Thesis advisor   Author

Professor Markus Greiner   **Widagdo Setiawan**

Fermi Gas Microscope

Abstract

Recent advances in using microscopes in ultracold atom experiment have allowed experimenters for the first time to directly observe and manipulate individual atoms in individual lattice sites. This technique enhances our capability to simulate strongly correlated systems such as Mott insulator and high temperature superconductivity. Currently, all ultracold atom experiments with high resolution imaging capability use bosonic atoms. In this thesis, I present our progress towards creating the fermionic version of the microscope experiment which is more suitable for simulating real condensed matter systems.

Lithium is ideal due to the existence of both fermionic and bosonic isotopes, its light mass, which means faster experiment time scales that suppresses many sources of technical noise, and also due to the existence of a broad Feshbach resonance, which can be used to tune the inter-particle interaction strength over a wide range from attractive, non-interacting, and repulsive interactions. A high numerical aperture objective will be used to image and manipulate the atoms with single lattice site resolution. This setup should allow us to implement the Hubbard hamiltonian which could describe interesting quantum phases such as antiferromagnetism, d-wave superfluidity, and high temperature superconductivity. I will also discuss the feasibility of the Raman sideband cooling method for cooling the atoms during the imaging process.

We have also developed a new electronic control system to control the sequence of the experiment. This electronic system is very scalable in order to keep up with the increas-

ing complexity of atomic physics experiments. Furthermore, the system is also designed to be more precise in order to keep up with the faster time scale of lithium experiment.

Contents

Acknowledgments

Physics had been my one and only passion for a long time. I came to Harvard many years ago to learn more about the one thing that I know best: physics. Then came the unexpected. During my years in graduate school, I found my new passion: engineering. I could have never imagined the amount of engineering knowledge I would acquire in physics graduate school. Graduate school has been a very rewarding experience for me. For this, I am really grateful for my supervisor Markus Greiner, for allowing me to explore and improve my engineering skills openly in the lab.

First, I want to thank Tobias Schuster, the first diploma student in our lab for building many of the early infrastructure for the lab. Also to Jonathan Mailoa who designed and implemented the early prototypes for our control electronics and software which becomes the foundation of the Spectron electronic control system. To Kate Wooley-Brown for building a variety of equipment in the lab over the years. Special thanks to Florian Huber who shares or even exceeds my passion in engineering, for continuously challenging me to improve every component we design. Without him, the Center for Cool Technologies could not have existed. I also want to thank Dylan Cotta for building a large number of electronics, mechanical, and optical components in the lab in a ridiculously short amount of time. He clearly would have been a perfect addition to our experiment. One of my only regrets in graduate school is my failure to make him stay in our lab as a full time graduate student.

The apparatus that we have been building is now in the capable hands of the next generation graduate student Maxwell Parsons, and postdoc Sebastian Blatt. Although I was initially a little bit worried about the continuation of the experiment, it turns out that my worry was totally unwarranted since they are now fully integrated with the experiment. I fully expect them to do great work with this experiment over the next few years.

I also want to thank the members of the Rubidium Awesome Quantum Gas Mi-

croscope experiment: Jonathon Gillen, Amy Peng, Waseem Bakr, Simon Foelling, Eric Tai, Alex Ma, Jonathan Simon, and Phillip Preiss. There is no doubt that the directions of our experiment has been greatly influenced by their experience in running the Rubidium experiment.

It's been Legendary.

Dedicated to myself.

Yes, seriously.

Chapter 1

Introduction and Summary

Classical Computation / Simulation

The Church-Turing thesis states that everything algorithmically computable is computable by a Turing machine[1]. This conjecture also encapsulates quantum mechanical systems since quantum theory is also algorithmically computable[2]. Therefore, in principle, we could simulate any quantum mechanical system with any Turing machine including, for example a computer.

