n + f(u_{n+1})h, \tag{5.13}$$

where the only difference is that we have replaced $f(u_n)$ with $f(u_{n+1})$. In general, each iteration must be solved with an iterative technique such as Newton's method.

Compared to the explicit Euler method, the implicit method has superior energy conservation and stability behavior. Implicit methods typically have superior performance compared to their explicit counterparts for a given step size. However, the extra task of solving Equation 5.13 may outweigh the benefits of simply reducing the step size of the explicit method.

5.2.3 Velocity Verlet

The method chosen for the simulation is the velocity Verlet algorithm, a common algorithm for solving Newtons' equation of motion. It is given by

$$x_{n+1} = x_n + \frac{p_n}{m}h + \frac{1}{2}F(x_n)h^2 \tag{5.14a}$$

and

$$p_{n+1} = p_n + \frac{F(x_n) + F(x_{n+1})}{2}h. \tag{5.14b}$$

Determining the local and global error is more involved than for the Euler method, though it follows a similar procedure of Taylor expansions. The results

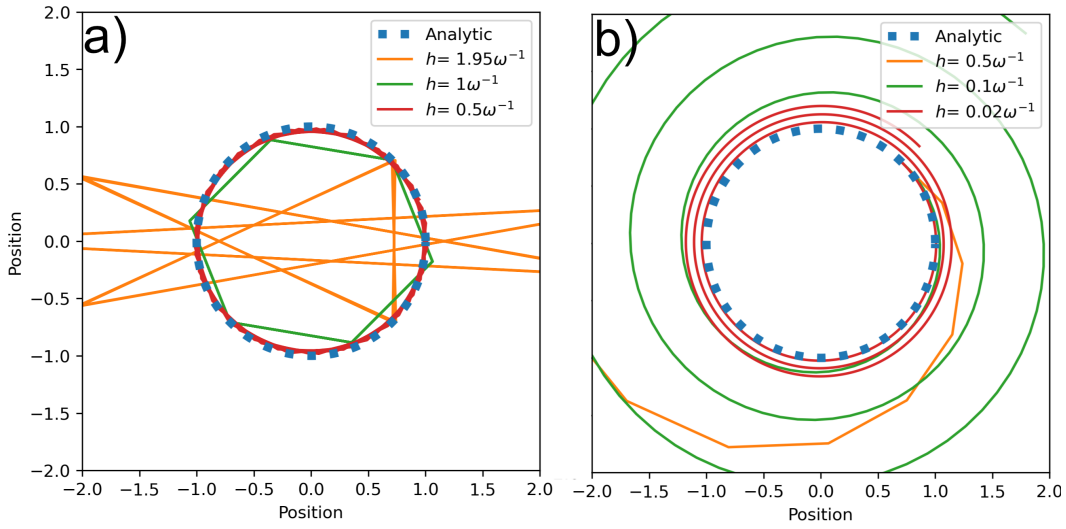

Figure 5.1: Trajectories in positions space for a radial harmonic oscillator with oscillation frequency ω. The timestep h is the specified fraction of ω. (a) Trajectories produced by the Euler method. (b) Trajectories produced by the velocity Verlet method.

yield a local error in position and velocity of $O(h^4)$ and $O(h^2)$ respectively, and a global error of $O(h^2)$. The stability criteria for a harmonic oscillator is given by the same procedure we used in Section 5.2.1 and yields

$$h\omega \leq 2, \tag{5.15}$$

which is conditionally stable. A comparison between the Euler method and the velocity Verlet method is given in Figure 5.1 for a harmonic oscillator. The velocity Verlet method also possesses the desired feature of being symplectic. Loosely speaking, an integrator is symplectic if it conserves an approximate form of the Hamiltonian.

5.3 Magnetic Forces

The gradient of the norm of the magnetic field is computed by forward or central difference of the norm of the magnet field. This is then used to compute the force with Equation 2.16. Magnetic fields are calculated depending on the element. In the case of an idealized component Equation 2.20 is used for the magnetic field norm.

For ideal cuboidal permanent magnets analytic solutions exist. Thus the magnetic field produced by a Halbach array of cuboid magnets can be solved for exactly. This is fortunate because the computation time for the same results by the Finite Element Method (FEM) can be significantly longer and implementation would be more cumbersome. An excellent library exists for generating magnetic fields from analytical results [97]. Fields can be generated from dipoles, rectangular and cylindrical sheets of current, and lines of currents.

The FEM is typically used to deal with magnetic models that include soft materials. However, one can account for soft materials with the magnetostatic Method of Moments (MoM) [98]. This method works by modeling the magnetized material of the system as N domains with each domain represented by a magnetic dipole. In the case of cuboid magnets, the magnet can be divided into smaller cuboids or simply left as a single domain. The magnetization of each dipole is then adjusted to satisfy Maxwell's equations and the B-H curve which in general requires an iterative solver. As discussed in Section 2.6.1, a permanent magnet can be more accurately modeled by modifying the square B-H curve with the addition of a linear permeability term. In this case the system results in an NxN matrix that can be

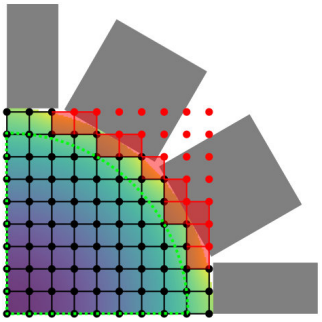

Figure 5.2: A diagram of the linear interpolation scheme and the limit it sets on the usable interior region. The green semicircle indicates the radius of the valid interpolation region. The radius of the semicircle increases with decreasing grid spacing

easily solved to give the new magnetization values. The results are valid as long as the demagnetizing fields are small enough that the B-H curve is still in a linear regime. A downside of this method is that the complexity scales as N^2, whereas FEM complexity scales as N. However, MoM is more easily implemented

The material and dimensional tolerances of permanent magnets are accounted for by modifying the relevant values of each cuboid. The deviation is sampled from a uniform distribution bounded by the tolerance.

Magnetic force calculations are computationally expensive and the simulation would be unusable if they were required at each time step. Thus an interpolation scheme is used. A natural and efficient choice is linear interpolation which in 3D is referred to as trilinear interpolation. A Numba compliant version was implemented.

A downside of trilinear interpolation is that a grid is required rather than a scattering of points that can conform to a geometry. An implication of this is that the grid spacing sets the maximum usable region of the element's bore, otherwise

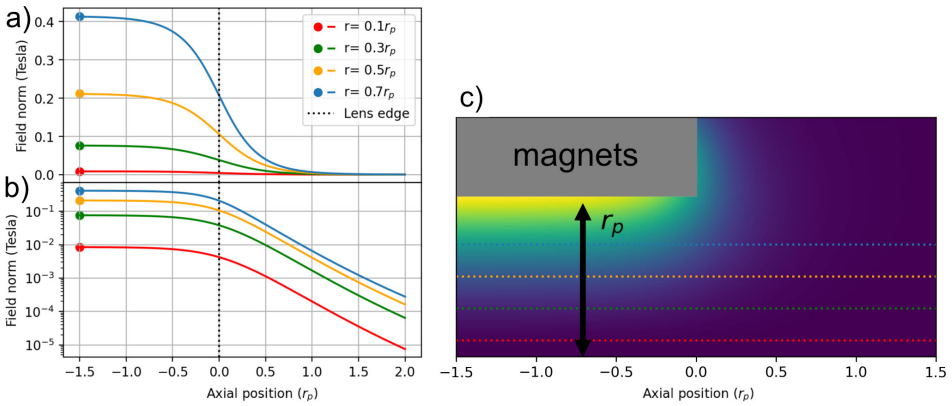

Figure 5.3: Magnetic field norm versus distance through a hexapole Halbach array lens. Distances are normalized to the bore radius r_p. The total length of the lens is $10r_p$. (a) Magnetic field norm plotted against the distance along 4 lines of different radial displacement from the bore centerline. The dots indicate the field value at the center of the lens. (b) A log scaled plot of (a). (c) A heatmap of the magnetic field norm of the lens. The dashed lines correspond to the lines in plots (a) and (b).

the interpolation would include fields inside the magnetic material, as shown in Figure 5.2. Fortunately, feasible grid sizes result in regions that are larger than restrictions imposed by vacuum tubes. The grid cell size must be chosen such that it is sufficiently small compared to the bore radius of the element. An alternative interpolation scheme could be implemented as an improvement, such as barycentric interpolation, but I found trilinear to be adequate.

In principle a single magnetic component's fields affect a particle's trajectory everywhere in space and thus a fully accurate model would include the effect of all magnetic components at every interpolation point. This would result in unreasonable computation times. In practice, at sufficiently large distances the impact becomes negligible. Magnetic field values for a hexapole lens drop to less than a percent of their peak values for distances larger than $1.5r_p$ from the lens face, as shown

in Figure 5.3. This can be used to exploit symmetry and reduce computation time. Elements are placed such that their interpolation region extends to a distance $1.5r_p$ and no overlap between interpolation regions occurs. Thus two elements of bore radius r_{p1} and r_{p2} would be separated by a distance of at minimum $1.5(r_{p1} + r_{p2})$. Since each element is considered independent of its neighbors, symmetry within each element is used. this works out nicely because gaps between magnetic elements exist for other reasons such as vacuum or optical access.

The minimum cross sectional symmetry of a 12 magnet hexapole Halbach array is a wedge of 15° which can reproduce the field everywhere in the cross section using mirroring, reflecting, and rotating. However, this symmetry can only be exactly used inside of an infinitely long lens. If long enough, a lens can be broken into two regions of symmetry to a high degree of accuracy. This is done using two regions: an outer region representing the input/output of the lens that extends $2r_p$ in and $1.5r_p$ out of the lens, and a cross section that is applied inside for depths $\geq 2r_p$. Additionally, since trilinear interpolation is being used, 15° would not produce the required interpolation grid point placement. Instead, a 90° symmetry is used as shown in Figure 5.2.

In the case of a bending element a similar approach is used, except that the interior is not accurately modeled by a cross section. The unit cell of half of a bending magnet slice is exploited for interior symmetry instead. The fringing region is accounted for by an interpolation region that extends from $1.5r_p$ outside the first bending magnet to 2 unit cells inside.

Approximating magnetic elements as being independent of each other to

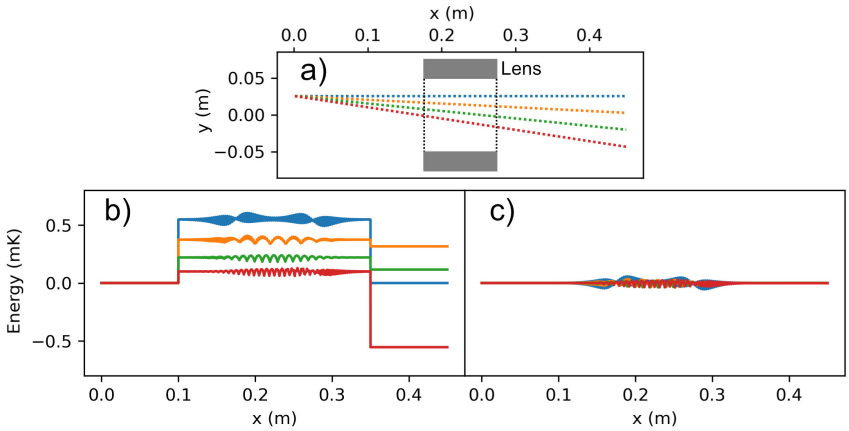

Figure 5.4: A comparison of the energy conservation behavior of the two models of the magnetic field. The total energy is calculated along 4 lines through the lens shown in (a). The energy at each point is calculated by taking the cumulative sum of the total work and potential energy of the previous points. (b) The energy along the 4 lines for the model where the magnetic fields are truncated 1.5 bore radii from the lens face. (c) The same plot as (b) except the longer range energy model is used.

utilize symmetry does negatively affect simulation realism, particularly in the form of particles gaining or losing excessive amounts of energy. However, the approximation is useful for optimization. Results can be corroborated by using a more realistic interpolation model that does not use symmetry. Instead, the magnetic field is interpolated everywhere inside the element and includes the fields of neighboring elements up to some maximum distance. Results such as total survival time or revolutions of a swarm of particles agree within a factor of two typically for simulation times of 100 seconds. Results diverge more strongly for longer times. When analyzing heating or phase space behavior the results can differ significantly. A simple comparison of the two models is shown in Figure 5.4 for a single lens.

5.4 Building Lattices

Assembling the lattice is conceptually straight forward. In the case of a linear lattice this is done by defining the initial location and angle of the first element, then adding the following elements. Elements line up by connecting their design orbit so that the ideal particle will follow this orbit exactly (barring numerical or other deviations). For lenses and drifts the design orbit is straight down the bore center. For a bender the design orbit is the orbit that a particle is in equilibrium with the centrifugal pseudo-force and the magnetic force (Section 4.2.1). For an ideal bender this is calculated analytically, while for the segmented bender this is the orbit that results in minimal RMS particle amplitude. This is determined by a minimization routine. There is always some oscillation because of the segmented nature of the bender, but a particle following the design orbit oscillates with an amplitude on the order of microns.

For closed lattices the procedure is the same as for linear lattices with the added constraint that the final element's position and output angle aligns with the initial element. For certain configurations this can be done rather easily analytically. For example, the simplest realistic closed lattice of two ideal benders and a single combiner can be reduced to a geometry problem of aligning 3 triangles. This requires leaving parameters of the lattice unspecified such as bending radii and element lengths to be later determined by constraints. Any convex closed lattice of ideal benders and combiners can be solved analytically. However, this can be tedious and does not extend to the inclusion of segmented benders because they cannot take on any bending radius and angle once the segment length is defined.

The problem is framed as a minimization problem to reduce a cost function instead. The cost is defined as the misalignment of the beginning and end of the lattice. An additional cost is applied to enforce that the final bending radii are very similar to the desired bending radii. The parameters that are varied are the bending radii, number of lenses in the benders, and the length of lenses. The user can specify in which elements the parameters can be varied to achieve a closed lattice. Typically there are multiple similar configurations that achieve a closed lattice.

5.5 Optimization

A reasonable figure of merit for a storage ring lattice is the flux multiplication F which is how many times the input flux (atoms/second) is amplified within the ring. This is defined as

$$F = \frac{\sum_i^N R_i}{N}, \tag{5.16}$$

where R_i is the number of revolutions of the ith particle before loss and N is the total number of particles. Probabilistic loss mechanisms such as vacuum and optical pumping can be included by weighting R_i which is discussed in Section 6.7. This value is defined for a specific collection of particles so a representative sample size must be used to characterize the lattice.

The efficiency of a lattice can be defined to be

$$\epsilon = F/F_{max}, \tag{5.17}$$

where

$$F_{max} = \frac{\sum_i^N v_i T / L_t}{N}, \tag{5.18}$$

T is the maximum simulation time, v_i is the longitudinal velocity of the ith particle and L_t is the total length of the lattice.

Neglecting losses not intrinsic to the ring such as vacuum or optical pumping, there are three phenomena that will prevent $\epsilon = 1$ for a given initial distribution:

1. Issues with the simulation such as too coarse of time stepping or magnetic field interpolation.

2. Particles starting in regions that collide with the vacuum walls almost immediately and would have been lost even with perfect stability.

3. Particles being lost due to instability. In practice this could apply to every particle as the effects of small errors in the lattice accumulate.

Mechanism 1 is addressed by iteratively reducing the relevant parameters until the model results clearly converge. Parameters are selected that result in a balanced trade-off between simulation time and precision. Mechanisms 2 and 3 can be reduced by an optimization routine as was alluded to in Section 4.12.2. The two mechanisms are in tension with each other and one dominates the other at certain time scales. For example, if the initial swarm is spread out in phase space such that many particles are initially lost to mechanism 2, then with a short simulation time the optimization routine may produce a lattice configuration that does not yield

long particle lifetimes. On the other hand a very long timescale will be biased in the other direction.

A reasonable choice for the timescale is to use the vacuum lifetime as an upper limit. Typical trap lifetimes for alkali atoms are 10s to 100 seconds. I chose a simulation time of 100 seconds when optimizing. The current design results in about a 50 second vacuum lifetime (Appendix B).

There are numerous optimization schemes to choose from and it there is not a general method for selecting one. The most straight forward and naive method is a grid search. Assuming a model requires 5 minutes to evaluate, there are 10 parameters, and we use a fairly coarse grid of 5 points per dimension, then the total evaluation time would be about 13,500 CPU-hours. For 5 parameters the total time is only 4.3 CPU-hours. This is known as the curse of dimensionality.

We saw in Section 4.6.1 the existence of distinct regions of stability. This implies the existence of numerous local minima. The problem is therefore one of global minimization in contrast to local minimization. Many local optimizer schemes can be shown to eventually converge to the minima. This usually cannot be done with a global optimizer. Many global optimizers are heuristic strategies inspired by the natural world such as "ant colony optimization" or "simulated annealing".

A powerful and general global optimization strategy is Differential Evolution (DE), shown in Figure 5.5. The strategy is inspired by natural selection. A population is initialized in the parameter space with each member's DNA being the parameter values at its respective location in parameter space. The member's fitness is inversely related to the value of the cost function at that location. At each

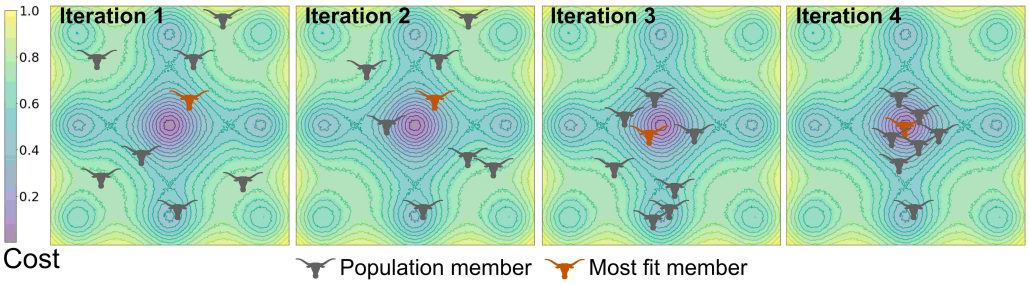

Figure 5.5: A simple example of the Differential Evolution (DE) algorithm. Members explore a cost function landscape and compete with each other to find a global minima.

iteration every member's DNA (the parent) is mixed with other members DNA and a new member (the offspring) is produced. If the offspring is more fit (less cost) than its parent it replaces the parent. A variety of recipes for mixing exist. A standard recipe is to randomly select one member and then take the difference of its DNA with the most fit member to produce new DNA to mix with the parent's DNA. The new DNA is randomly mutated and then segments are randomly inserted into the parent's DNA. Because we took the difference in DNA the population will eventually converge to a single DNA. Multiple generations are modeled by sweeping through the entire population. An example of DE is shown in Figure 5.5.

In the original DE algorithm the cost function for each member is evaluated serially for each generation. All the members must produce children before proceeding to another sweep through the population. The SciPy [99] library offers an implementation in which multiple cost functions are evaluated in parallel for improved performance. However, this parallel evaluation is applied for each generation. This can bottleneck efficiency for problems where the cost function evaluation

time varies across orders of magnitude. In the extreme all but one member of the population is rapidly evaluated. The remaining member uses one processor of the CPU and the rest are idled. This occurs when optimizing the storage ring system. For the initial iterations the vast majority of configurations result in either invalid geometries or the loss of all particles nearly immediately. A single long lived solution will then bottleneck the computer. To prevent this an asynchronous version of DE was implemented wherein members can evolve independently of the rest of population and different generations exist at the same time. This has proven very helpful when using the Lonestar6 super computer.