Unfortunately, in quantum mechanics if we want to simulate the full wavefunction, the number of states that have to be simulated grows exponentially with the number of particles in the system. This can easily be seen by considering a two-spin system. Consider the case where we have two distinguishable spin one-half particles, each fixed in space. The unnormalized wave function of each individual particle can be written as (neglecting the overall phase)

$$|\Psi_1\rangle = a_1 |\downarrow\rangle + |\uparrow\rangle$$

$$|\Psi_2\rangle = a_2 |\downarrow\rangle + |\uparrow\rangle,$$

where a_1 and a_2 are complex number parameters that determine the state of each particle. Note that I do not put any coefficient on the spin up state since that coefficient is constrained by the normalization condition. The general wavefunction of both particles combined is

$$|\Psi_{12}\rangle = b_1 |\uparrow\uparrow\rangle + b_2 |\uparrow\downarrow\rangle + b_3 |\downarrow\uparrow\rangle + |\downarrow\downarrow\rangle$$

Again, the last term is constrained by the normalization condition. Here, the number of free parameters increases non-linearly with the number particles. In general, we need $2^N - 1$ complex numbers to describe the spin state of N distinguishable particles. This is in contrast with a classical system where the number of parameters increases only linearly with the number of particles.

This exponential growth practically limits the size of the quantum mechanical system that can be simulated for a given computational resource. Note that this limitation is really severe, even with the best possible computer. Assuming a classical particle can store a single complex number, and we use every single particle in the universe as the storage for our computer or simulator, and assuming the universe has 10^{100} particles, the number of quantum spins that can be fully simulated is 332. In practice, our computer is far smaller than the size of the universe, so fully simulating even 50 spin one-half particles is already impossible.

Note that the analysis above only applies for the case where the wavefunction is highly entangled. If we can write the wavefunction as a product of subsystem wavefunctions, the exponential growth in the number of states can be reduced or even eliminated[3, 4]. Furthermore, the simulation can be simplified if we only need to compute some particular observable from the system, such as the average spin direction in non-entangled system[5].

Strongly Correlated Systems

In strongly correlated systems, we cannot write the multi-particle wavefunction as a product of single particle wavefunctions[6, 7]. The origin of the correlation can be as simple as a Coulomb repulsion between the particles. Here, Coulomb repulsion prevents us from finding two electrons too close to each other due to the high energy cost of such a configuration. This implies that the position of one electron depends on the position of all other electrons in the system.

In solid-state materials, strongly correlated electron systems can usually be found in materials with open d and f electrons in transition metals such as iron, vanadium, and Lanthanum, for example [8]. One interesting property of some strongly correlated systems with important potential applications is the high temperature superconductivity found in Cuprate transition metal compounds such as LaBaCuO, LSCO, YBCO, BSCCO, and TBCCO[9, 10, 11].
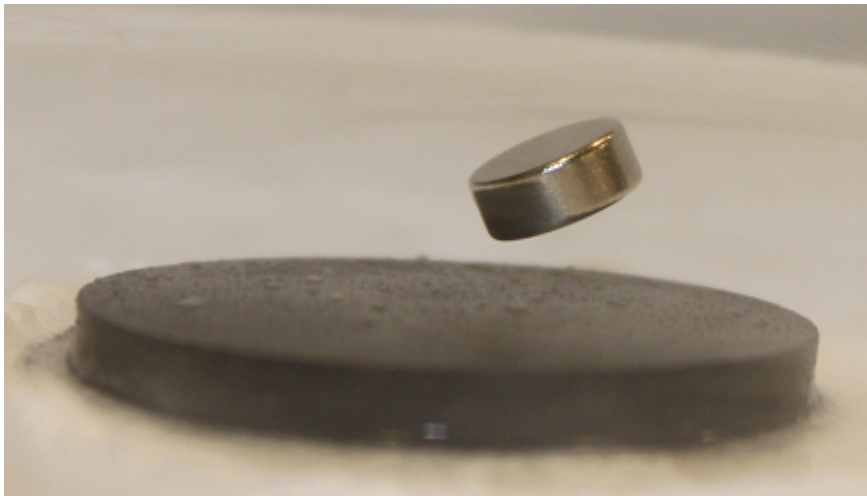
Figure 1.1: A magnet floating above a high temperature superconductor sample.

High Temperature Superconductor One open question in solid-state physics is high temperature superconductivity, first observed in 1986. These superconductors are qualitatively different from typical low temperature superconductors which can be described by the BCS theory[12]. Due to the strongly correlated nature of these high temperature superconductors, the theory of their superconductivity is still not well understood[13].
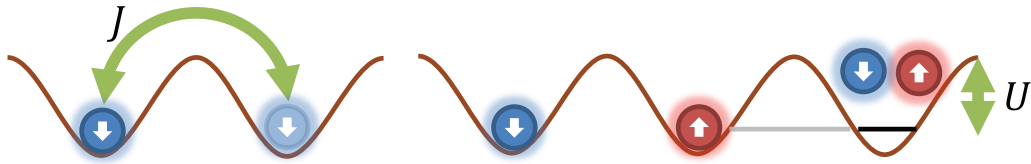
Figure 1.2: Illustration of Hubbard model. The model only contains two terms: J, which is the tunneling energy to a nearest neighbour lattice site, and U, which is the onsite interaction energy between two particles.

Hubbard Model One simple model that has the possibility of capturing this superconducting behavior is the Hubbard model. The Hubbard model was first proposed in 1963 as an improvement to the tight-binding model[14]. The model only includes on-site interaction energy terms and nearest neighbor tunneling terms as illustrated in Figure 1.2. Despite its simplicity, the Hubbard model captures complex physical processes, including potentially high temperature superconductivity[13]. Unfortunately, the Hubbard model is also very difficult to solve. The ground state of the two-dimensional Hubbard model is still not well understood. Understanding the Hubbard model, both theoretically and experimentally, is one of the most active field of research over the past few years[15, 16, 17].

Quantum Simulation

As mentioned before, simulating highly correlated systems such as near-degenerate Fermi gasses or the Hubbard Hamiltonian using classical systems is difficult. This is where quantum simulation comes in. Richard Feynman proposed the idea of quantum simulation in 1982[18]. Instead of using a classical system to simulate a quantum system, Feynman proposed to simulate a quantum system with another, ideally simpler and more accessible, quantum system. Such a device is called a quantum simulator. A quantum simulator would be able to simulate some classes of quantum systems efficiently without using exponentially large numbers of particles[19].

This does not necessarily mean that quantum simulation is exponentially more powerful than classical simulation. The problem is that we cannot probe the wavefunction of a quantum system due to the Heisenberg uncertainty principle. The best that we can do is to project the wavefunction into some basis and measure the projected value on that basis. In contrast, a full classical simulation of a quantum system would give us the full wavefunction of that simulated quantum system. Nevertheless, quantum simulation is still very useful since we can, in principle, choose the basis of our measurement or even repeat the same experiment with measurements in different bases to learn about the property of the quantum system in question.

Note that quantum simulation is qualitatively distinct from quantum computation. Quantum computation deals with qubits which are the quantum version of the classical digital bits, while quantum simulation deals with simulating quantum system without transforming the system into the digital domain.

Ultracold Atoms as Quantum Simulator
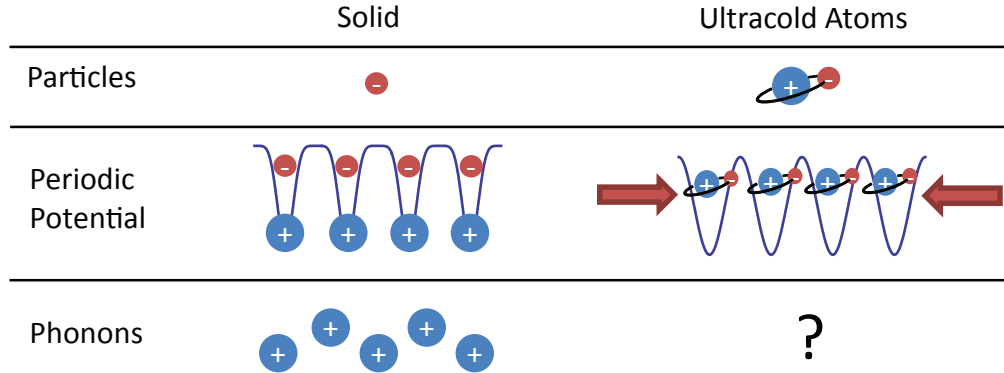
Figure 1.3:   Quantum simulation of solid-state material with ultracold atoms.

Recent progress in atomic physics has opened up the possibility of quantum simulation[25, 26, 27, 28, 29, 30, 31, 32, 33, 34] using ultracold atoms. In particular, we are interested in simulating solid-state materials using ultracold atoms in an optical lattice. The typical mapping between solid-state system and ultracold atom system is shown on Figure 1.3 and Table 1.1.