On a typical personal computer, the optimization time can still be quite long, on the order of days or even weeks. An alternative approach is to separate the problem into two problems, injection and circulation, which can significantly reduce the computation time. Additionally, simulating the injection system is more than an order of magnitude faster than simulating the ring system.

To simulate the injection system a reduced model, or surrogate model, of the ring system is still used to enforce the constraint that the injection system must not overlap with the ring system. The surrogate model also includes a lens after the combiner to encourage a more realistic result. A reasonable guess of element sizes is used. The figure of merit for the injector is then simply what fraction of particles survive through to the end of the combiner.

The optimal parameters from the injection system are then used to optimize the ring system. Finally, both systems are coupled together and the last step is to "polish" the results where the full system is optimized. This can be another global

optimization routine that starts from the optimal point and explores a parameter space with bounds shrunk around the optimal point, or a local optimizer strategy. Separating the problem in this way reduces the computation time to 2 or 3 days though it may get stuck in a local minima.

5.6 Miscellaneous

5.6.1 Performant Python

Python is an interpreted language rather than compiled, which loosely speaking means each line is recompiled each time. This makes the language a poor choice for problems which require many loops and functions, which is exactly what a time stepping simulation requires. However, many simulation problems are bottle necked by a small portion of the logic. In this case only a few hundred lines of code are the culprit amongst the roughly 10,000 lines. Rather than abandoning Python, and losing the features which have made it so popular, the problematic Python code can be replaced with high performance code. There are a few methods to do this. Historically the most popular is Cython which allows for python functions to be written in a mixed C/Python language. "Boost.Python" is another option that converts C++ modules to functions that can be called by Python.

An exciting alternative is the Numba library [95]. The library works by using a Just In Time (JIT) compiler that reduces Python code to high performance low-level code the first time the function is called. The library is used by simply adding a single line above a compliant python function and results in performance similar to code written in C. For reference, in a simple simulation of tracing a particle

through an ideal lens, the evaluation time per step is about 2500 ns without Numba and about 10 ns with Numba. For more information on high performance Python see [100].

5.6.2 Vacuum

A simple model of the vacuum in the ring and injector is used for analysis and optimization. Pressure profiles are calculated throughout the vacuum system resulting from outgassing, pumping, and gas flow. Surprisingly, there appears to be no publicly available Python libraries that do this. I implemented a simple version based on the fundamental vacuum equations briefly reviewed below. The fundamental vacuum equations are

$$\begin{aligned} &Q_{12} = C_{12}P_1 - C_{21}P_2 \\ &\text{and} \\ &Q = SP, \end{aligned} \tag{5.19}$$

where P is pressure, Q is the gas flow, C is the conductance, and S is pumping speed [79]. S and C have the same units of volume/time. The subscripts 1 and 2 indicate regions of vacuum separated by some restriction to gas flow. For an interconnected system of N vacuum regions this can be expressed with matrices as

$$\mathbf{Q} = (\mathbf{C} + \mathbf{S})\,\mathbf{P}, \tag{5.20}$$

where $\mathbf{Q}$ and $\mathbf{P}$ are length N vectors and $\mathbf{S}$ and $\mathbf{C}$ are $N \times N$ matrices. The system can solved directly if $\mathbf{S}$ and $\mathbf{C}$ are independent of $\mathbf{P}$, otherwise an iterative method

is required. The pressure profile in tubes can be determined with this method by modeling them as many shorter tubes. The effects of outgassing can be included in a similar fashion

5.6.3 Acknowledgements

I would like to acknowledge the following open source packages for playing a significant role in enabling this simulation: SciPy, Numpy, Matlotlib, Scikit-Optimize, Shapely, and Pathos [99, 101–104]. I would also like to thank the open source community at large for their efforts that are enabling better scientific research. Finally, I would like to thank Dr. Michael Ortner of Silicon Austria Labs (SAL) for his help with the magnetostatic method of moments and insights into modeling permanent magnets.

Chapter 6

Simulation Results

6.1 Introduction

This chapter describes an optimized storage ring design. Various aspects of the design are characterized by simulation including phase space behavior, expected flux, feasibility of evaporative cooling, and use with other species. Loss mechanisms are discussed, including practical considerations such as material and alignment errors. The design described achieves about 30% of the maximum flux possible considering vacuum limited lifetime and optical pumping efficiencies.

6.2 Overview

An overview of the lattice is shown in Figure 6.1. Magnetic elements are labeled and their details are given in Appendix C. The round-trip length of the design orbit of the ring is about 8.04 m and the design speed is 210 m/s. Particles originate from a 1 cm diameter aperture surrounding the existing atomic focus. Their phase space distribution is a model of the focus and is described in the following sections. To reduce Majorana spin-flips a solenoidal winding is present in every magnetic element to prevent a zero field at the center of the element. The field generated is around 5 gauss and has a negligible effect on simulation results.

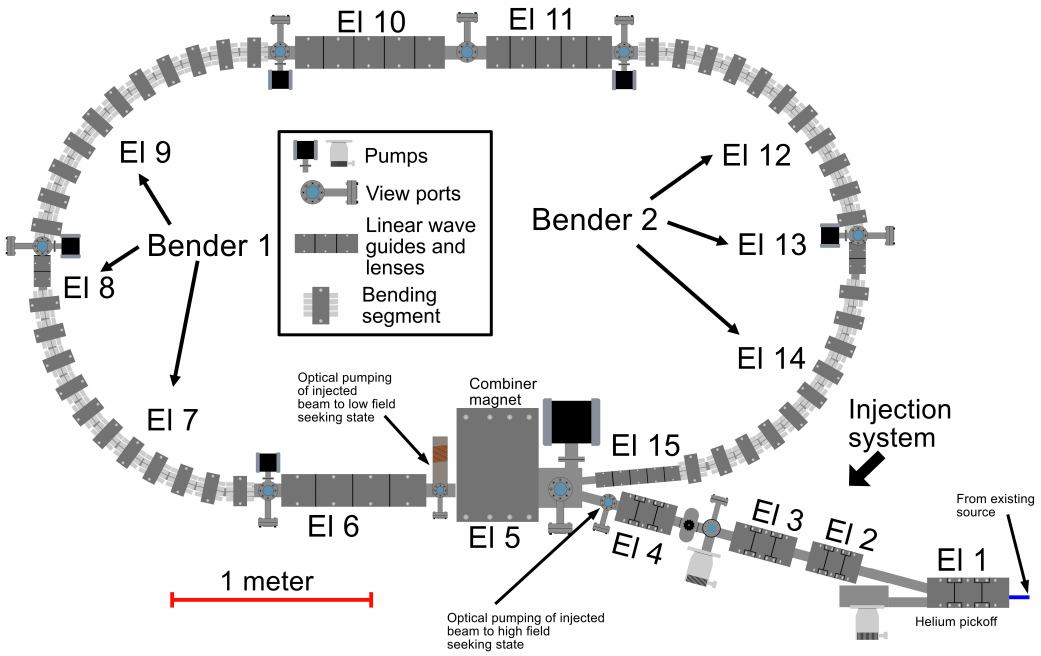

Figure 6.1: Graphical overview of the injector/storage ring system with vacuum system. The view is from above directed vertically downward. Particles originate from the source and travel through the injection system into the ring where they begin to circulate and build up a large flux. Dimensions are approximate. "El N" is the element label for later reference where "N" is a number assigned to each element.

Atoms from the atomic focus enter a lens (El 1) with a horizontally transverse offset of 11.4 mm. The offset shifts the lithium beam out of the helium beam, and the helium beam continues to a dump chamber where it is pumped away. The lithium beam then travels through a series of 3 lenses (El 2, 3, and 4). These lenses allow for small adjustments to the injected beam to tune the mode matching into the ring as well as allowing space for differential pumping. The pressure drop from the existing source to the ring needs to be about 3 orders of magnitude for adequate lifetime. In this region He diffusing from the existing chamber will dominate the pressure.

As atoms leave the last lens of the injector (El 4) they are optically pumped from a low field seeking state to a high field seeking state (). They then enter the combiner magnet (El 5) and are deflected in the opposite direction of the circulating beam. Both beams join at the output colinear with each other and the injected beam is pumped into the low field seeking stretched state $2^2S_{1/2}\,|F=2, m_F=2\rangle$. A large gap is present between the output of the combiner and the next element so that fringe fields will have fallen to low enough values to allow high field optical pumping. Details of the optical pumping setup are described in Appendix A.

The atoms then begin circulating in the ring periodically traveling through elements 5-15. Elements 7-9 and 12-14 form bender 1 and 2 respectively. Each bender is split at its apex to allow for vacuum pump access. Between each split is a short lens (El 8 and 13) to couple the atoms between the half benders. Elements 10 and 11 transfer atoms from one bending section to another with space for pumping and observation. Elements 15 and 6 mode match atoms into and out of the combiner

respectively.

6.3 Flux

In Section 5.5 the flux multiplication, F, was introduced as a possible figure of merit. It describes how many times the input flux (atoms/second) is amplified within the ring. It is possible that a storage ring would support an infinite number of revolutions for some regions of initial phase space, which would result in $F = \infty$ without any additional losses. In reality there are several loss mechanisms, including imperfect optical pumping, vacuum, material imperfection, misalignments, and Majorana spin-flip losses. The addition of a solenoidal coil to the vacuum system is expected to nearly eliminate Majorana spin-flips so we will not consider this limit.

At the output of the combiner both the injected and circulating beam pass through the optical pumping region. For an atom in the circulating beam there is some probability of being inadvertently pumped into the untrapped high field seeking state and being lost. The optical pumping efficiency, ϵ, is defined to be the fraction of particles remaining in the trapped low field state. The optical pumping efficiency for the circulating beam is estimated to be 99.96% for a single pass through the optical pumping region. This sets an upper limit on F of about 2500 from Equation 6.6. This depends strongly on the exact value of efficiency as shown in Table 6.2. However, this is the value for the low atom density limit neglecting radiation trapping effects. Including radiation trapping in the model results in an upper limit of about 1600 times. When other loss mechanisms are included, the differences in efficiency gives results that are negligibly different, though this

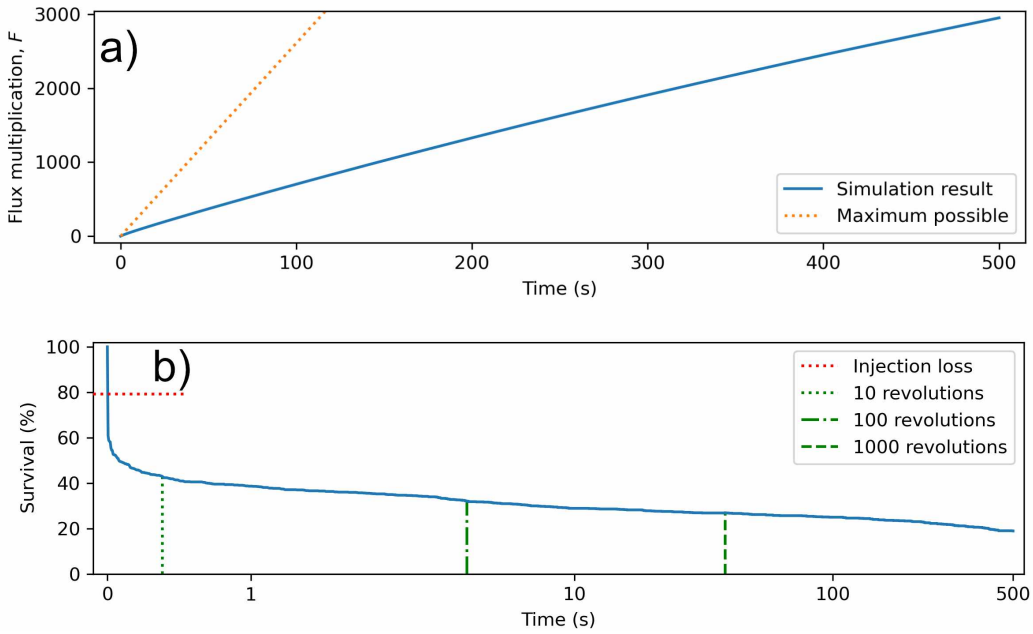

Figure 6.2: Results of simulating injection and circulation in the storage ring where the only loss mechanism is particles colliding with vacuum walls. (a) Flux multiplication versus simulation time. (b) Percent survival of particles versus simulation time. About 20% of particles are lost through the injection system, and another 35% are lost in the next 10 revolutions. About 20% of particles are very long lived.

assumption begins breaking down for $F > 1000$. Another limit is set by vacuum lifetime. Appendix B describes a possible vacuum system that would achieve a lifetime of about 50 seconds. From Equation 6.4 this results in a vacuum limited F of about 1300, assuming only vacuum losses. Perhaps the most concerning is the limit from the introduction of magnet errors described in Section 2.6.2. This has the possibility of reducing F to less than 50. In Section 6.7.2 a solution is presented that largely alleviates this concern.

Figure 6.2 shows the result of a simulation where the only loss mechanism

is particles colliding with vacuum walls. 1000 particles were injected from the focus and allowed to circulate for up to 500 seconds. About 20% of particles are long lived and survive over 13,000 revolutions. Combining the simulated results with the vacuum lifetime and optical pumping efficiency using Equation 6.7 gives a flux multiplication of about 240 times. Table 6.1 presents the flux multiplication for different values of vacuum lifetime and optical pumping efficiency for the simulated results as well as the zero vacuum wall collision loss case. The proposed design is performing at about 30% of the optimal value given realistic vacuum and optical pumping constraints. The maximum flux of the source is about $2{\times}10^{12}$ 1/s at the focus and with the anticipated flux multiplication this results in a flux of 4.8×10^{14} atoms/s circulating in the storage ring. This corresponds to about 1.8×10^{13} total atoms circulating in the ring. As discussed in Section 3.3, we expect the flux at the focus to increase up to 10 fold if the chamber pressure is reduced.

The simulation was performed with a 1 μs time step and a dense field interpolation requiring 30 GB of memory. A particle which survives to 500 seconds will perform around 500 million time steps. This raises the possibility of small errors compounding and producing unrealistic results, but less refined models were not significantly different. In the end, the weighted impact of particles surviving past 100 seconds is largely negligible for realistic vacuum and optical pumping limits.

6.3.1 Initial Phase Space Distribution

The phase space distribution of the source of particles used for optimization comes from the combined jet seeding particle tracing simulation used to model the

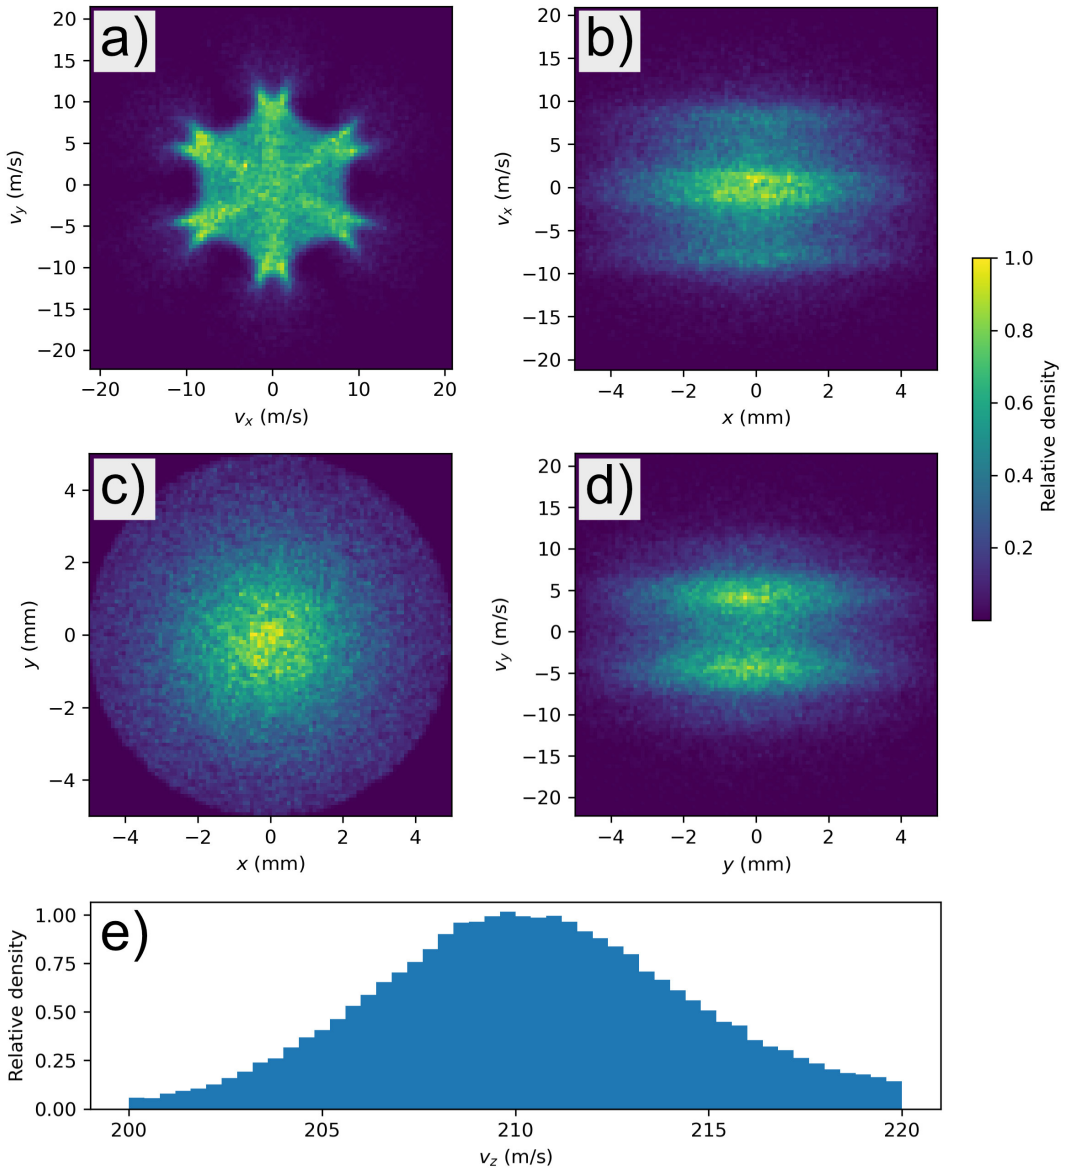

Figure 6.3: Initial phase space distribution used for modeling and optimizing the storage ring. The distribution comes from the combined jet seeding particle tracing simulation with 3 mK temperature floor. The symmetry of the magnet is clearly seen in (a) and somewhat visible in (b) and (d). In (e) the long high velocity tail of the distribution is apparent.

τ (s) \ ϵ (%)	100	99.99	99.96	99.9	99
∞	? (∞)	1700[a] (10000)	650 (2500)	280 (1000)	35 (99)
200	1200[a] (5200)	850 (3400)	460 (1700)	240 (840)	34 (98)
100	670 (2600)	550 (2100)	350 (1300)	210 (720)	34 (96)
50	360 (1300)	320 (1200)	240 (860)	170 (570)	33 (92)
30	220 (780)	210 (730)	180 (600)	130 (440)	32 (88)
10	83 (260)	81 (250)	76 (240)	67 (210)	26 (72)

Table 6.1: Simulated flux multiplication under the influence of optical pumping efficiency, ϵ, and vacuum lifetime, τ, losses. Values in parenthesis are the joint optical pumping/vacuum limited only case and represent the upper limit. The low atom density limit of optical pumping efficiency is assumed. The yellow cell is the anticipated value. [a]These values may be dubious because they depend more strongly on very long lived particles and the low atom density approximation for optical pumping is violated.

focus described in Section 3.2.2. The 3 mK temperature floor model is used instead of the adiabatic condition because of its close fit to observed data. The initial phase space distribution of this model is shown in Figure 6.3. Notice that the longitudinal velocity is artificially truncated. This is because the distribution has a very long tail for higher velocities while only a small portion near the peak velocity arrive to the atomic focus and can enter the ring. A range of 10 m/s on either side of the peak velocity is used.