Due to the differences between solid-state system and ultracold atom system, ultracold atom experiments have several advantages over the corresponding solid-state system:

- The lattice in ultracold atom experiments is generated by interfering laser beams[35]. The lattice structure is typically free of defects as long as the laser beam profile generating this lattice is clean. In solid-state experiments, these defects usually contribute to decoherence in the quantum state. However, some physical effects, such as Anderson localization, greatly depend on the impurities and defects in the system[36]. In ultracold atom experiments, these defects and impurities can be simulated by projecting disordered light beams through a high NA lens[22].

Property	Solid	Ultracold Atoms
Particles	Electrons	Neutral atoms
Periodic Potential	Atomic nuclei crystal structure	Optical lattice
Phonons	Nuclei displacements	Usually not simulated
Inter-particle interactions	Coulomb Phonon mediated	Van der Waals
Magnetic field	Magnetic Field	Rotating optical lattice Rotating atomic sample Raman coupling
Momentum detection	X-Ray, neutron, and electron diffractions	Time of flight images Laser beam diffraction[20]
Spatial detection	Scanning tunneling microscopy Atomic force microscopy Transmission electron microscopy	Direct optical imaging Electron microscopy[21]
Defects	Impurities Crystal defects	Noise in optical lattice beams Added optically on purpose[22, 23]
Filling factor	Fixed by charged neutrality Tuned by doping	Add or remove atoms[24] Change overall confinement

Table 1.1: Analogy between solid-state material and their ultracold atom quantum simulations.

- There is no phonon in typical optical lattice experiments since the lattice structure is generated using interfering laser beams. The back coupling of the atoms on the lattice beams is usually negligible, with some exceptions[37, 38]. In some cases, this small back coupling is very beneficial since phonons usually contribute to unwanted decoherence in the system. However, some physics such as BCS superconductivity, heavily relies on the phonons[14], making low-level quantum simulation of BCS physics in atomic physics experiments more difficult.

- The tunneling rate can be easily tuned in real-time during an experiment[39]. The tunneling rate of an atom (or an electron in solid-state systems) to a different lattice site is determined by the height and spacing of the periodic potential. In solid-state systems, this periodic potential is usually fixed by the crystal structure of the material. In ultracold atom systems however, we can change the height of the periodic potential by changing the laser power generating the periodic potential, changing the tunneling rate in realtime.

- The on-site interaction energy can be tuned in real-time during an experiment using Feshbach resonance[40].

- Highly coherent quantum dynamics. Since there is no thermal in typical ultracold atoms experiment, the quantum dynamics of the system is highly coherent.

- The length scale of the lattice is on the order of the optical wavelength, which is far larger than the typical interatomic spacing in solid-state materials. This allows for the direct observation of occupation number on a single lattice site level[41, 42, 43]. This capability is the focus of this thesis.

- The time scale of the dynamics of the atoms in the optical lattice is slow (10^{-3} s)

compared to the time scale in solid-state materials $(10^{-12}\,\mathrm{s})$[44]. This allows us to observe the dynamics of the system directly in the experiment.

- The filling factor is easily tunable by changing the number of atoms which could be both an advantage and a disadvantage. I will elaborate more on this in the next section.
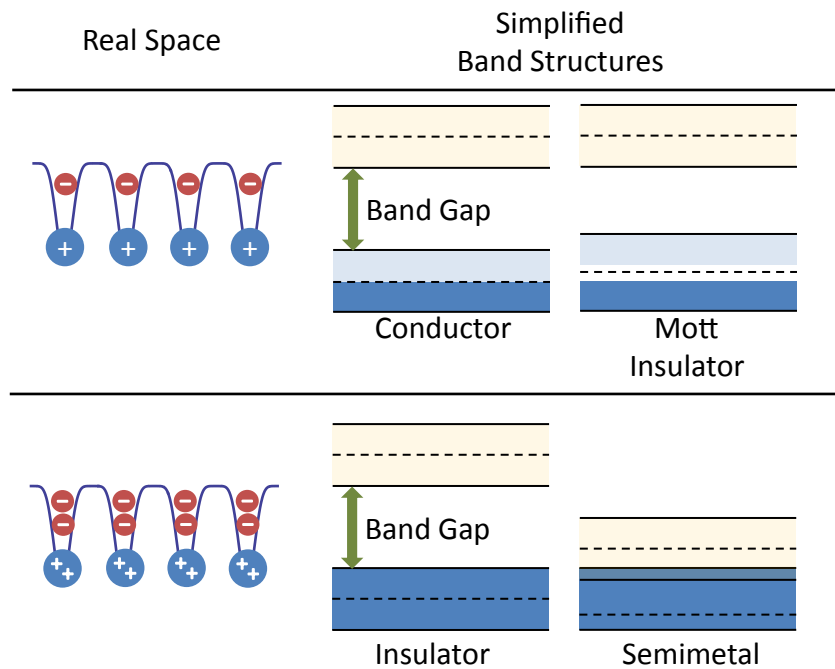
Figure 1.4: The filling factor in a pure solid-state system determines the macroscopic behaviour of the material. A half-filled band material is usually a conductor with the exception of Mott-insulating materials where the interaction energy between the electrons creates an effective energy gap. A fully-filled band material is usually an insulator unless the next band overlaps with the valence band.