6.4 Stability

An important characteristic of a storage ring is its stability. In our linearized theory derived in Chapter 4, stability is only dependent on the longitudinal velocity of a particle. In reality the presence of nonlinear forces introduce a dependence on other initial phase space coordinates. This dependence is stronger for particles

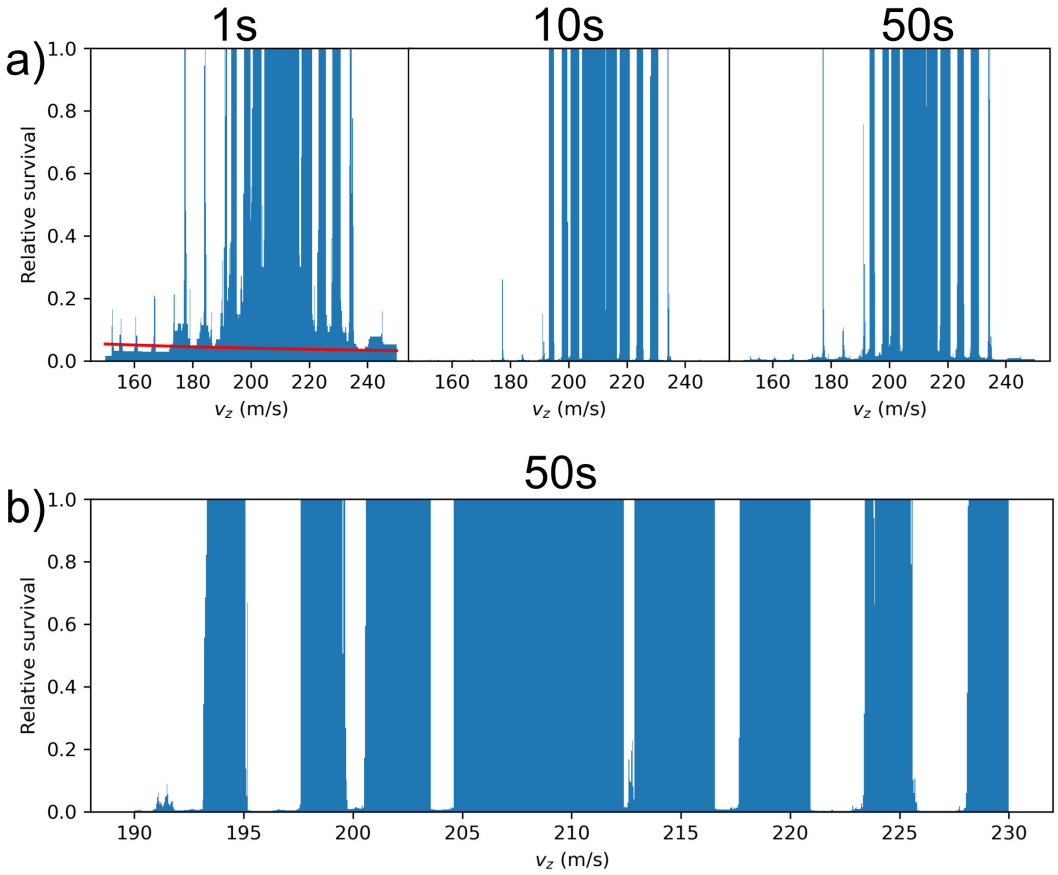

Figure 6.4: Relative survival of particles versus initial longitudinal velocity. A value of 1 indicates the particle surviving to the simulation time. (a) Stability behavior over a range of 150-250 m/s for 3 simulation times. For reference, the red line in the first figure is the survival value of one revolution around the ring. (b) Stability behavior over 190-230 m/s with smaller velocity spacing for a simulation time of 50 seconds.

that more strongly violate the linear assumptions which occurs for trajectories in phase space that are further from the ideal. For particles that exist near the ideal trajectory in phase space, the linear stability assumption is more accurate. If cooling is applied particles will move towards the ideal phase space trajectory which make its stability of central importance. Figure 6.4 depicts the stability characteristics of the ring. Particles are launched from within the ring with zero transverse momentum and position, and allowed to circulate for 50 seconds. As we saw in Section 4.12.2, deadbands separate regions of stability.

The drop off of particle survival at higher velocities occurs because magnetic forces cannot balance the centrifugal pseudo forces. Equivalently, their value of δ is so large that their dispersive orbit carries them into a collision with an aperture. The reason for the drop off at lower velocities is less clear. One explanation is that particles are lost after passing through the combiner because they are more strongly deflected than faster particles and fail to be aimed into the following element. However, nearly every particle that is lost before making one full revolution survives well past the combiner. Another explanation is simply that the oscillations that the slower particles experience from the imbalance of the centrifugal pseudo force and magnetic force build in amplitude for each pass through a bender.

Figure 6.5 shows the simulated stability characteristics versus the theoretical predictions. Notice that the number of deadbands increase and regions of stability shrink at lower velocities as we already saw in Section 4. The general sparsity of stable regions is a possible explanation for the lack of survival at lower velocities. Around the central velocity region there is reasonable agreement between predicted

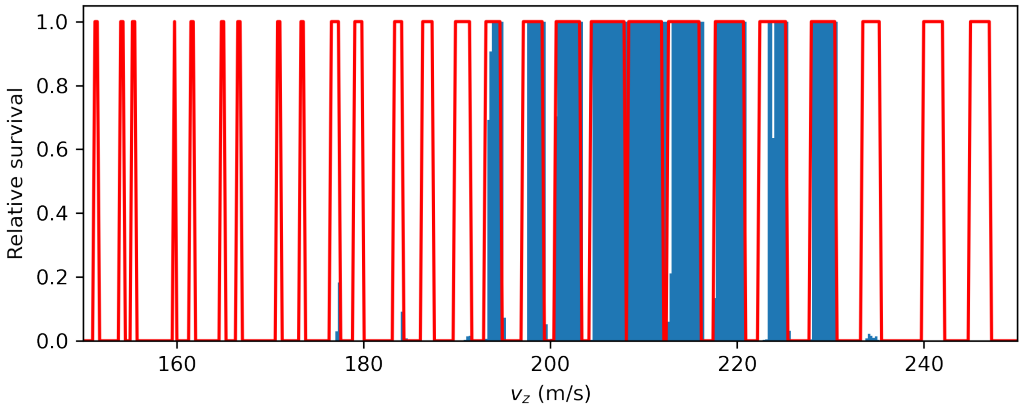

Figure 6.5: Relative survival of simulated particles and predicted survival versus longitudinal velocity. A value of 1 indicates the particle surviving to the simulation time. Particles have zero initial transverse momentum and position displacement.

and simulated survival.

6.5 Phase Space behavior

The trajectories of particles through phase space can shed light on a dynamical system. Figure 6.6 shows Poincaré maps at the output of the combiner for an ellipsoidal swarm of 12,536 particles. The swarm was launched from within the ring at the combiner output. No particle collided with vacuum apertures. The statistical definition of emittance, Equation 4.83, is given for each figure. Recall that the emittance is proportional to phase space area.

The ellipse begins to deform after a few tens of revolutions and starts to fold in on itself. Eventually the ellipse reaches an equilibrium with an elliptical interior and exterior. The emittance increases by about 60% and nearly reaches its final

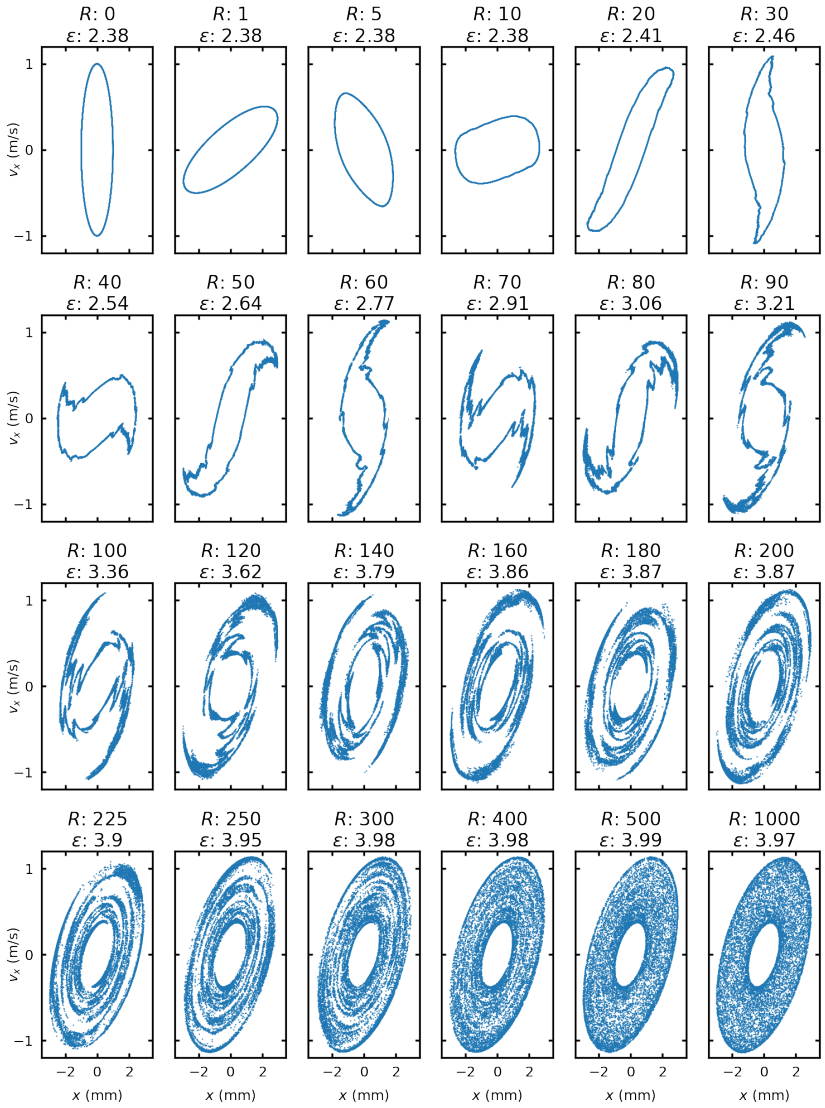

Figure 6.6: Poincaré maps for a range of revolutions in the storage ring. R is the number of revolutions and ϵ is the emittance in units of mm·mrad/s.

value by 160 revolutions. Interestingly, the phase space distribution continues to evolve and mix substantially after this point. This is a well known phenomena in accelerator physics called phase space filamentation. It is possible for a lattice to preserve transverse phase space area yet distort the phase space distribution into an undesired shape. Typically this occurs by the distribution revolving around itself with "filaments" exactly as seen here.

In a linear system the evolution through phase space is an affine transformation such as a rotation, translation, scaling etc. Clearly the transformation from $R = 0 \rightarrow R = 1000$ is not affine, while $R = 0 \rightarrow R = 20$ is approximately affine. In analogy with optics one could prescribe a phase space coherence length of approximately 10 revolutions which corresponds to about 400 ms.

We can also analyze the phase space acceptance of the storage ring/injector system to gain insights on what distributions can be used. To do this a swarm of 1 million particles is initialized in phase space, launched from the focus into the injector, and allowed to circulate for 10 seconds. Results are shown in Figure 6.7. The effects of longitudinal instability are clearly seen in Figure 6.7(e), but there is less contrast between stable and unstable regions than in Figure 6.4. The asymmetries of the injection system can be seen in 6.7(a-d). Unsurprisingly, there is a large amount of overlap between Figures 6.7 and 6.3.

6.6 Different Species

In Section 4.13 we saw that an ideal storage ring could accept different species as long as they had the same species factor $\xi = 2\mu/mv^2 = \mu/E_k$. For reference, the

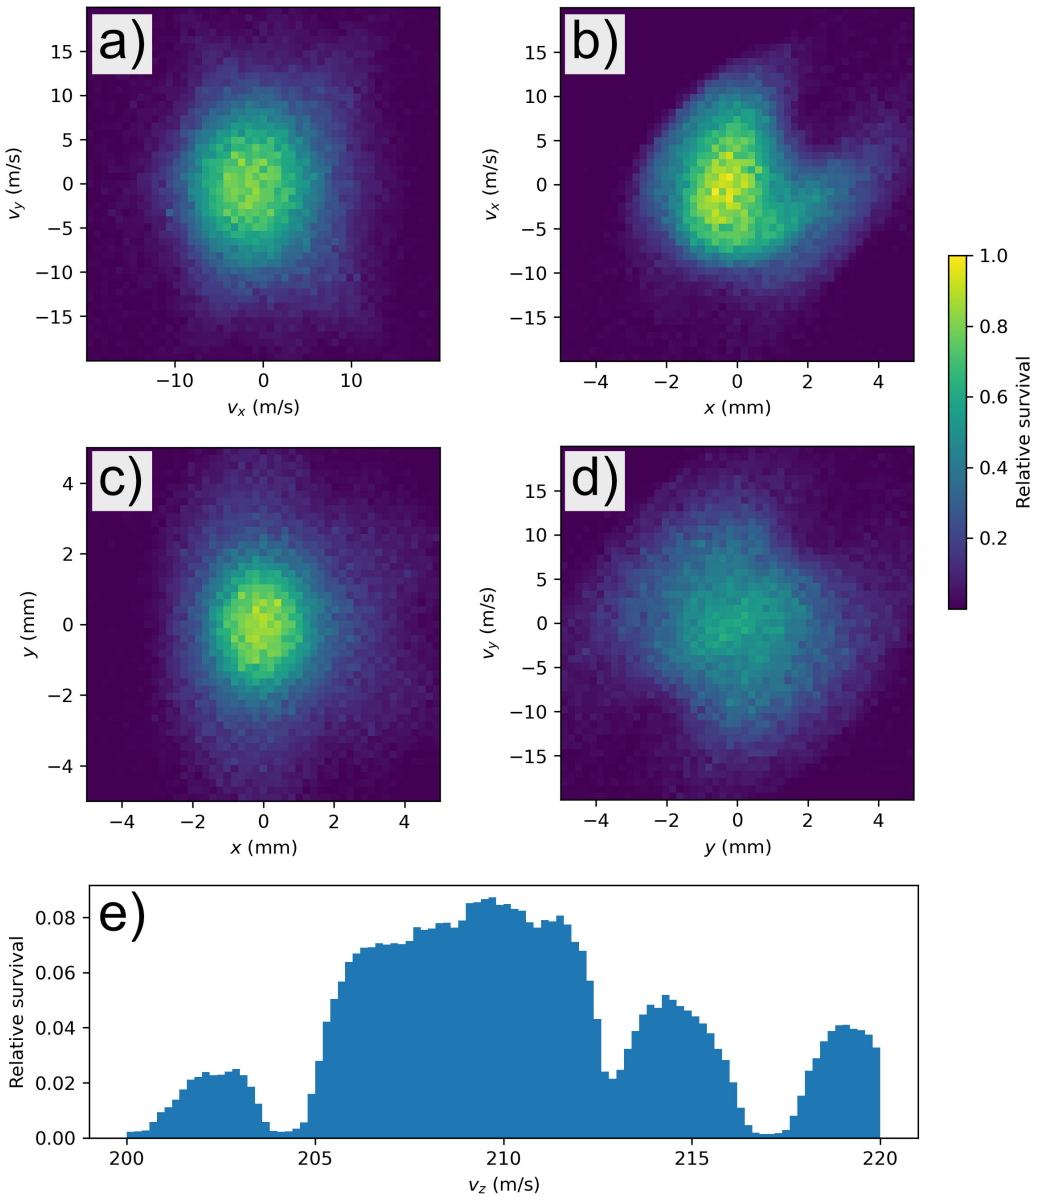

Figure 6.7: Acceptance of particles versus initial phase space location for a 6 dimensional distribution. A relative survival of 1 indicates that the particles survived to the simulation cutoff time. All figures are histograms, so the relative survival is the average in the pin/pixel.

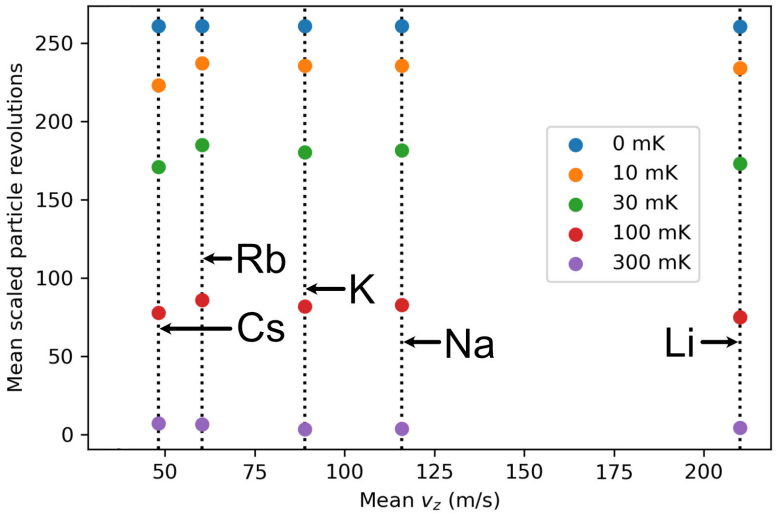

Figure 6.8: Mean particle revolutions versus velocity for the stable alkali atoms. Various temperatures are shown. Revolutions are scaled by $210/v_z$ for comparison. The maximum revolutions possible is 261.1 for Li.

species factor of our source is $\xi_0 \approx .036$ T^{-1} at 210 m/s. We also found if a thermal distribution is used as the source of loading then as long as the magnetic moment and temperature are identical, any species can be used.

The simulated results agreed for a lattice using only ideal elements with no combiner or injection system. To test whether the agreement still holds, a similar simulation was carried out. 100 particles were injected from a thermal point source at the focus and traced through the ring for up to 10 seconds. Results are shown in Figure 6.8. Once again there is good agreement over a wide range of masses.

Our source's nozzle produces a He beam with forward velocity of 210 m/s, but this can be intentionally increased using the nozzle heater. The species factor can be rearranged to give $\mu/m = v^2\xi/2$. Because v can be increased, any species

can be used with this design as long as it satisfies

$$\frac{\mu}{m} \geq \frac{1}{2} v_{min}^2 \xi_0 \approx 796, \tag{6.1}$$

where v_{min} is the minimum forward velocity of the jet which is currently 210 m/s. For alkali atoms this is only satisfied with the stable isotopes ^{6}Li and ^{7}Li.

In principle the forward velocity can be reduced with further cooling. Temperatures of about 1.5 K can be achieved using a 2 stage PTR, such as the one used in this experiment, and evaporative cooling of ^{4}He [105, 106]. From Equation 2.6 the new jet forward velocity would be $210\sqrt{1.5/4.5} \approx 121$ m/s. The next lightest alkali atom is sodium which from Equation 6.1 requires $v \approx 116$ m/s. It may be possible to use sodium, but a more careful analysis is required. Other alkali atoms could be used, but this would require temperatures below 1K where it is challenging to maintain the necessary high cooling power [107].

Another approach for applicable species such as alkali atoms would be the use of laser slowing which results in velocities of tens of m/s [108–110]. This could be done continuously if an offset lens is used to shift the slowed atoms out of the laser and into the injection system.

6.7 Loss Mechanisms

In this section loss mechanisms are investigated such as collisions with background gas, optical pumping inefficiencies, misalignments, and material imperfections. Li-Li collision induced losses are discussed in Section 6.8.