Filling Factor In a solid-state material, the particles are electrons while the periodic potential is generated by the crystal structure of the nuclei. An electron has two possible spin states, spin up and spin down. The total number of states in a given band is then given by $2N$ where N is the number of lattice sites. Since both the number of lattice sites

and the number of electrons are determined by the number of nuclei (due to charge neutrality), the number of electrons per lattice site is solely determined by the number of electrons per unit cell in the lattice as illustrated in Figure 1.4.

For an odd number of electrons, the valence band is generally half-filled since each band can take an even number of electrons per lattice site. Here, we expect the solid to be in a conducting state[1]. On the other hand, if there are an even number of electrons per atom, the valence band is fully-filled. If the next band is not overlapping with the valence band, we expect the solid to be an insulator, and the insulator is called a band insulator.
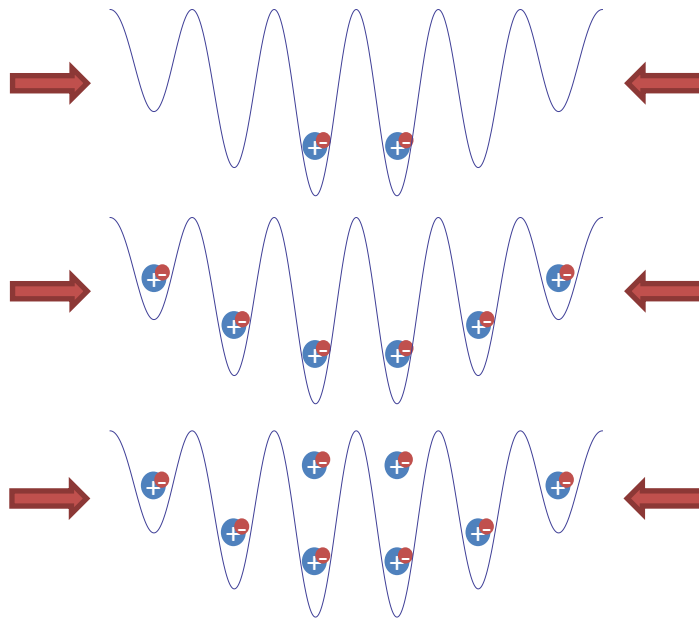
Figure 1.5: In ultracold atoms experiments, we can change the filling factor by simply changing the total number of particles in the system.

In a solid-state material, the number of electrons can be changed from either full-filling or half-filling by doping the material with other atoms. In optical lattice experiments, the filling factor can be changed by changing the number of atoms as illustrated in Figure 1.5.

---

[1]There is some exception to this rule, for example, the Mott insulating materials[45] where materials with half filled valence band still acts like an insulator due to strong correlations between the electrons.

This filling factor degree of freedom is both a good thing and a bad thing depending on the situation. In some cases, such as reaching the Dirac point in graphene[46], it is important that the lower band is fully filled and the upper band is completely empty. This is very natural in pure graphene since the number of electrons is locked to the number of lattice sites by charge neutrality. In atomic physics experiments however, we have to control the atom number with enough precision to reach this Dirac point[24].