6.7.1 Vacuum and Optical Pumping

The vacuum lifetime sets an upper limit on the possible flux multiplication. Without losses the flux multiplication can be written as

$$F = \int dF = \frac{1}{N}\sum_{i}^{N}\int_{0}^{R_i} dR = \frac{1}{N}\sum_{i}^{N}\int_{0}^{T_0}(v_i/L_t)dt, \tag{6.2}$$

where Equation 5.16 was used and L_t is the total length of the lattice. Including vacuum losses yields

$$\begin{aligned} F &= \frac{1}{N}\sum_{i}^{N}\int_{0}^{R_i} e^{-t/\tau}dR = \frac{1}{N}\sum_{i}^{N}\int_{0}^{R_i} e^{-RL_t/\tau v_i}dR \\ F &= \frac{1}{N}\sum_{i}^{N}\frac{\tau v_i}{L_t}\left(1 - e^{-R_i L_t/\tau v_i}\right), \end{aligned} \tag{6.3}$$

where τ is the vacuum lifetime. Assuming each particle has approximately the same velocity v and reaches the same number of revolutions R_0, the limiting case of $R_0 \to \infty$ yields

$$F = \frac{\tau v}{L_t}, \tag{6.4}$$

which is the vacuum limited case.

The optical pumping sets another limit on the flux multiplication. We can characterize the optical pumping efficiency, ϵ, as the probability of a circulating particle remaining in the trapped low field state. After R revolutions this probability is ϵ^R. Inserting this into Equation 6.2 yields

$$F = \frac{1}{N}\sum_{i}^{N}\int_{0}^{R_i}\epsilon^R dR = \frac{1}{N}\sum_{i}^{N}\frac{\epsilon^{R_i}-1}{\log\epsilon}. \tag{6.5}$$

ϵ	95%	98%	99%	99.5%	99.9%	99.99%
F	19.5	49.5	99.5	199	999	9,999

Table 6.2: Flux multiplication versus optical pumping efficiency assuming only optical pumping losses.

Again assuming $R_i = R_0$ and taking $R_0 \to \infty$ gives

$$F = \frac{1}{|\log \epsilon|}, \tag{6.6}$$

which is the optical pumping limited case. Table 6.2 shows F for several values of ϵ. There is a strong dependence of the exact value of F on ϵ.

We can combine both losses to get

$$F = \frac{1}{N} \sum_{i}^{N} \omega_i \frac{1 - \epsilon^{R_i} e^{-R_i/\omega_i}}{1 - \omega_i \log \epsilon}, \tag{6.7}$$

where $\omega_i = v_i \tau / L_t$ is the number of revolutions made in one vacuum lifetime. Once again assuming $R_i = R_0$ and taking $R_0 \to \infty$ gives

$$F = \frac{\omega}{1 - \omega \log \epsilon}, \tag{6.8}$$

where $\omega = v\tau / L_t$.

6.7.2 Permanent Magnet Imperfections

As mentioned, magnet imperfections have the possibility of severely constraining the performance of the design. In this section we will analyze this issue and discuss a few methods to circumvent it. Modeling these effects is computationally

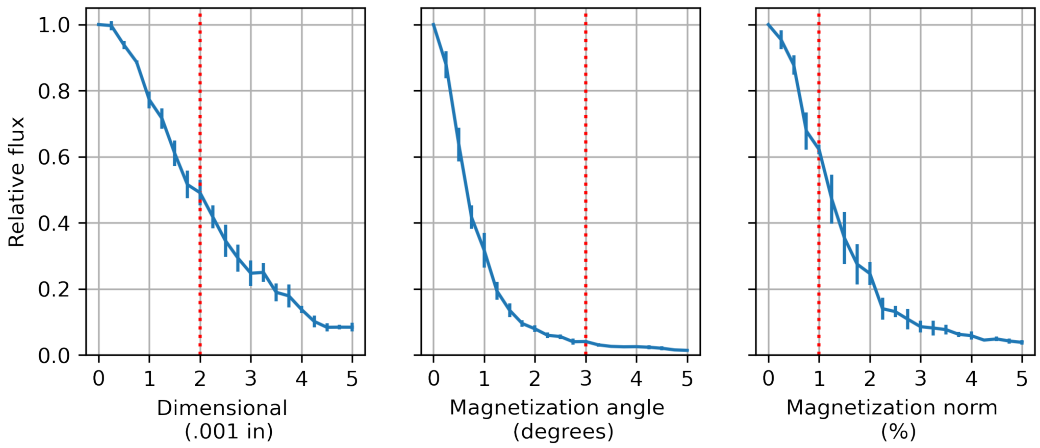

Figure 6.9: Flux multiplication versus tolerance for permanent magnet errors. There are 4 runs at each model and error bars are the RMS value. While one tolerance value was swept, the others were held at zero. The dashed red lines specify the tolerances in Table 2.1.

expensive so to ease analysis, a less refined model is used. The flux multiplication of this model without imperfections is used as a baseline to which the results of models with magnet imperfections are compared. With no magnet imperfections, infinite vacuum lifetime, and perfect optical pumping efficiency, the model results in a flux multiplication of about 480. The model of the focus is used with 1000 particles. Very long permanent magnets are impractical because they are fragile and handling them is difficult. In this model magnets have a length to width ratio of two, or the nearest viable value. In a bender the length of each magnet is the length of the individual segment.

Randomly applying the tolerances specified in Table 2.1 for 30 different models results in a flux multiplication of 18 ± 3. This is a reduction of about 95%

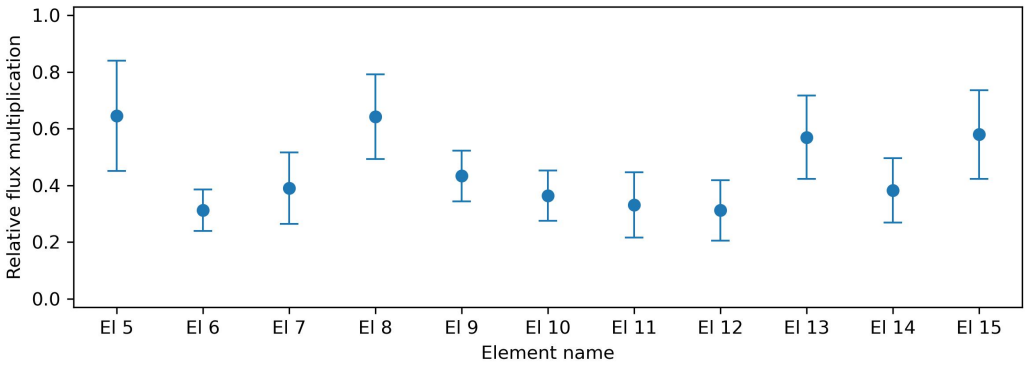

Figure 6.10: Flux multiplication from errors applied to one element at a time. There are 10 runs per element and errorbars are the RMS value.

from the baseline. The dependence of flux on the tolerance is shown in Figure 6.9. The biggest impact is clearly the magnetization angle. The results shown are for a design with slightly different dimension, but the difference is minor.

The effects of errors applied to one element at a time are shown in Figure 6.10. Interestingly, El 5 and 6 (the combiner and the lens after the combiner) are both large hexapole lenses of similar length yet have very different effects on results. What is happening is that the beam is traveling through the edge of the combiner while it is traveling through the center of the following lens. As explained in Section 3.4.4, the impact of magnet imperfections is relatively stronger near the center of the lens.

It would be hard to justify building the system with such a low level of performance. However, there are a few possible solutions to the problem. The first is to simply purchase magnets with a magnetization angle tolerance of less than 1

degree. From speaking with the supplier Dexter Magnetics, the total magnet cost, which is already on the order of 10,000 US dollars, would increase by 5-10 times. However, it may be possible to find lower cost vendors.

Another solution is inspired by a similar problem that occurs with ferromagnets in accelerators [92]. The characteristics of steel vary from manufacturer to manufacturer and even batch to batch which could result in the performance of electromagnets varying by unacceptable amounts. Rather than requiring challenging and expensive tolerances to ensure consistency, the ferromagnets are assembled out of thin sheets of steel that have been randomly selected from the entire stock. This results in the deviations between batches and manufacturers averaging out. For permanent magnets this would entail using many thin magnets. A preliminary investigation of this approach was not promising because tens of thousands of magnets would be required, which is impractical.

Yet another solution from accelerator physics is the use of shimming which is the addition or subtraction of iron to shape the magnetic field [92]. This is typically done to extend the usable field region and improve fringe fields. The method works because magnetic field lines are mostly perpendicular to interfaces between materials when the difference in permeability is large between the materials. This condition is met at the interface of air and unsaturated steel which allows field lines to be "aimed".

One version of this approach is to apply shimming to individual magnets. Figure 6.11 depicts this with a 1x1 cm N52 grade neodymium cuboid. The simple addition of a 1 mm shim to each end redirects the field towards the intended di-

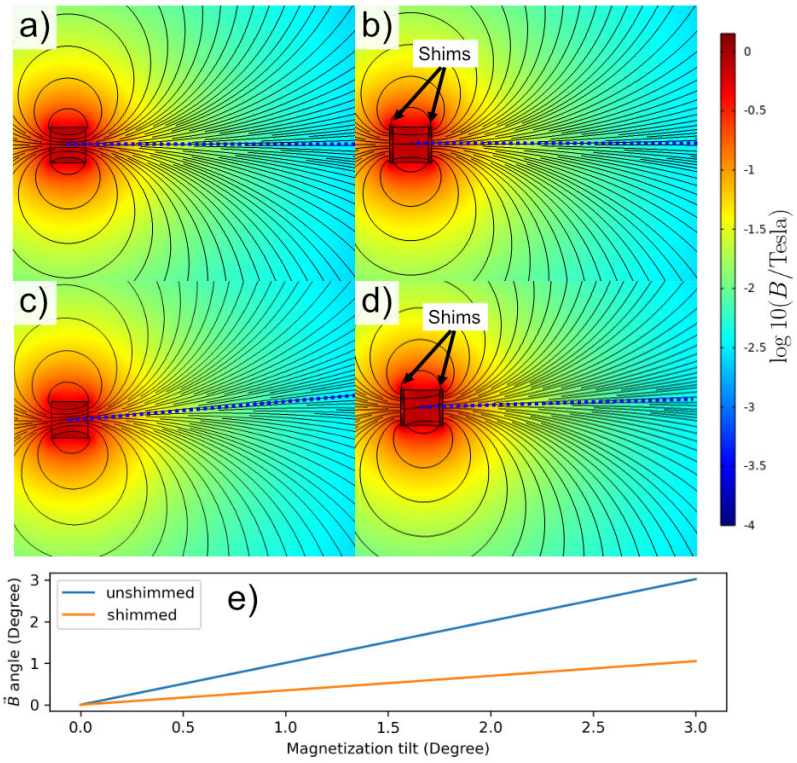

Figure 6.11: Results of shimming a 1x1 cm infinitely long permanent magnet with a 1 mm steel shim. Figures (a) and (c) show the magnet with the ideal magnetization and a tilt of 5° respectively. Figures (b) and (d) show the same magnets with the shim added to each end. The magnetic field lines are clearly straightened with the addition of the shim. Figure (e) plots the magnetic field angle as a function of magnetization angle tilt for a point located 1 cm from the magnet face for the shimmed and unshimmed magnet.

rection. This was simulated with COMSOL in 2D and assumes material extends infinitely in and out of plane. Shimming could also be used in the bore of the element. This is somewhat similar to an accelerator magnet design known as a hybrid magnet, which is effectively a ferromagnet where the coils have been replaced by permanent magnets [111, 112].

A more clever and possibly simpler solution is to exploit a detail of the manufacturing process. A basic understanding of the manufacturing process of neodymium magnets will make this clear [113]. Raw materials are combined together in a furnace to react and form chunks of NdFeB. The chunks are reduced to a fine powder then pressed into a block under a magnetic field to align the powder particles. At this point the block is not appreciably magnetized. The block is cut into the desired shapes, often by an Electrical Discharge Machine (EDM) and then coated. Finally, the shapes are magnetized. The details of the process vary depending on the size and shape of the final product. Depending on the manufacturer, a lot of magnets may have been all cut from the same block and magnetized together at the same time. This is an important detail because it implies that the magnetization errors within a batch are all be nearly *identical*. The dimensional errors may be identical as well because of the nature of EDM cutting. Even if the errors vary throughout the batch, they will be very correlated for magnets that were close together during the process.

Assuming that each element is assembled from permanent magnets with the same errors and the errors vary from element to element, the flux multiplication of 30 different models is 410 ± 80. This is a 15% reduction from baseline. This

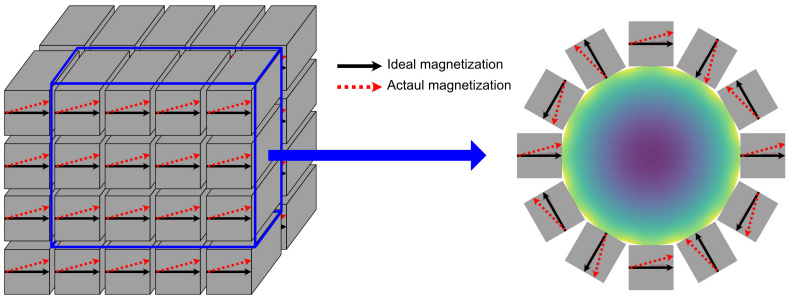

Figure 6.12: To overcome material errors, particularly magnetization direction errors, each lens must be assembled from magnets that are near each other in the batch so that the errors are correlated.

assumes that the magnets are installed such that their magnetization vectors all deviate by the same direction as shown in Figure 6.12. The results are now quite adequate and magnets with reasonable tolerances can be used. To reiterate, this only works if the magnets are from the same batch and delivered in a pattern that can be used to place the magnets correctly. Dexter Magnetics has stated that they can provide magnets which all originate from the same batch.

6.7.3 Assembly Misalignment

Another source of loss is assembly tolerances. One form of assembly misalignment is the misplacement of individual elements. To understand the impact of this, each element was randomly perturbed in 3D up to some maximum deviation. The element is randomly translated and rotated such that the maximum separation from the ideal position is up to the maximum deviation. The same model mentioned in the previous section was used. Figure 6.13 shows the relationship between flux multiplication and this tolerance. Benders are neglected in this analysis.

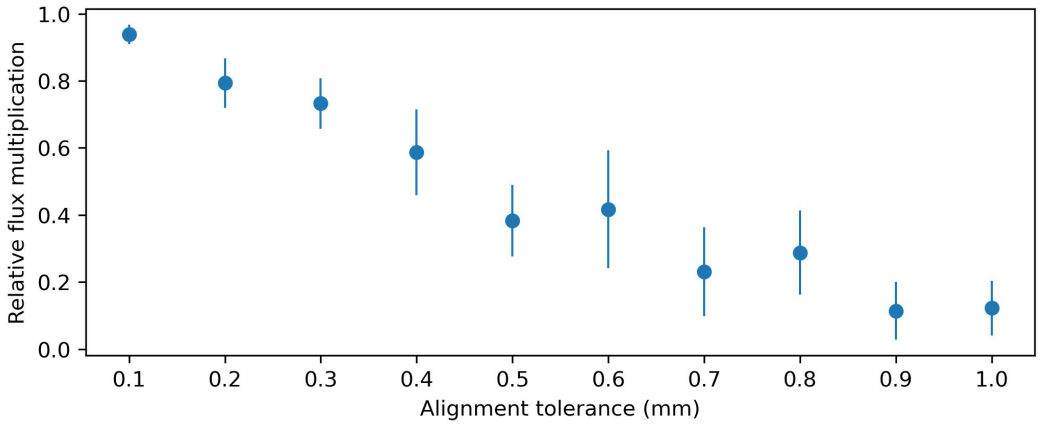

Figure 6.13: Impact of assembly alignment tolerances on performance. At each tolerance value the system was modeled 10 times. Error bars are the RMS value.

A fairly tight tolerance is required to achieve nearly optimal results. Particle accelerators have the same problem and a variety of solutions exist [92,93,114]. One technique is to first assemble the system as precisely as reasonably possible, then perturb each element until the desired performance is achieved. This can be done systematically by placing a beam monitor after the first element and making adjustments until the desired alignment is achieved. The beam monitor is then placed after the next element and the process is repeated. It is likely that something like this will be required for the storage ring. Possibly a third step involving perturbing each element to maximise total flux will be required. Another form of misalignment is the placement of individual magnet cuboids in their mounts. This is approximately the same problem as the material errors, and thus it is essential that the angular deviations remain below 1 degree.

Particles could be lost during injection if the focus moves from the design

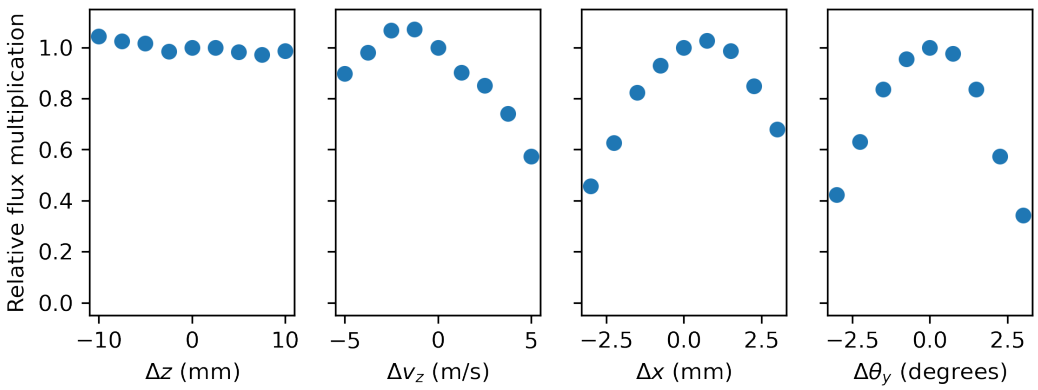

Figure 6.14: Relative flux multiplication versus deviation of the focus from the ideal case. Orbit coordinates are specified.

location or otherwise changes. Based on our experience longitudinal position and velocity variations of a few millimeters and m/s should be expected. Improvements to the nozzle temperature control would likely improve this. Transverse behavior has been observed to be much more consistent with position and angle variations of less than a millimeter and degree. Figure 6.14 shows the dependence of survival on deviations of the focus from ideal. 1000 particles were injected from the focus and simulated for 10 seconds. Over a reasonable range of values the design still performs well. At the focus a 1 cm aperture will be present and is not included in this analysis.

6.8 Cooling

One application of a neutral atom storage ring is to produce a degenerate or near degenerate gas with a nonzero momentum vector. This would likely require the

use of evaporative cooling. The temperature of our focus is around 10 mK, while Bose-Einstein Condensation (BEC) in alkali atoms has been achieved in the range of tens of nanoKelvin to a few microKelvin depending on the density [115–119]. The design must be resistant to intrabeam Li-Li collisions across this temperature range or evaporative cooling will not be possible.

This stability can be tested with a simple collisional model similar to the seeding simulation. A constant collision rate Γ and beam temperature T is assumed everywhere in the ring. The beam is assumed to be traveling at 210 m/s. At each time step the probability of collisions is $p = \Gamma h$, where h is the time step and $p \ll 1$. A collision occurs by first sampling a 3D velocity vector from a thermal distribution of temperature T with the appropriate mean velocity. This velocity vector is then used as a collision partner with the circulating atom. A random scattering direction is chosen in the center of mass frame. Typically a few hundred collisions per trap lifetime are required for evaporative cooling [120]. For our anticipated vacuum lifetime of 50 seconds, this corresponds to $\Gamma \approx 5$ s^{-1}.

The collision rate in the storage ring can be estimated with the Li-Li cross section, relative Li-Li velocity, and mean density. The mean density in the ring can be determined from the simulated density profile which is given in Figure 6.19. The predicted profile from the Twiss parameters from Equation 4.84 is given as well and shows good agreement. From this plot the mean density is $n \approx 6 \times 10^{16}$ m^{-3}. Assuming a lithium temperature of about 10 mK, the relative velocity is approximately $v_{rel} = \sqrt{16KT/\pi m} \approx 7.8$ m/s. This neglects relative velocity originating from beam dynamics. Unfortunately, data on the ^{7}Li-^{7}Li cross section

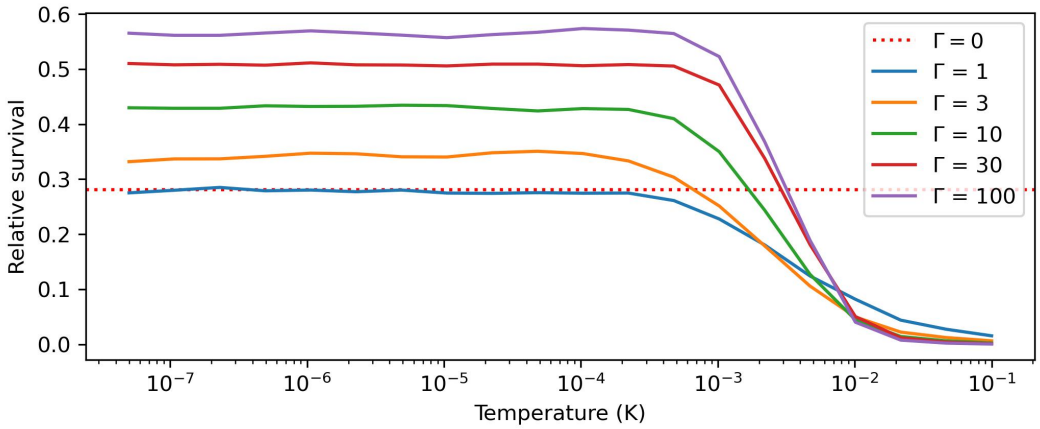

Figure 6.15: Relative survival versus temperature for several collision rates. Relative survival is the mean particle revolutions of the swarm divided by the maximum possible.

is not available at 10 mK. The best we can do is use the S-wave scattering cross section which is $\sigma_{Li-Li} = 8\pi a^2 \approx 6 \times 10^{-17}$ m^2 [121]. The collision rate is then $\Gamma = v_{rel} n \sigma_{Li-Li} \approx 30$ s^{-1}.

Figure 6.15 shows the survival of 1000 particles originating from the simulated focus for a range of Γ and T values. The particles were traced for 50 seconds. From 50 nK to about 1 mK the survival of particles is increased with greater collisionality. Around 10 mK the trend reverses.

The results indicate that there would be substantial losses and the circulating flux would be reduced. However, this is not likely to be the case for two reasons. First, the S-wave scattering cross section is likely an overestimate because it is only valid for temperatures much less than 10 mK, and cross sections typically fall with increasing temperature for alkali atoms. For example, the He-Li cross section de-

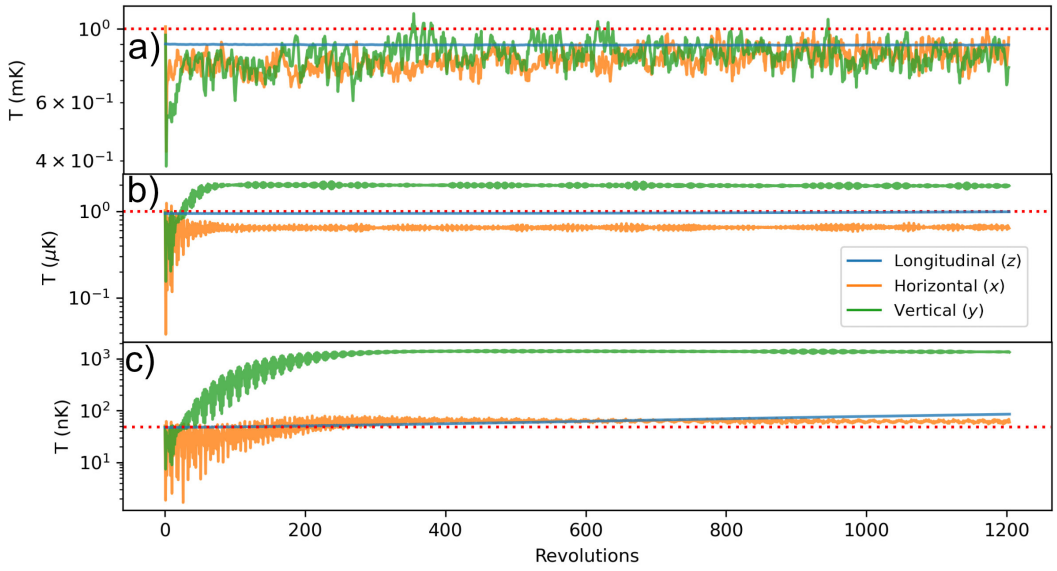

Figure 6.16: Heating of a swarm versus revolutions for 3 different initial temperatures: 1 mK, 1 μK, and 50 nK.

creases by at least 60% from the S-wave limit to 10 mK [74, 75]. Second, and most importantly, the losses result in evaporative cooling. If the temperature falls by a factor of 3, which is quite small for evaporative cooling, then losses cease. Therefore, collisional losses will probably not be a substantial loss mechanism. Further investigation is needed.

If the atoms were instead cooled before entering the ring, then the ability of the ring to maintain low temperatures is relevant. To determine what degree of heating may be occurring, a swarm of 1000 particles with temperatures of 1 mK, 1μK and 50 nK was launched from within the ring and allowed to circulate for 50 seconds. Three particles were lost to collision with the vacuum tubes for the 1 mK swarm but otherwise there were no losses. The swarm was launched from between

El 14 and 15 and the temperature was measured at that point as well. Results are shown in Figure 6.16.

For milliKelvin initial temperature the lattice introduces negligible heating. For microKelvin and nanoKelvin temperature, this is also true for the x dimension but not the y dimension. By 200 revolutions the y dimension heats up to about 2 mK. Disabling the gravitational force results in this effect disappearing. It is not clear what conclusions can be drawn from this. On one hand it is possible that the oscillations produced by gravity induce this heating. This could occur by the swarm being moved into regions with more non linear forces that shear the swarm in velocity space. On the other hand, it is possible that this is a limitation of the simulation because only small changes in velocity can result in heating. The use of a less refined model did not appreciably change results.

Assuming this temperature increase is a real effect, it can be considered a form of heating. In the y dimension the 50 nK swarm reaches equilibrium in about 20 seconds. In the first second the temperature fits reasonably well to a straight line with a slope of about 30 nK/s. The first 5 seconds fit well to a second order polynomial giving $T \approx 14.4 + 29.9t + 19.3t^2$, where t is time.

Notice the strong periodic behavior for 50 nk and the initial dip for 1 mK and 1 μK. This is because the coherence length for such low temperature point swarms is larger than one revolution. Table 6.3 gives the coherence length for the different temperature swarms from Equation 4.78, where Δv is one standard deviation of the thermal distribution. For very low temperature swarms the measured temperature is strongly affected by the characteristics of the lattice.

Temperature	1 mK	1 μK	100 nK
N_c	8.2	260	820

Table 6.3: Coherence length for different Li atom temperatures. N_c is the coherence length in terms of revolutions through the lattice.

Measuring the temperature at different points in the lattice from the swarm's origin results in different temperature measurements. From Figure 4.9 the particle's maximum transverse velocity is proportional to the Twiss parameter γ as $v_\perp \propto \sqrt{\gamma}$. From the Maxwell-Boltzman distribution we have $\sigma_v \propto \sqrt{T}$ which yields for the transverse temperature

$$T_\perp \propto \gamma. \tag{6.9}$$

The particle's measured temperature depends on the value of γ which varies through the lattice as shown in Figure 6.17. Figure 6.18 shows the excellent agreement between γ and temperature in the y direction for a $1\mu K$ swarm. The x direction fits similarly well. There are no interactions between particles so they are not in thermal equilibrium. Therefore, these measurements are not measurements of temperature strictly speaking.

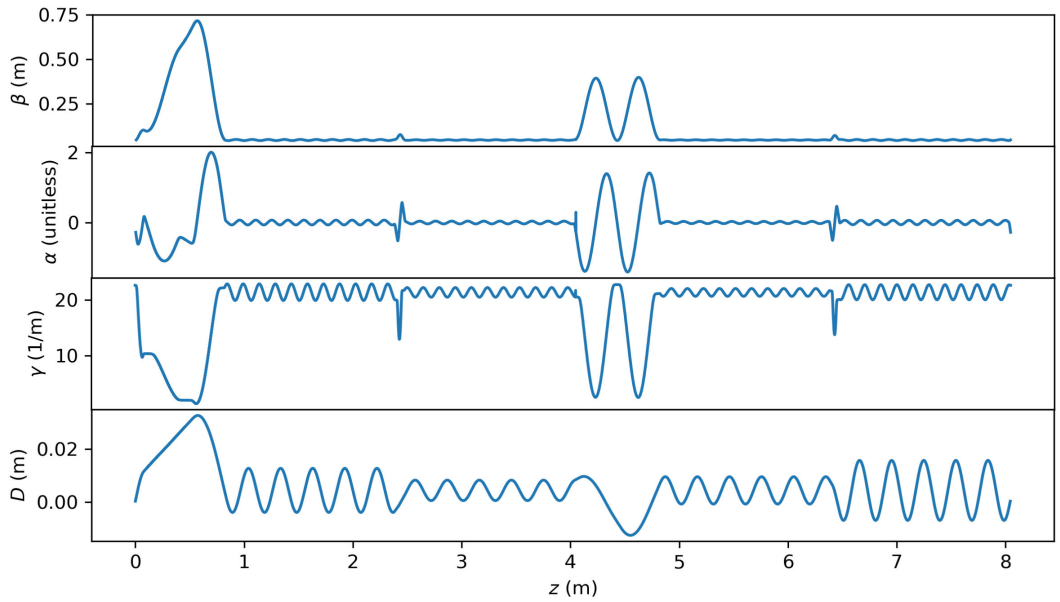

Figure 6.17: Twiss parameters (β, α, γ) and dispersion (D) plotted against z in the ring. $z = 0$ shortly after the combiner (El 5) and increases clockwise around the ring. Some numerical artifacts are present in the form of sharp peaks.

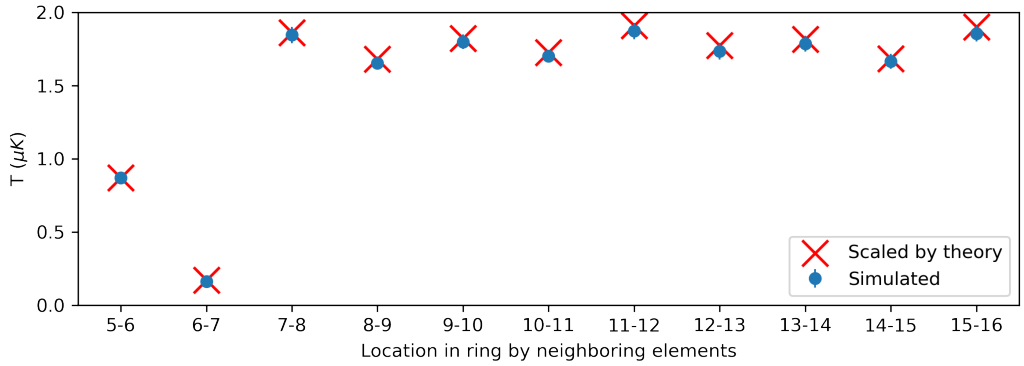

Figure 6.18: Simulated temperature in y and temperature predicted from theory by scaling γ relative to the first data point. The locations of measurement are between the elements specified on the x-axis. Magnetic fields are very weak at these locations.

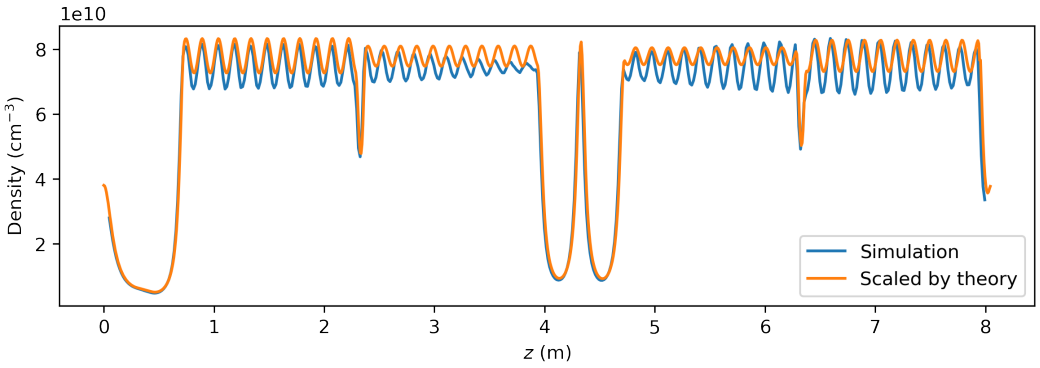

Figure 6.19: Simulated density profile and predicted density profile from Twiss parameters. The predicted profile is $\sqrt{\beta(z)}$ scaled by the peak simulated value. The density is the mean value over a circle that contains 90% of particles. A circulating particle flux of 4.8×10^{14} s^{-1} at 210 m/s is assumed.

Chapter 7

Conclusion

In this thesis the design and characterization of a meter scale permanent magnet storage ring has been presented. The improvement and characterization of a cold continuous source of ^{7}Li atoms that will be used to load the ring has been described as well. The source achieves a flux of 2×10^{12} atoms/s, a forward velocity of 210(2) m/s, a brightness of $1.8(6)\times10^{19}$ $\mathrm{m^{-2}s^{-1}sr^{-1}}$, and a longitudinal temperature of 7 mK in the moving frame. Improvements to the vacuum are expected to increase flux by an order of magnitude.

With the existing source the storage ring is anticipated to build up a flux of about 5×10^{14} atoms/s and atom number of about 2×10^{13}. Continuous loading will be achieved with optical pumping. The theory of charged particle accelerators has been extended to paramagnetic species and shown to agree well with an extensive simulation. Simulations indicate that the design is amenable to evaporative cooling. Permanent magnet imperfections are a major design concern, but a technique to circumvent these issues was described. Although the storage ring design is based on the existing source, it can be used with other atomic and molecular species as long as a combination of the velocity, magnetic moment, and mass of the species satisfy a constraint. Beside the construction of the ring, future work could consider the

introduction of an output coupler to produce an intense pulse out of the circulating flux. The use of variable mechanical apertures within the ring to perform forced evaporative cooling should be investigated as well.

Appendices

Appendix A

Optical Pumping

This section describes the optical scheme to enable the continuous injection of atoms as described in Section 4.2.2. The majority of the work described here was done by colleague Jeremy Glick. We optically pump atoms seeding into the jet in the low field limit and achieve an efficiency of 90-95% by careful control of laser polarization. However, the same efficiency in the storage ring would result in 5-10% of atoms being pumped into the high field seeking state and being lost. This would severely limit our possible flux as shown in Table 6.2. A value of at least 99.5% is needed which requires a different approach.

A viable alternative is high field optical pumping. For field above a few tens of gauss the level splitting become appreciable as shown in Figure 2.10. Transitions into unwanted states are now suppressed by both polarization selection rules and large splitting. The pumping can be applied with a laser for each transition, which would be technically challenging and costly. Alternatively, a specific field value can be chosen that minimizes the number of required lasers by introducing dark states. The goal is to minimize the transition rate for the $|2,2\rangle \rightarrow |2,1\rangle$ and $|2,2\rangle \rightarrow |1,1\rangle$ states and minimize the number of laser needed to pump into the $|2,2\rangle$ state.

The magnetic field will be produced by a ferromagnetic dipole magnet shown

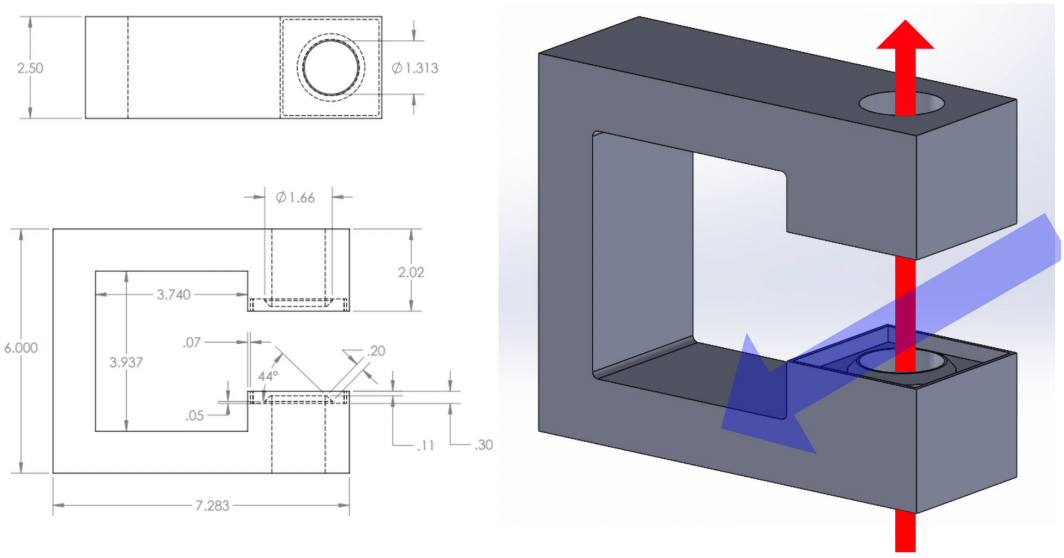

Figure A.1: Schematic of the proposed optical pumping magnet. The unusual shape of the pole face is from the addition of shims to improve the field quality. The red arrow depicts the laser and the blue depicts the beam.

Ground States (F,mF)	Excited States (F,mF)	Δf (MHz)
$\|1,-1\rangle$	$\|1,0\rangle$ $\|2,0\rangle$	−180 −16
$\|1,0\rangle$	$\|1,1\rangle$ $\|2,1\rangle$	−95 108
$\|1,1\rangle$	$\|2,2\rangle$	201
$\|2,-2\rangle$	$\|1,-1\rangle$ $\|2,-1\rangle$	97 217
$\|2,-1\rangle$	$\|1,0\rangle$ $\|2,0\rangle$	−74 89
$\|2,0\rangle$	$\|1,1\rangle$ $\|2,1\rangle$	−194 4
$\|2,1\rangle$	$\|2,2\rangle$	−64

Figure A.2: Relative splitting for $\sigma+$ transitions for $2^2S_{1/2} \rightarrow 2^2P_{1/2}$ with a 145 G external field.

in Figure A.1. The magnet pole face is shimmed to maximize the region of acceptable magnetic field [92]. The design uses 1010 steel with a driving current of 885 ampere-turns. A similar grade will work with small changes to the current . It is important that neighboring hexapole lenses be far enough away so that fringe fields have fallen to an acceptable level (Figure 5.3).