Simulating the Hubbard Model   In ultracold atom experiments, we can easily tune the lattice potential and interaction energy. This makes ultracold atoms in an optical lattice a good quantum system to simulate the Hubbard model[47, 48, 49, 50]. We can achieve the Hubbard regime by making sure that the tunneling is weak enough that we can ignore tunneling other than nearest neighbor, while making sure that the band gap is large enough that only the first band is filled with atoms. One of the long term goals of our experiment is to simulate high temperature superconductivity using ultracold atoms.

Lithium for Quantum Simulation

Out of all possible atoms for ultracold atom quantum simulation, lithium has several advantages:

- Lithium is an alkali atom. Alkali atoms have one valence electron which greatly reduces the number of states that have to be controlled in an experiment. This simplifies the laser cooling and imaging step considerably compared to other non-alkali atoms[2].

---

[2]Another fermionic Alkali atom is the Potassium atom.

- Lithium has both fermionic (lithium-6) and bosonic (lithium-7) isotopes, which allows us to simulate both the Bose-Hubbard and Fermi-Hubbard Hamiltonian[3] using almost identical setups.

- Lithium has a relatively light mass compared to other alkali atoms which allows for faster experiments.

- Lithium-6 has a broad Feshbach resonance that can be used to tune the on-site interaction energy in the Hubbard model[51].

Faster Time Scale Even though the slower time scale of atoms in optical lattices compared to the solid-state systems is one of the primary reasons why we are doing quantum simulations, we do not want the time scale to be too slow. The typical lifetime of an atom in an ultra-high vacuum chamber is in the order of 1 minute due to background gas collisions[52], so we want to perform the physics experiments in a much shorter time. Lithium, with its light mass, has the clear advantage of faster time scales for an atom. The natural energy scale of atoms in an optical lattice is the recoil energy $E_r = \hbar^2 k^2 / 2m$[39], where k is the lattice wavevector. Plugging in the mass for both Rubidium (87) and lithium(6), we get that lithium has an advantage of $87/6 = 14.5$ compared to Rubidium in time scale, for the same lattice wavevector.

Feshbach Resonance Another advantage of lithium is the existence of a low-loss Feshbach resonance that can be used to tune the scattering length between the atoms[53]. The idea is to use a magnetic field to change the relative potential between atomic state and molecular state in and out of resonance. The scattering length can be parameterized

---

[3]Fermi-Hubbard model is usually called Hubbard model since the Hubbard model was originally made for electrons and electrons are fermions.

as a function of magnetic field (B) as[52]:

$$
\begin{aligned}
a(B) &= a_{bg}\left(1 - \frac{\Delta B}{B - B_0}\right), \\
a_{bg} &= 1405\, a_0 \\
\Delta B &= 300\,\text{Gauss}
\end{aligned}
$$

where a_{bg} is the background scattering length, B_0 is the resonance magnetic field, and ΔB is the characteristic width of the resonance. Using this Feshbach resonance, we can tune the interaction such that the on-site interaction energy (U) in a lattice site can be as high as $10\,\text{kHz}$ while typical Rudibium experiment has typical U of around $500\,\text{Hz}$.

Control System

The complexity of atomic physics experiments has been increasing considerably over the years. The number of steps required to perform a typical experiment has also increased substantially. Furthermore, the faster time scale of lithium in optical lattice experiments demands a more precise experimental control system. We decided from early on to focus some of our efforts in engineering a next generation control system that is a little bit more up to date with current technological development[4]. If there is one thing that differentiates our experiment to other BEC type experiments, it is the control system. I will dedicate a significant part of this thesis to describing the improvements that we made with this new generation control system, in hopes that other groups could use our system as a starting point for designing their own system.

---

[4]In reality, our control system is still around 10 years behind current technology. Keeping up with the state of the art electronics is both totally unnecessary and prohibitively expensive for our experiments.

Overview

This thesis is organized as the follows:

- The second chapter is an overview of the basic theoretical framework for optical lattice experiments with fermionic lithium atoms.

- The third chapter describes Spectron, our electronic control system, in detail.

- The fourth chapter explains the experimental setup and the progress that we have made.

- The fifth chapter gives an outlook for possible future directions of the experiment.

- There is an appendix on basic electronics background knowledge that might be relevant for some atomic physics experimentalists.

Chapter 2

Imaging Lithium Atoms in an Optical Lattice

One defining feature of this experiment is the capability to image fermionic lithium atoms in an optical lattice with single site resolution. This chapter describes some of the theory that is relevant for the experiment.