A field value of 145 G is an optimal value for σ^+ polarization. This results in the shifts for ground state to excited state transitions shown in Figure A.2. The existing near field optical pumping laser will be used to drive this scheme. Since this laser's line-width is broadened to about 20 MHz and has 803.504 MHz sidebands from an EOM, ground state transitions in Figure A.2 can be driven with the addition of a 80 MHz EOM except for the $|1,1\rangle \rightarrow |2,2\rangle$ transition. For this a singlet AOM driven at 200 MHz will be used. This results in 10 different laser frequencies, but 3 are off resonance of any transition so in effect there are 7. Accounting for the

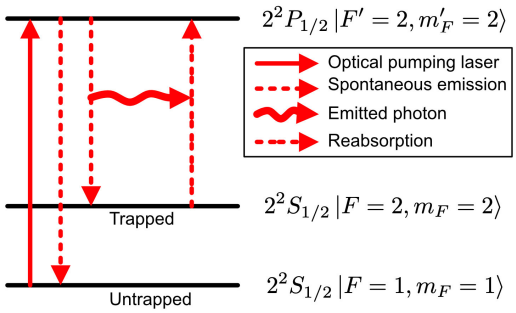

Figure A.3: Diagram of the 3 level system used to model the effects of radiation trapping in [122].

.

spectrum of side-bands and line-width of the laser the closest $\sigma-$ transition for driving an atom from the $|2,2\rangle$ ground state is about 50 MHz off resonance. The optical pumping efficiency is estimated by numerically solving the system of rate equations. There are 16 rate equations and 7 lasers in the model. The models indicate that the efficiency of optically pumping injected atoms is 99.62% and the efficiency for the circulating atoms is 99.96%. States $|2,2\rangle$ and $|2,1\rangle$ are considered trapped. These results are for a laser intensity of 1 mW/cm^2 being sent through the EOM's with a modulation depth of 1.12 rad for the 800 MHz sidebands and 1.45 rad for the 80 MHz sidebands. For the output of the AOM a laser intensity of 4 mW/cm^2 is used. The effects of laser noise is negligible. These values were optimized for maximum optical pumping efficiency of the circulating atoms.

A serious concern is the efficiency of the optical pumping as the atom density increases. When injected atom are pumped from high field seeking to low field seeking their final transition into the trapped ground state emits a photon that can excite an atom out of the trapped ground state. With low atom density this photon

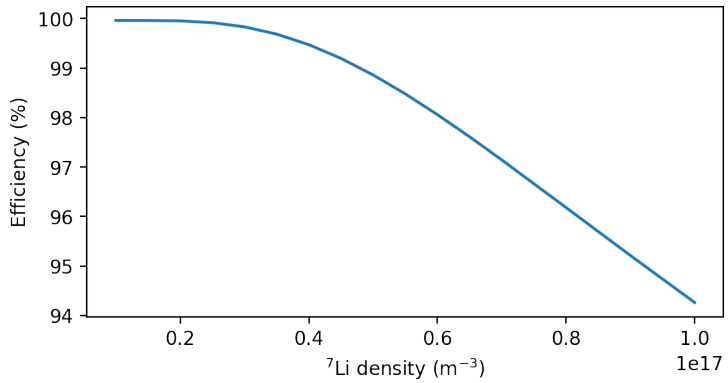

Figure A.4: Optical pumping efficiency versus atom density. The efficiency decreases with greater density because of the radiation trapping effect.

is not likely to interact with another atom. At higher density however it may excite a trapped ground state atom into an excited state. At this point the atom may fall back into the trapped ground state and emit the same frequency photon or fall into an untrapped ground state. As density increases the probability that a new trapped atom results in an untrapped atom increases towards unity. This effect can be estimated with a 3 level system, shown in Figure A.3, according to reference [122].

The results are shown in Figure A.4 and are scaled by the low field optical pumping efficiency. Density refers to the density in the optical pumping region. This can be used to determine the maximum possible flux multiplication. The density in the optical pumping region is approximately given by

$$n = \frac{F\dot{N}_0}{vA} = \frac{\dot{N}_0}{vA|\log\epsilon(n)|}, \tag{A.1}$$

where A is the area of the optical pumping region, $\dot{N}_0$ is the loaded flux of atoms, v is the speed of the atoms and Equation 6.6 was used. Solving this requires a

numerical method. Circulating atoms are spread out over a region of about 3 cm in diameter in the optical pumping magnet, the loaded flux is about $1{\times}10^{12}$ s^{-1}, and the atoms travel at about 210 m/s. Using these parameters results in $n \approx 2 \times 10^{16}$, $F \approx 1600$, $\epsilon \approx 0.9995$, and a circulating flux of about 3×10^{15}. It was assumed that the only loss mechanism is optical pumping.

Appendix B

Vacuum System

A central design goal is for the storage ring to build up a large value of atomic flux and a possible limit on this is vacuum lifetime. A typical long vacuum lifetime in a trap ranges from 1 to 100 seconds. From Equation 6.4 this spans a range of about 30 to 3000 times flux multiplication. A reasonable goal of 1000 times flux multiplication corresponds to about a 40 second vacuum lifetime. Figure B.2 and Figure B.1 show schematics of possible vacuum system for the injector and ring respectively[1]. There are 3 major possible gas loads; Hydrogen outgassing in the storage ring vacuum chamber, background helium diffusing into the storage ring from the existing chamber, and directed helium from the nozzle entering the storage ring. The ring is particularly challenging because the long bending segments are conductance limited so even with infinite pumping speeds the pressure from out gassing can only be reduced so much. To asses the impact of these sources we must first relate pressure to lifetime.

Vacuum lifetime can be expressed as

$$\tau = \frac{1}{nK_{x-y}}, \tag{B.1}$$

[1] I would like to thank Jeremy Glick for the initial work he did on the vacuum system design.

where K_{x-y} is the thermally averaged rate constant for a collision between the trapped particle x and a background gas particle y, n is the background gas density, P is the background gas pressure, and T is the background gas temperature.

Values of K_{x-y} typically assume that the trapped particle is stationary which our lithium atoms are not. However this is not a significant concern because the FWHM of the velocity distributions of Helium and Hydrogen is over 5 times larger than 200 m/s so lithium atoms appear nearly stationary. K_{x-y} can also be given for a specific trap depth because not all collisions will result in a particle being lost. It's not clear what value of trap depth is most appropriate for our situation. Collisions which do not result in a Li atom being lost immediately may still excite oscillation that result in it being lost shortly afterwords. Even if the atom is not lost, this is a form of undesired heating. A conservative choice is to use the *total* collision rate constant.

It can be difficult to find reliable values of K_{x-y}. Fortunately NIST is working on developing a vacuum standard using trapped Lithium atoms [123]. They have carried out extensive theoretical calculations of thermally averaged rate constants for Li-He and Li-H_2 [124, 125] resulting in the following values:

$$\begin{aligned} K_{Li-He} &= (1.461 \pm .013) \times 10^{-9} \text{ cm}^3/\text{s} \\ K_{Li-H_2} &= (6.0 \pm .1) \times 10^{-9} \text{ cm}^3/\text{s} \end{aligned} \tag{B.2}$$

Because the collisions of both background gases add we need to double our lifetime requirement for each to 80 seconds. This results in the following require-

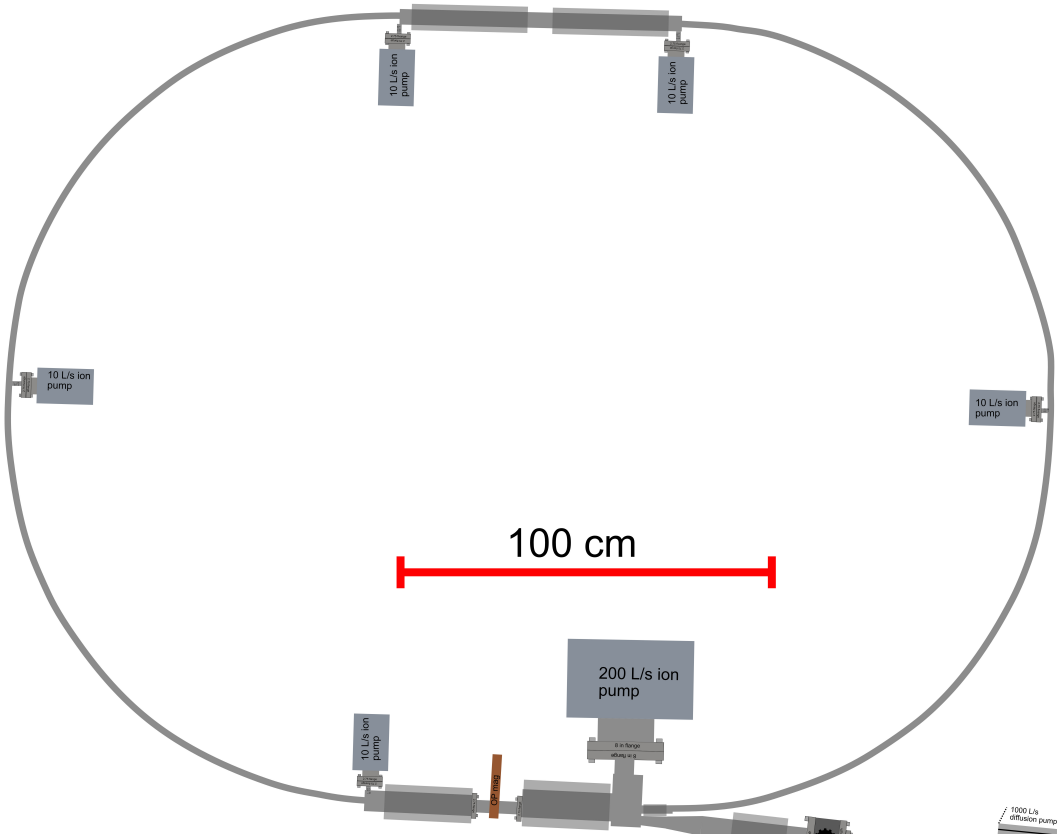

Figure B.1: Schematic of the ring vacuum system. The transparent grey squares represent lens material. The material is not shown for the bending sections.

ments:

$$P_{He} < 2.6 \times 10^{-10} \text{ torr}$$
$$P_{H_2} < 6.3 \times 10^{-11} \text{ torr}. \tag{B.3}$$

This is the Ultra High Vacuum (UHV) regime.

Hydrogen outgasses from stainless at appreciable rates in a UHV system with a nearly infinite lifetime at room temperature. A recent paper studied the

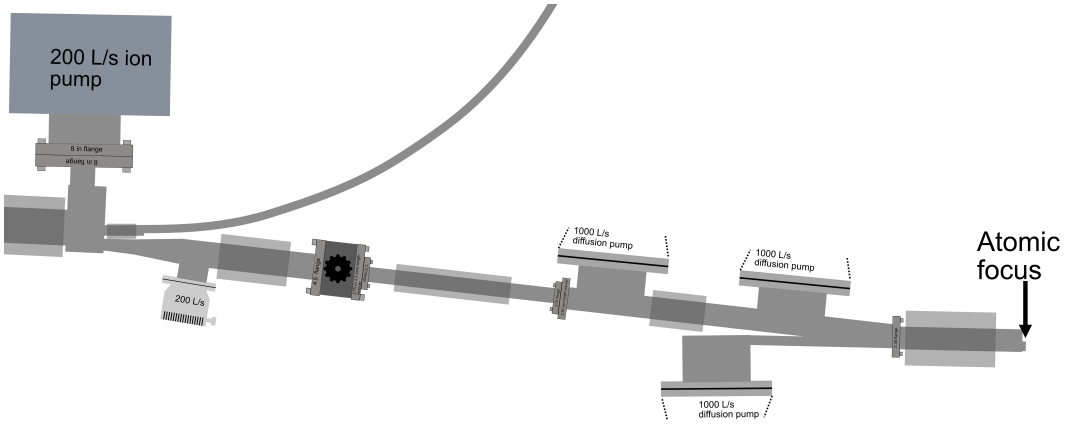

Figure B.2: Schematic of the injector vacuum system.The transparent grey squares represent lens material.

outgassing rate of H_2 from 7 chambers made of different material and/or different treatments [126]. They measured an outgassing rate of 5×10^{-12} torr L/s cm^2 for 316 stainless steel. This agrees within a factor of two with the value from CERN [127]. Assuming the vacuum system of the ring is made entirely of 316 stainless steel this would result in a mean pressure of about 5×10^{-10} torr and a lifetime of about 11 seconds. As previously mentioned, the issue is primarily the bending sections building up excessive pressure. One solution is to divide the bending sections and introduce pumps at each split until an adequate pressure is achieved. The downside to this is that additional pumps are required, the complexity of the design increases, and performance of the design may decrease. One alternative approach is to use a different material such as titanium or aluminum which both have very low H_2 outgassing. Aluminum UHV vacuum chambers are a newer technique and require hardening so that CF flanges can still be used. We elected to not

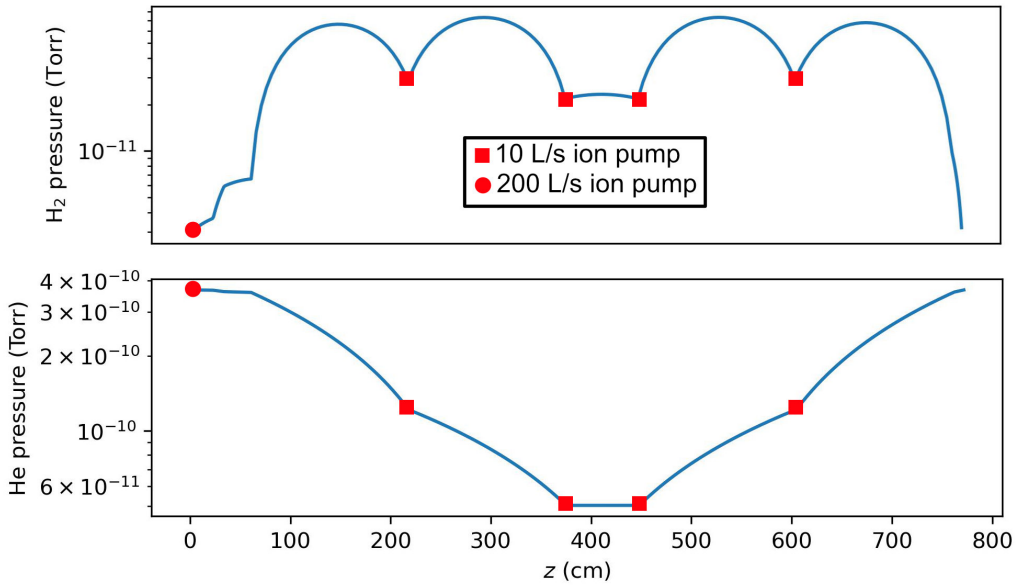

Figure B.3: He and H_2 pressure throughout the ring. The red dot indicates the location of the 200 L/s ion pump near the combiner. The He pressure is maximum at the combiner because of diffusion through the injector, and is minimum on the other side of the ring. H_2 pressure is maximum in the bending sections where the pumps are conductance limited.

consider aluminum primarily because of its novelty. Titanium is a more proven option. However experience in our group and discussion with other groups leave us concerned that is may not perform as expected.

Another promising approach is to heat treat the stainless steel under vacuum and rapidly outgas the H_2. This is known as *vacuum firing*. CERN operates a 6 meter tall 1 meter diameter cylindrical chamber for this purpose and describes that vacuum firing at 950C for 2 hours results in an outgassing rate of 1.5×10^{-13} torr L/s cm^2 for "austenitic stainless steel". They report a lower outgassing rate of "$\approx 3.8 \times 10^{-15}$ torr L/s cm^2" but assume a higher value for design purposes. Reference [126] reports that vacuum firing 316 stainless steel at 950C for results in an outgassing rate of 4×10^{-14} torr L/s cm^2.

Conservatively assuming an outgassing rate of 5×10^{-13} torr L/s cm^2 with vacuum firing results in a mean pressure of about 4.6×10^{-11} torr and a lifetime of about 110 seconds which meets our design goals. The resulting pressure profile is shown in Figure B.3. Reference [126] reports an outgassing rate of about 2×10^{-14} torr L/s cm^2 which is acceptable as well but as mentioned is more suspect.

The existing chamber has a pressure of about 2×10^{-6} torr above the diffusion pump which is dominated by He. This pressure need to be reduce by 4 orders of magnitude which requires multiple stages of differential pumping. The system begins with a 1 cm aperture at the focus. The background He gas load through the aperture is $Q = PC \approx 5 \times 10^{-7}$ torr L/s. The system branches into a line that the directed He from the nozzle follows into a dump chamber and a line that enters the storage ring vacuum. The line to the storage ring has in the following order: two 1000 L/s

diffusion pumps, a valve, and a 200 L/s turbo pump follow. The valve is to allow the storage ring vacuum to be isolated from the existing chamber when needed such as when the existing chamber needs to be vented. There is some concern using a diffusion pump in proximity to an UHV chamber because of the possibility of oil entering the UHV system. Oil back streaming cannot be completely prevented, but it may be possible to reduce it to tolerable levels. Reference [128] describes some techniques. A common solution are optically dense cooled baffles above the diffusion pump, however this reduces pump speed by abut 50%. Agilent sells a diffusion pump, the VHS-4, that an "extended cold cap" can be used with which reduces the backstreaming to the levels of an optically dense baffle while preserving 80 % of the pumping speed. For He the speed is 1200 L/s and the flange has the same dimension as diffusion pump flanges shown in the vacuum diagram.

The injection vacuum system ends in the combiner which is pumped by a 200 L/s ion pump. Unfortunately, ion pump speeds are reduced to 10% of their rated value for Helium. The system was modeled with the technique described in Section 5.6.2. The net He gas load into the storage ring is about 8×10^{-9} torr L/s. The resulting pressure profile is shown in Figure B.3. The mean pressure is about 1.7×10^{-10} torr and the vacuum life is about 120 seconds which satisfies our requirements with room to spare. There will also be a small gas load of He diffusing through glass windows.

To analyze the impact of the directed flow we can use Equation 2.7 to get the infinitesimal gas load

$$dQ \approx \gamma \cos(\theta)^2 / r^2 da, \tag{B.4}$$

where Q is the gas load in units of torr L/s, γ is an undetermined constant, and the result is approximate because we dropped the 1.15 factor. What we want is a function like $Q = Q_0 \cos(1.15\theta)$ where Q is the gas load into half angle θ and Q_0 is the total gas load of the nozzle. Some algebra yields

$$Q = Q_0(1 - \cos(\theta_0)^3), \tag{B.5}$$

where the proof is left to the reader.

The half angle of the aperture is about $0.5/176 \approx .003$ radians, and the total gas flow through the nozzle at 100 SCCM is 1.25 torr L/s. From Equation B.5 this corresponds to about 2×10^{-5} torr L/s which is substantial. If this gas load was allowed to travel with the lithium beam directly into the storage ring a vacuum speed of about 80,000 L/s would be required to meet our He pressure requirement.

An alternative approach is to move the lithium beam out of the helium beam with a shifted lens. The helium beam can then be dumped into a chamber where it is pumped away. A possible design for this is a 10 cm long pickoff tube with 1.2 cm diameter, so that it is slightly larger than the helium beam, that ends in a chamber with a 1000 L/s pump. The pressure in the dump chamber is then about 2×10^{-8} torr which results in about 1×10^{-7} torr L/s of gas load diffusing back through the pickoff tube. This is a small fraction of the gas load from background He in the existing chamber entering the injector.

Appendix C

Design Dimensions

The dimensions of the design are given in Figure C.1. The meaning of each term is described below and additional details are given. Keep in mind that all elements are either a hexapole lens or composed of hexapole lenses.

- "El #" refers to the element labeling in Figure 6.1.
- "Length" is the longitudinal material length of the element.
- "Bore ID" is the inside diameter of the element from magnet face to magnet face
- "Magnet width" is the transverse widths of the cuboidal magnets that form the element.
- "Vacuum ID" is the inside diameter of the vacuum tube through the element.
- "Input position" describes the location of the transverse radial center of the input plane of the element with x, y displacements. In this coordinate system the focus is the origin and the nozzle is aligned along -$\hat{x}$ as shown in Figure 6.1.
- "Angle" is the angle the input of the element makes with $+\hat{x}$.

- "Bender radius" is the radius of the bender measured at the center of a single hexapole lens segment.
- "Bend angle" is the total angle that the particle is deflected by traveling through the bender.
- "Centrifugal offset" is the radial offset of the design orbit because of the centrifugal pseudo-force. This number is added to the "Bender radius" to get the radius of the design orbit.
- "Number of segments" is the number of full length hexapole lens segments. The beginning and end of a bender is a *half* length segment. Thus N segments mean $N-1$ full length segments and 2 half length segments.
- "Segment length" is the longitudinal length of the cuboidal magnets that compose each hexapole lens segment in the benders.
- The thinnest vacuum wall is 40 thousandths of an inch. For some elements this may be larger so that the interpolation scheme shown in Figure 5.2 works correctly without using an unnecessary number of interpolation points. The slightly restricted aperture does not affect results meaningfully in those elements. In the bender section however it is important to use a wall thickness near this value.
- In the combiner the wall thickness is larger because it is assumed the vacuum system in that region will be more like a vacuum chamber than a circular vacuum tube.

- More details of the combiner are given in Figure C.2.

Lenses

El #	Length	Bore ID	Magnet width	Vacuum tube ID	Input position	Angle
1	19.04	4.84	1.27,1.905	4.555	-3,1.14	180
2	11.43	4.84	1.27	4.555	-64.151,5.14	174.614
3	24.327	3.649	.952	3.434	-105.108,9.001	174.614
4	14.788	6.0	1.582	5.646	-151.721,13.395	174.614
6	23.199	6.0	1.582	5.646	-235.891,19.482	178.633
8	3.876	1.366	.340	1.286	-362.318,124.102	89.250
10	30.815	4.5	1.18	4.235	-256.983,229.309	359.868
11	30.815	4.5	1.18	4.235	-219.418,229.222	359.868
13	3.876	1.366	.340	1.286	-83.619,127.431	270.634
15	6.115	2.0	.51	1.882	-183.923,19.596	181.399

Combiner

El #	Length	Bore ID	Magnet width	Vacuum tube ID	Input position	Angle
5	23.456	8	2.118	7.531	-198.598,17.571	178.633

Benders

El #	Bender radius	Bend angle	Input position	Angle
7	199.449	89.3829	-264.794,20.394	178.633
9	199.449	89.3829	-362.010,130.678	89.250
12	199.777	89.2341	-183.523,228.916	359.868
14	199.777	89.2341	-83.769,120.848	270.634

Bore ID	Magnet width	Vacuum tube ID	Centrifugal offset	Number of segments	Segment length
1.624	.409	1.416	.222	121	1/2 in

Figure C.1: Dimensions and parameters of each element in the design. Distance units are centimeters and angle units are degrees unless otherwise specified.

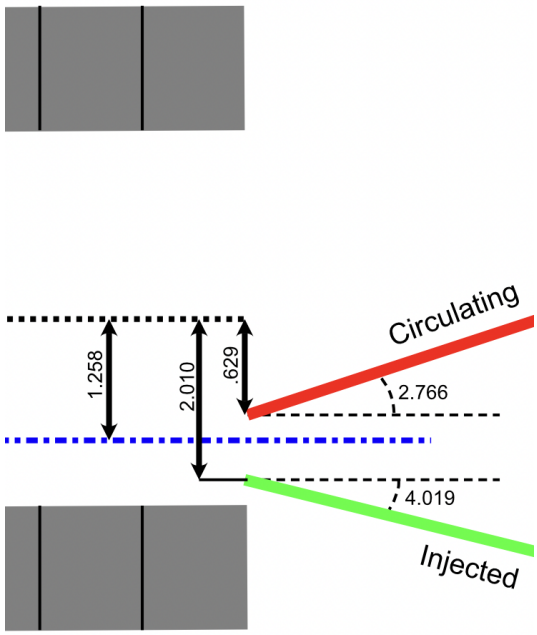

Figure C.2: Dimensions and parameters of the loading scheme. Distance units are centimeters and angle units are degrees. The black dashed line is the centerline of the hexapole lens. The blue dot dashed line represents the offset from the centerline where the circulating and injected beam join at the output.

Appendix D

Lithium D Line

Figure D.1 depicts the Lithium D line transition and relevant parameters for reference.

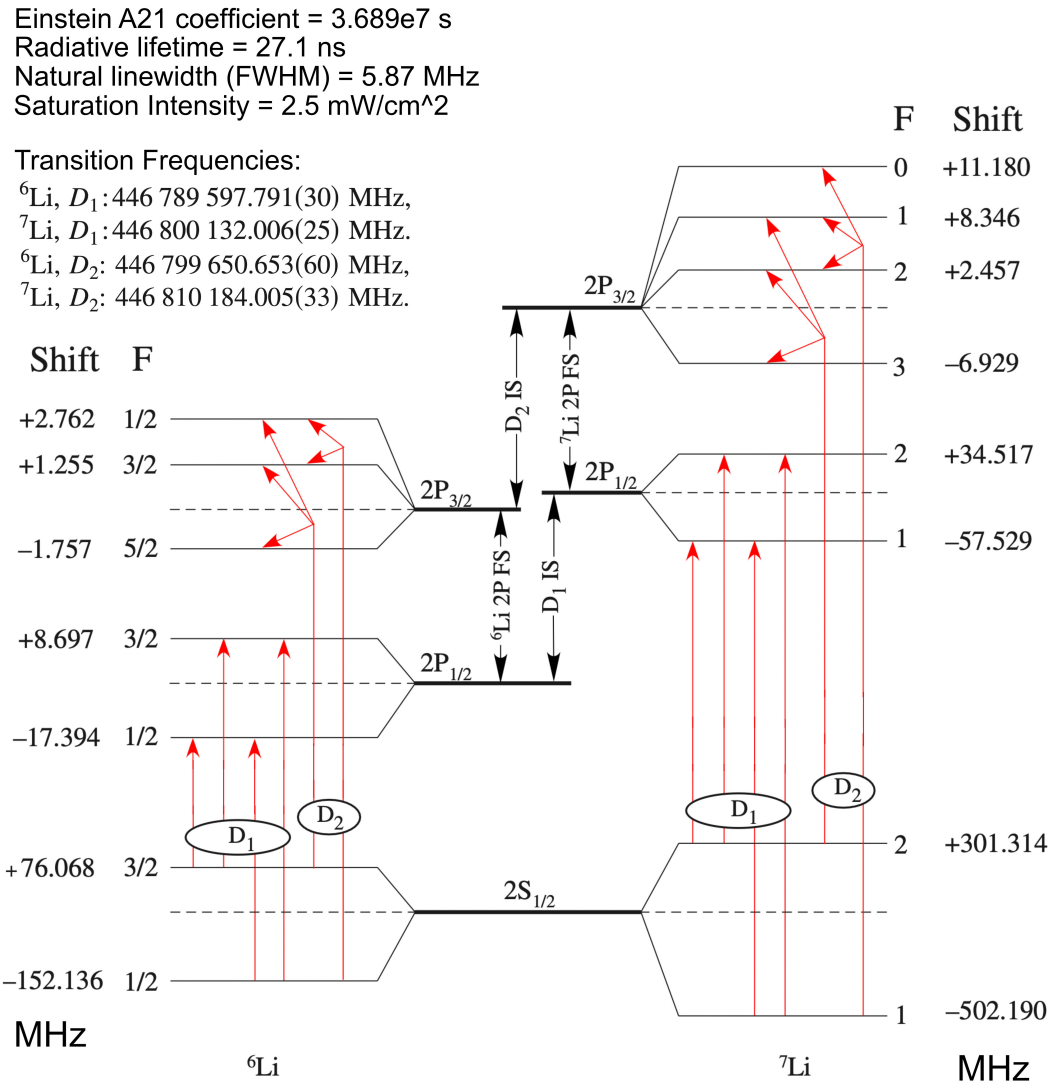

Figure D.1: Figure modified from [129].

Appendix E

Stern-Gerlach Magnet

We initially considered using a Stern-Gerlach style magnet for continuous injection. With this magnet the horizontal bending force is nearly constant. It can be constructed by the combination of a dipole and quadrupole multipole term (Equation 2.17). The magnetic field is given by

$$\vec{B} = \vec{B}_{\text{dipole}} + \vec{B}_{\text{quadrupole}} = c_2 y\hat{x} + (c_1 + c_2 x)\hat{y}, \tag{E.1}$$

where c_1 is the strength of the dipole term, c_2 is the strength of the quadrupole term, and x and y are transverse coordinates. This results in a force of

$$F_x = \frac{-\mu_B c_2(c_1 + c_2 x)}{|\vec{B}|}\hat{x} + \frac{-\mu_B c_2 y}{|\vec{B}|}\hat{y} \tag{E.2}$$

where it is assumed the atom is in a low field seeking state. The bending and combination of both beams is in the x direction. For small values of x and y this gives a constant bending force of

$$F_x \approx -\mu_B c_2 \tag{E.3}$$

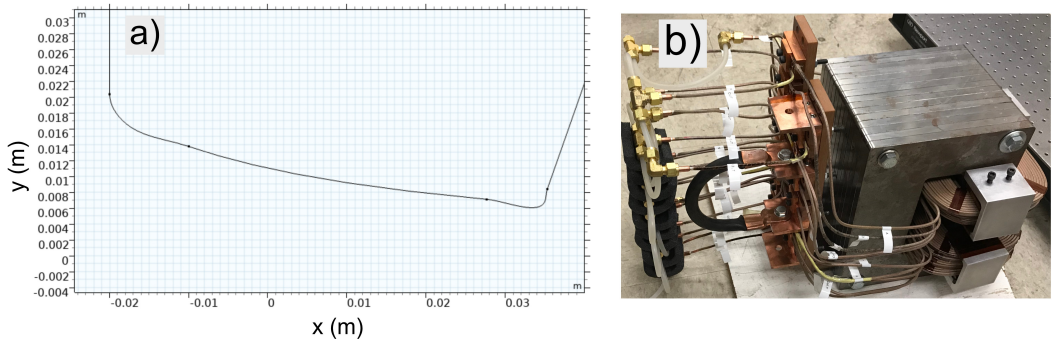

Figure E.1: a) Profile of pole face of Stern-Gerlach style combiner magnet. The origin (0,0) is the center of the combiner and location of optimal bending. b) Image of the completed design. The magnet is composed of 15 7.5 in long sheets.

We designed and built a Stern-Gerlach combiner magnet [1], shown in Figure E.1 . The required field is produced by a "c" shaped ferromagnetic and water cooled current carrying coils. The multipole values are $c_1 \approx 1$ T and $c_2 \approx 17.5$ T/m. A major downside of this approach is that there is very minimal confinement in x and thus particles effectively drift through this region compared to the design orbit. Additionally, the specific design suffers from a small bore aperture of 1.2 cm vertical height. Simulations indicated poor performance. A permanent magnet hexapole lens yielded better results.

[1] This magnet was designed by a previous undergraduate student Collin Diver as a senior project.

Bibliography

[1] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, "Search for new physics with atoms and molecules," *Rev. Mod. Phys.*, vol. 90, p. 025008, Jun 2018.

[2] R. Wynands and S. Weyers, "Atomic fountain clocks," *Metrologia*, vol. 42, no. 3, pp. S64–S79, 2005.

[3] K. Bong, M. Holynski, J. Vovrosh, P. Bouyer, G. Condon, E. Rasel, C. Schubert, W. P. Schleich, and A. Roura, "Taking atom interferometric quantum sensors from the laboratory to real-world applications," *Nature Rev. Phys.*, vol. 1, no. 12, pp. 731–739, 2019.

[4] G. M. Tino, "Testing gravity with cold atom interferometry: results and prospects," *Quantum Sci. Technol.*, vol. 6, p. 024014, 2021.

[5] 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, and X. Wu, "Improved limit on the electric dipole moment of the electron," *Nature*, vol. 562, no. 7727, pp. 355–360, 2018.

[6] J. Jankunas and A. Osterwalder, "Cold and controlled molecular beams: Production and applications," *Ann. Rev. Phys. Chem.*, vol. 66, no. 1, pp. 241–262, 2015.

[7] B. R. Heazlewood and T. P. Softley, "Towards chemistry at absolute zero," *Nature Rev. Chem.*, vol. 5, no. 2, pp. 125–140, 2021.

[8] W. G. Kaenders, F. Lison, I. Müller, A. Richter, R. Wynands, and D. Meschede, "Refractive components for magnetic atom optics," *Phys. Rev. A*, vol. 54, pp. 5067–5075, Dec 1996.

[9] P. L. Walstrom, M. Di Rosa, *et al.*, "Beam transport and storage with cold neutral atoms and molecules," tech. rep., Los Alamos National Laboratory, 2012.

[10] W. D. Phillips and H. Metcalf, "Laser deceleration of an atomic beam," *Phys. Rev. Lett.*, vol. 48, no. 9, pp. 596–599, 1982.

[11] C. Slowe, L. Vernac, and L. V. Hau, "High flux source of cold rubidium atoms," *Rev. Sci. Instrum.*, vol. 76, p. 103101, 2005.

[12] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman, "Low-velocity intense source of atoms from a magneto-optical trap," *Phys. Rev. Lett.*, vol. 77, no. 16, pp. 3331–3334, 1996.

[13] K. I. Lee, J. A. Kim, H. R. Noh, and W. Jhe, "Single-beam atom trap in a pyramidal and conical hollow mirror," *Opt. Lett.*, vol. 21, no. 15, pp. 1177–1179, 1996.

[14] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven, "Two-dimensional magneto-optical trap as a source of slow atoms," *Phys. Rev. A*, vol. 58, no. 5, pp. 3891–3895, 1998.

[15] N. Fitch and M. Tarbutt, *Laser-cooled molecules*, vol. 70, pp. 157–262. Academic Press, 2021.

[16] X. Yu, J. Mo, T. Lu, T. Y. Tan, and T. L. Nicholson, "Magneto-optical trapping of a group-III atom," *Physical Review A*, vol. 105, jun 2022.

[17] R. E. Smalley, L. Wharton, and D. H. Levy, "Molecular optical spectroscopy with supersonic beams and jets," *Accounts Chem. Res.*, vol. 10, no. 4, pp. 139–145, 1977.

[18] M. R. Tarbutt, J. J. Hudson, B. E. Sauer, E. A. Hinds, V. A. Ryzhov, V. L. Ryabov, and V. F. Ezhov, "A jet beam source of cold ybf radicals," *J. Phys. B*, vol. 35, no. 24, pp. 5013–5022, 2002.

[19] B. Yan, P. F. H. Claus, B. G. M. van Oorschot, L. Gerritsen, A. T. J. B. Eppink, S. Y. T. van de Meerakker, and D. H. Parker, "A new high intensity and short-pulse molecular beam valve," *Rev. Sci. Instrum.*, vol. 84, no. 2, pp. 023102–023102, 2013.

[20] P. Aggarwal, H. L. Bethlem, A. Boeschoten, A. Borschevsky, K. Esajas, Y. Hao, S. Hoekstra, K. Jungmann, V. R. Marshall, T. B. Meijknecht, M. C. Mooij, R. G. E. Timmermans, A. Touwen, W. Ubachs, L. Willmann, Y. Yin, and A. Zapara, "A supersonic laser ablation beam source with narrow velocity spreads," *Rev. of Sci. Instrum.*, vol. 92, no. 3, pp. 033202–033202, 2021.

[21] M. Hillenkamp, S. Keinan, and U. Even, "Condensation limited cooling in supersonic expansions," *The Journal of chemical physics*, vol. 118, no. 19,

pp. 8699–8705, 2003.

[22] K. S. Melin, P. I. Nagornykh, Y. Lu, L. E. Hillberry, Y. Xu, and M. G. Raizen, "Observation of a quasi-one-dimensional variation of the stern-gerlach effect," *Phys. Rev. A*, vol. 99, p. 063417, Jun 2019.

[23] H. L. Bethlem, G. Berden, and G. Meijer, "Decelerating neutral dipolar molecules," *Phys. Rev. Lett.*, vol. 83, no. 8, pp. 1558–1561, 1999.

[24] E. Narevicius, C. G. Parthey, A. Libson, J. Narevicius, I. Chavez, U. Even, and M. G. Raizen, "An atomic coilgun: using pulsed magnetic fields to slow a supersonic beam," *New J. Phys.*, vol. 9, no. 10, pp. 358–358, 2007.

[25] N. R. Hutzler, H.-I. Lu, and J. M. Doyle, "The buffer gas beam: An intense, cold, and slow source for atoms and molecules," *Chem. Rev.*, vol. 112, no. 9, pp. 4803–4827, 2012.

[26] M. Borysow, *A high-intensity cold atom source.* PhD thesis, University of Texas at Austin, 2012.

[27] J. Wang, V. A. Shamamian, B. R. Thomas, J. M. Wilkinson, J. Riley, C. F. Giese, and W. R. Gentry, "Speed ratios greater than 1000 and temperatures less than 1 mk in a pulsed he beam," *Phys. Rev. Lett.*, vol. 60, pp. 696–699, Feb 1988.

[28] R. R. Chaustowski, V. Y. F. Leung, and K. G. H. Baldwin, "Magnetic hexapole lens focusing of a metastable helium atomic beam for uv-free lithography," *Applied physics. B*, vol. 86, no. 3, pp. 491–496, 2007.

[29] E. Anciaux, G. Stratis, and M. G. Raizen, "Brightening of a supersonic beam of neutral atoms," *Phys. Scripta*, vol. 93, no. 12, pp. 124009–, 2018.

[30] M. Benakli, S. Raghavan, A. Smerzi, S. Fantoni, and S. R. Shenoy, "Macroscopic angular-momentum states of bose-einstein condensates in toroidal traps," *Europhysics Letters*, vol. 46, p. 275, may 1999.

[31] E. Nugent, D. McPeake, and J. F. McCann, "Superfluid toroidal currents in atomic condensates," *Phys. Rev. A*, vol. 68, p. 063606, Dec 2003.

[32] F. Bloch, "Superfluidity in a ring," *Phys. Rev. A*, vol. 7, pp. 2187–2191, Jun 1973.

[33] L. Garay, J. Anglin, J. Cirac, and P. Zoller, "Sonic analog of gravitational black holes in bose-einstein condensates," *Physical review letters*, vol. 85, no. 22, pp. 4643–4647, 2000.

[34] S. Eckel, J. G. Lee, F. Jendrzejewski, N. Murray, C. W. Clark, C. J. Lobb, W. D. Phillips, M. Edwards, and G. K. Campbell, "Hysteresis in a quantized superfluid 'atomtronic' circuit," *Nature (London)*, vol. 506, no. 7487, pp. 200–203, 2014.

[35] B. T. Seaman, M. Krämer, D. Z. Anderson, and M. J. Holland, "Atomtronics: Ultracold-atom analogs of electronic devices," *Phys. Rev. A*, vol. 75, p. 023615, Feb 2007.

[36] L. Amico, M. Boshier, G. Birkl, A. Minguzzi, C. Miniatura, L.-C. Kwek, D. Aghamalyan, V. Ahufinger, D. Anderson, N. Andrei, A. S. Arnold, M. Baker,

T. A. Bell, T. Bland, J. P. Brantut, D. Cassettari, W. J. Chetcuti, F. Chevy, R. Citro, S. De Palo, R. Dumke, M. Edwards, R. Folman, J. Fortagh, S. A. Gardiner, B. M. Garraway, G. Gauthier, A. Günther, T. Haug, C. Hufnagel, M. Keil, P. Ireland, M. Lebrat, W. Li, L. Longchambon, J. Mompart, O. Morsch, P. Naldesi, T. W. Neely, M. Olshanii, E. Orignac, S. Pandey, A. Pérez-Obiol, H. Perrin, L. Piroli, J. Polo, A. L. Pritchard, N. P. Proukakis, C. Rylands, H. Rubinsztein-Dunlop, F. Scazza, S. Stringari, F. Tosto, A. Trombettoni, N. Victorin, W. v. Klitzing, D. Wilkowski, K. Xhani, and A. Yakimenko, "Roadmap on atomtronics: State of the art and perspective," *AVS Quantum Science*, vol. 3, no. 3, p. 039201, 2021.