2.1 Lithium-6 Atom

Lithium has three electrons where two of the inner electron from a closed $n = 1$ inner shell, where n is the principle quantum number, while the third valence electron contributes to most of the low energy electronic structure[1]. The figure above shows the $n = 2$ energy levels of lithium[54]. At zero magnetic and electric field, the energy splitting can be divided into three energy scales as shown in Figure 2.1:

---

[1] The inner shell energy scales is in the order of $10\,\mathrm{eV}$
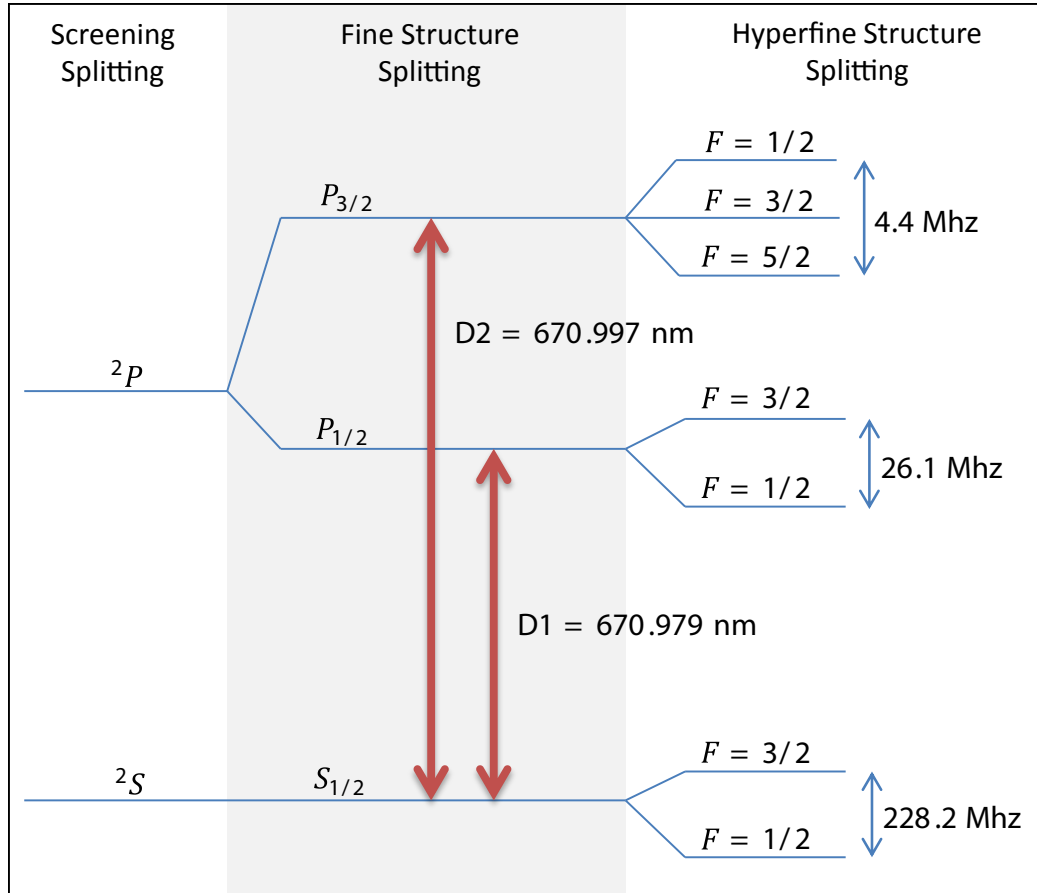
Figure 2.1: Energy structure of lithium atom in the absence of external magnetic and electric field

- Screening structure splitting. The ^{2}S to ^{2}P transition is caused by different effective inner electron screening due to different valence electron orbital shapes (S and P orbitals). This splitting provides the majority of the energy splitting for optical transition at 670.98 nm: MOT, Zeeman slower, imaging, and dipole trap transitions[55].