[37] C. Ryu, P. W. Blackburn, A. A. Blinova, and M. G. Boshier, "Experimental realization of josephson junctions for an atom squid," *Phys. Rev. Lett.*, vol. 111, p. 205301, Nov 2013.

[38] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt, "Atom interferometry with bose-einstein condensates in a double-well potential," *Phys. Rev. Lett.*, vol. 92, p. 050405, Feb 2004.

[39] O. Garcia, B. Deissler, K. J. Hughes, J. M. Reeves, and C. A. Sackett, "Bose-einstein-condensate interferometer with macroscopic arm separation," *Phys. Rev. A*, vol. 74, p. 031601, Sep 2006.

[40] S. Hou, B. Wei, L. Deng, and J. Yin, "A novel molecular synchrotron for cold collision and edm experiments," *Scientific Reports*, 2016.

[41] W. H. Heathcote, E. Nugent, B. T. Sheard, and C. J. Foot, "A ring trap for ultracold atoms in an rf-dressed state," *New journal of physics*, vol. 10, no. 4, pp. 043012–, 2008.

[42] B. E. Sherlock, M. Gildemeister, E. Owen, E. Nugent, and C. J. Foot, "Time-averaged adiabatic ring potential for ultracold atoms," *Physical review. A, Atomic, molecular, and optical physics*, vol. 83, no. 4, 2011.

[43] G. D. Bruce, J. Mayoh, G. Smirne, L. Torralbo-Campo, and D. Cassettari, "A smooth, holographically generated ring trap for the investigation of superfluidity in ultracold atoms," *Physica Scripta*, vol. T143, p. 014008, feb 2011.

[44] E. Courtade, O. Houde, J.-F. Clement, P. Verkerk, and D. Hennequin, "Dark optical lattice of ring traps for cold atoms," *Physical review. A, Atomic, molecular, and optical physics*, vol. 74, no. 3, pp. 031403(R)–, 2006.

[45] T. A. Bell, J. A. P. Glidden, L. Humbert, M. W. J. Bromley, S. A. Haine, M. J. Davis, T. W. Neely, M. A. Baker, and H. Rubinsztein-Dunlop, "Bose-einstein condensation in large time-averaged optical ring potentials," *New journal of physics*, vol. 18, no. 3, pp. 35003–, 2016.

[46] S. Beattie, S. Moulder, R. J. Fletcher, and Z. Hadzibabic, "Persistent currents in spinor condensates," *Phys. Rev. Lett.*, vol. 110, p. 025301, Jan 2013.

[47] A. Ramanathan, K. C. Wright, S. R. Muniz, M. Zelan, W. T. Hill, C. J. Lobb, K. Helmerson, W. D. Phillips, and G. K. Campbell, "Superflow in a toroidal

bose-einstein condensate: An atom circuit with a tunable weak link," *Phys. Rev. Lett.*, vol. 106, p. 130401, Mar 2011.

[48] F. Crompvoets, H. Bethlem, R. Jongma, and G. Meijer, "A prototype storage ring for neutral molecules," *Nature (London)*, vol. 411, no. 6834, pp. 174–176, 2001.

[49] P. C. Zieger, C. J. Eyles, S. Y. van de Meerakker, A. J. van Roij, H. L. Bethlem, and G. Meijer, "A forty-segment molecular synchrotron," *Zeitschrift für physikalische Chemie (Neue Folge)*, vol. 227, no. 11, pp. 1605–1645, 2013.

[50] C. Heiner, D. Carty, G. Meijer, and H. Bethlem, "A molecular synchrotron," *Nature physics*, vol. 3, no. 2, pp. 115–118, 2007.

[51] K.-J.Kügler, K.Moritz, W.Paul, and U.Trinks, "Nestor — a magnetic storage ring for slow neutrons," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 228, no. 2, pp. 240–258, 1985.

[52] D. Thompson, R. V. E. Lovelace, and D. M. Lee, "Storage rings for spin-polarized hydrogen," *J. Opt. Soc. Am. B*, vol. 6, pp. 2227–2234, Nov 1989.

[53] W. Ketterle and D. E. Pritchard, "Trapping and focusing ground state atoms with static fields," *Applied Physics B*, 1992.

[54] N. D. Bhaskar, J. Pietras, J. Camparo, W. Happer, and J. Liran, "Spin destruction in collisions between cesium atoms," *Phys. Rev. Lett.*, vol. 44, pp. 930–933, Apr 1980.

[55] S. Gupta, K. W. Murch, K. L. Moore, T. P. Purdy, and D. M. Stamper-Kurn, "Bose-einstein condensation in a circular waveguide," *Physical Review Letters*, vol. 95, no. 14, 2005.

[56] J. A. Sauer, M. D. Barrett, and M. S. Chapman, "Storage ring for neutral atoms," *Physical review letters*, vol. 87, no. 27 Pt 1, pp. 270401–270401, 2001.

[57] A. S. Arnold, C. S. Garvie, and E. Riis, "Large magnetic storage ring for bose-einstein condensates," *Phys. Rev. A*, vol. 73, p. 041606, Apr 2006.

[58] S. Wu, W. Rooijakkers, P. Striehl, and M. Prentiss, "Bidirectional propagation of cold atoms in a "stadium"-shaped magnetic guide," *Phys. Rev. A*, vol. 70, p. 013409, Jul 2004.

[59] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard, "Optics and interferometry with atoms and molecules," *Reviews of Modern Physics*, vol. 81, pp. 1051–1129, jul 2009.

[60] I. Vázquez, M. P. Russell, D. R. Smith, and R. Radebaugh, *Helium Adsorption on Activated Carbons at Temperatures between 4 and 76 K*, pp. 1013–1021. Boston, MA: Springer US, 1988.

[61] D. R. Miller, *Free jet sources*, vol. 1, pp. 14–53. Oxford Univ. Press, 1988.

[62] H. Ashkenas and F. S. Sherman, "The Structure and Utilization of Supersonic Free Jets in Low Density Wind Tunnels," in *Rarefied Gas Dynamics, Volume 2*, vol. 2, p. 84, Jan. 1965.

[63] O. F. Hagena, "Condensation in free jets: Comparison of rare gases and metals," *Z. Phys. D*, vol. 4, no. 3, pp. 291–299, 1987.

[64] P. W. Stephens and J. G. King, "Experimental investigation of small helium clusters: magic numbers and the onset of condensation," *Phys. Rev. Lett.*, vol. 51, no. 17, pp. 1538–1541, 1983.

[65] H. Bichneau, E. L. Knuth, J. Northby, J. P. Toennies, and C. Winkler, "Mass spectra and time-of-flight distributions of helium cluster beams," *J. Chem. Phys.*, vol. 92, no. 11, pp. 6875–6889, 1990.

[66] L. W. Bruch, W. Schöllkopf, and J. P. Toennies, "The formation of dimers and trimers in free jet 4he cryogenic expansions," *J. Chem. Phys.*, vol. 117, no. 4, pp. 1544–1566, 2002.

[67] J. Wörmer, V. Guzielski, J. Stapelfeldt, and T. Möller, "Fluorescence excitation spectroscopy of xenon clusters in the vuv," *Chem. Phys. Lett.*, vol. 159, no. 4, pp. 321–326, 1989.

[68] R. A. Smith, T. Ditmire, and J. W. G. Tisch, "Characterization of a cryogenically cooled high-pressure gas jet for laser/cluster interaction experiments," *Rev. Sci. Instrum.*, vol. 69, no. 11, pp. 3798–3804, 1998.

[69] R. E. Grisenti, W. Schollkopf, J. P. Toennies, G. C. Hegerfeldt, T. Kohler, and M. Stoll, "Determination of the bond length and binding energy of the helium dimer by diffraction from a transmission grating," *Phys. Rev. Lett.*, vol. 85, no. 11, pp. 2284–2287, 2000.

[70] M. Chrysos, "Dirac bubble potential for he–he and inadequacies in the continuum: Comparing an analytic model with elastic collision experiments," *J. Chem. Phys.*, vol. 146, p. 024106, 2017.

[71] H. Pauly, *Atom, Molecule, and Cluster Beams I.* Berlin: Springer, 2000.

[72] M. Hillenkamp, S. Keinan, and U. Even, "Condensation limited cooling in supersonic expansions," *The Journal of Chemical Physics*, vol. 118, no. 19, pp. 8699–8705, 2003.

[73] G. Tejeda, B. Maté, J. M. Fernández-Sánchez, and S. Montero, "Temperature and density mapping of supersonic jet expansions using linear raman spectroscopy," *Phys. Rev. Lett.*, vol. 76, no. 1, pp. 34–37, 1996.

[74] C. Makrides, D. S. Barker, J. A. Fedchak, J. Scherschligt, S. Eckel, and E. Tiesinga, "Collisions of room-temperature helium with ultracold lithium and the van der waals bound state of heli," *Phys. Rev. A*, vol. 101, p. 012702, Jan 2020.

[75] E. Tiesinga and J. Klos. personal communication, 1 May 2021.

[76] Q. Shu, R. W. Fast, and H. L. Hart, "An experimental study of heat transfer in multilayer insulation systems from room temperature to 77 k," *Adv. Cryog. Eng.*, vol. 31, pp. 455–463, 01 1986.

[77] R. Li, A. Onishi, T. Satoh, and Y. Kanazawa, *Temperature Stabilization on Cold Stage of 4 K G-M Cryocooler*, pp. 765–771. Boston, MA: Springer US, 1997.

[78] A. L. Woodcraft, "Recommended values for the thermal conductivity of aluminium of different purities in the cryogenic to room temperature range, and a comparison with copper," *Cryogenics*, vol. 45, no. 9, pp. 626–636, 2005.

[79] M. C. John Moore, Christopher Davis, *Building Scientific Apparatus*. Perseus Books, 2002.

[80] D. A. Steck, *Quantum and Atom Optics*. John Wiley & Sons, Inc, 2022.

[81] J. Rossbach and P. Schmueser, "Basic course on accelerator optics," in *CAS-CERN Accelerator School: 5th general accelerator physics course* (S. Turner, ed.), Association for Computing Machinery, 1992.

[82] K. Halbach, "Design of permanent multipole magnets with oriented rare earth cobalt material," *Nuclear Instruments and Methods*, vol. 169, no. 1, pp. 1–10, 1980.

[83] D. M. Technologes, "B-h curves of nd-fe-b." https://www.dextermag.com/wp-content/uploads/2016/10/Dexter_Nd-Fe-B__Curves_Properties_Table.pdf. Copyright 2010. For Reference Only. V102.

[84] M. Katter, "Angular dependence of the demagnetization stability of sintered nd-fe-b magnets," *IEEE Transactions on Magnetics*, vol. 41, no. 10, pp. 3853–3855, 2005.

[85] J. H. E. Peter W. Milonni, *Laser Physics*. John Wiley & Sons, Inc, 2010.

[86] F. E. H. George Arfken, Hans Weber, *Mathematical Methods for Physicists*. Academic Press, 2012.

[87] C. Tsallis, "Possible generalization of boltzmann-gibbs statistics," *J. Stat. Phys.*, vol. 52, no. 1-2, pp. 479–487, 1988.

[88] P. Douglas, S. Bergamini, and F. Renzoni, "Tunable tsallis distributions in dissipative optical lattices," *Phys. Rev. Lett.*, vol. 96, no. 11, pp. 110601–110601, 2006.

[89] H. Beijerinck and N. Verster, "Absolute intensities and perpendicular temperatures of supersonic beams of polyatomic gases," *Physica B+C*, vol. 111, no. 2, pp. 327–352, 1981.

[90] R. Kersevan and J.-L. Pons, "Introduction to molflow+: New graphical processing unit-based monte carlo code for simulating molecular flows and for calculating angular coefficients in the compute unified device architecture environment," *J. Vac. Sci. Technol.*, vol. 27, no. 4, pp. 1017–1023, 2009.

[91] H. Wiedemann, *Particle Accelerator Physics*. Springer Berlin Heidelberg, 2015.

[92] J. Tanabe, "Iron dominated electromagnets design, fabrication, assembly and measurements," tech. rep., Stanford Linear Accelerator Center, 2005.

[93] K. Wille, *The Physics of Particle Accelerators*. Oxford University Press, 2000.

[94] M. Reiser, *Theory and Design of Charged Particle Beams*. Theoretical and Computational Chemistry Series, Wiley, 2008.

[95] S. S. S.K. Lam, A. Pitrou, "Llvm '15: Proceedings of the second workshop on the llvm compiler infrastructure in hpc," in *Numba: a LLVM-based Python JIT compiler*, Association for Computing Machinery, 2015.

[96] E. Vouga, "Physical simulation from geometric first principles : Course notes for physical simulation," September 2022.

[97] M. Ortner and L. G. Coliado Bandeira, "Magpylib: A free python package for magnetic field computation," *SoftwareX*, 2020.

[98] O. Chadebec, J.-L. Coulomb, and F. Janet, "A review of magnetostatic moment method," *IEEE Transactions on Magnetics*, vol. 42, no. 4, pp. 515–520, 2006.

[99] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.

[100] M. Gorelick and I. Ozsvald, *High Performance Python: Practical Performant Programming for Humans*. O'Reilly Media, 2020.

[101] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," *Nature*, vol. 585, pp. 357–362, Sept. 2020.

[102] J. D. Hunter, "Matplotlib: A 2d graphics environment," *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.

[103] "Manipulation and analysis of geometric objects in the cartesian plane." https://github.com/shapel

[104] M. McKerns and M. Aivazis, "pathos: a framework for heterogeneous computing," 2010-. https://uqfoundation.github.io/project/pathos.

[105] T. Hanrahan, M. Johnson, and C. Wang, "A closed cycle 1k refrigerator precooled by a 4k pulse tube cryocooler," tech. rep., Cryomech, Inc., 2018.

[106] C. Wang, B. Lichtenwalter, A. Friebel, and H. X. Tang, "A closed-cycle 1k refrigeration cryostat," *Cryogenics*, vol. 64, pp. 5–9, 2014.

[107] H. Zu, W. Dai, and A. de Waele, "Development of dilution refrigerators—a review," *Cryogenics*, vol. 121, p. 103390, 2022.

[108] F. Lison, P. Schuh, D. Haubrich, and D. Meschede, "High-brilliance zeeman-slowed cesium atomic beam," *Phys. Rev. A*, vol. 61, p. 013405, Dec 1999.

[109] C. Slowe, L. Vernac, and L. V. Hau, "High flux source of cold rubidium atoms," *Review of Scientific Instruments*, vol. 76, no. 10, p. 103101, 2005.

[110] G. Lamporesi, S. Donadello, S. Serafini, and G. Ferrari, "Compact high-flux source of cold sodium atoms," *Review of Scientific Instruments*, vol. 84, no. 6, p. 063102, 2013.

[111] T. Mihara, Y. Iwashita, M. Kumada, and C. Spencer, "Variable permanent magnet quadrupole," *IEEE Transactions on Applied Superconductivity*, vol. 16, no. 2, pp. 224–227, 2006.

[112] G. Tosin, P. Palma Sanchez, J. F. Citadini, and C. Castro Vergasta, "Super hybrid quadrupoles," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 674, pp. 67–73, 2012.

[113] J. Cui, J. Ormerod, D. Parker, and et al., "Manufacturing processes for permanent magnets: Part i—sintering and casting," *The Journal of The Minerals, Metals & Materials Society*, vol. 74, pp. 1279–1295, 2022.

[114] S. Turner, ed., *Measurement and Alignment of Accelerator and Detector Magnets*, CERN Accelerator School (CAS), (Geneva, Switzerland), CERN, 1998.

[115] C. C. Bradley, C. A. Sackett, and R. G. Hulet, "Bose-einstein condensation of lithium: Observation of limited condensate number," *Phys. Rev. Lett.*, vol. 78, pp. 985–989, Feb 1997.

[116] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, "Observation of bose-einstein condensation in a dilute atomic vapor," *Science*, vol. 269, no. 5221, pp. 198–201, 1995.

[117] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, "Bose-einstein condensation of cesium," *Science*, vol. 299, no. 5604, pp. 232–235, 2003.

[118] G. Modugno, G. Ferrari, G. Roati, R. J. Brecha, A. Simoni, and M. Inguscio, "Bose-einstein condensation of potassium atoms by sympathetic cooling," *Science*, vol. 294, no. 5545, pp. 1320–1322, 2001.

[119] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, "Bose-einstein condensation in a gas of sodium atoms," *Phys. Rev. Lett.*, vol. 75, pp. 3969–3973, Nov 1995.

[120] K. B. Davis, M. O. Mewes, and W. Ketterle, "An analytical model for evaporative cooling of atoms," *Applied Physics B Laser and Optics*, vol. 60, no. 2-3, p. 155–159, 1995.

[121] E. R. I. Abraham, W. I. McAlexander, J. M. Gerton, R. G. Hulet, R. Côté, and A. Dalgarno, "Triplet s-wave resonance in ^{6}Li collisions and scattering lengths of ^{6}Li and ^{7}Li," *Phys. Rev. A*, vol. 55, pp. R3299–R3302, May 1997.

[122] D. Tupa, L. W. Anderson, D. L. Huber, and J. E. Lawler, "Effect of radiation trapping on the polarization of an optically pumped alkali-metal vapor," *Phys. Rev. A*, vol. 33, pp. 1045–1051, Feb 1986.

[123] J. Scherschligt, J. A. Fedchak, D. S. Barker, S. Eckel, N. Klimov, C. Makrides, and E. Tiesinga, "Development of a new UHV/XHV pressure standard (cold atom vacuum standard)," *Metrologia*, vol. 54, pp. S125–S132, nov 2017.

[124] C. Makrides, D. S. Barker, J. A. Fedchak, J. Scherschligt, S. Eckel, and E. Tiesinga, "Elastic rate coefficients for $Li + h_2$ collisions in the calibration of a cold-atom vacuum standard," *Phys. Rev. A*, vol. 99, p. 042704, Apr 2019.

[125] C. Makrides, D. S. Barker, J. A. Fedchak, J. Scherschligt, S. Eckel, and E. Tiesinga, "Collisions of room-temperature helium with ultracold lithium and the van der waals bound state of heli," *Phys. Rev. A*, vol. 101, p. 012702, Jan 2020.

[126] J. A. Fedchak, J. K. Scherschligt, S. Avdiaj, D. S. Barker, S. P. Eckel, B. Bowers, S. O'Connell, and P. Henderson, "Outgassing rate comparison of seven geometrically similar vacuum chambers of different materials and heat treatments," *Journal of Vacuum Science & Technology B*, vol. 39, no. 2, p. 024201, 2021.

[127] P. Chiggiato, "Materials and properties iv outgassing," tech. rep., CERN Accelerator School (CAS) on Vacuum for Particle Accelerators, june 2017.

[128] J. R. Gaines, M. D. Healy, and L. Rankin, "Backstreaming of pump oil vapors in vacuum systems," *Vakuum in Forschung und Praxis*, vol. 33, no. 1, pp. 28–32, 2021.

[129] D. Das and V. Natarajan, "Absolute frequency measurement of the lithium d lines: Precise determination of isotope shifts and fine-structure intervals," *Phys. Rev. A*, vol. 75, p. 052508, May 2007.