- Fine structure splitting. This is caused by the two possible relative orientations between the electron spin and electron orbital angular momentum direction[56]. This splitting only exists on the P-orbital and higher since the S-orbital has no

orbital electron angular momentum. In ^{6}Li, this splitting is only $10\,\mathrm{GHz}$ which is much smaller than other alkali atoms. For comparison, Na fine structure splitting is $515\,\mathrm{GHz}$[57], while Rb fine structure splitting is $7123\,\mathrm{GHz}$[58]. The small splitting of lithium severely limits the maximum detuning that can be achieved for state dependent light shift making state dependent potential lossy[35].

- Hyperfine structure splitting. This splitting is caused by the different possible relative orientations between the electron total angular momentum and the nuclear spin. For the D_2 line, the excited state hyperfine state splitting is smaller than the transition linewidth, making standard m_F optical molasses difficult to implement[55].

2.2 Optical Dipole Potential

One very useful tool for manipulating atoms is the optical dipole potential[35]. When far off-resonant light interacts with atoms, we can treat the light as a classical oscillating electric field. At the same time, we treat the atom as an oscillator with natural frequency ω_0 and damping coefficient Γ[59]. This single oscillator approximation is usually good for computing the energy shift in the ground state. The oscillating electric field causes an AC Stark shift on the atom which is proportional to the light intensity[60]. If the light intensity varies as a function of position, the atoms effectively feel a varying potential proportional to this Stark shift. The potential caused by the light on an atom in the ground state is given by

$$U(r) = -\frac{3\pi c^2}{2\omega_0^3} \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega} \right) I(r),$$

where:

- $\omega_0 = 446.799\,677\,\text{THz}$ is the atomic transition angular frequency for ^6Li

- $\Gamma = 2\pi \times 5.872\,4\,\text{MHz}$ is the transition linewidth for ^6Li

- $c = 299,792,458\,{}^m\!/\text{s}$ is the speed of light

- ω is the applied light angular frequency

- $I(r)$ is the light intensity as a function of position

When ω_0 and ω are very close, we can ignore the $\frac{\Gamma}{\omega_0 + \omega}$ term. This is called the rotating wave approximation. For lithium however, we usually use far detuned light so this rotating wave approximation is not appropriate. For 1064 nm light, this correction is 23%.

The light also causes some spontaneous scattering with rate given by:

$$\Gamma_{sc}(r) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\omega}{\omega_0}\right)^3 \left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega}\right)^2 I(r)$$

For a red detuned light ($\omega < \omega_0$) and for a given Stark shift, it is generally beneficial to detune the light further in order to minimize the spontaneous scattering. When ω_0 and ω are very close, we can take approximation that $\left(\frac{\omega}{\omega_0}\right)^3 \approx 1$. However, for 1064 nm light in lithium, this approximation is not appropriate since $\left(\frac{\omega}{\omega_0}\right)^3 \approx 0.25$. The scattering rate in the typical experiment lattice of $V = 10E_r$ is around $\Gamma_{sc} = 0.02\,/\text{s}$.

2.2.1 Gaussian Beam

Since we are generating the light using lasers, the natural intensity profile of the light is Gaussian or very close to Gaussian. A Gaussian beam with waist w_0 at the origin going in the positive $\hat{z}$ direction has complex electric field profile:

$$E(r, z) \;\;=\;\; E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w^2(z)}\right) \exp\left(-ikz - ik\frac{r^2}{2R(z)} + i\zeta(z)\right)$$

$$
\begin{aligned}
w(z) &= w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \\
z_R &= \frac{\pi w_0^2}{\lambda} \\
\zeta(z) &= \arctan\left(\frac{z}{z_R}\right) \\
R(z) &= z\left[1 + (\frac{z_R}{z})^2\right] \\
E_0 &= E(0,0) = \sqrt{2\sqrt{\frac{\mu_0}{\epsilon_0}} I_0},
\end{aligned}
$$

where E_0 is the maximum complex electric field, r is the perpendicular distance to the axis of the beam propagation, z_R is the Rayleigh range, $R(z)$ is the radius of curvature which cause the radial phase shift, and $\zeta(z)$ is the Gouy phase shift which is the axial phase shift. In practice, the atoms are usually located at $r = 0$ and $z = 0$, so usually the radial phase shift and the Gouy phase shift can be ignored. The time averaged intensity is given by:

$$
I(r,z) = I_0 \left(\frac{w_0}{w(z)}\right)^2 \exp\left(-\frac{2r^2}{w^2(z)}\right)
$$

2.2.2 Optical Lattice
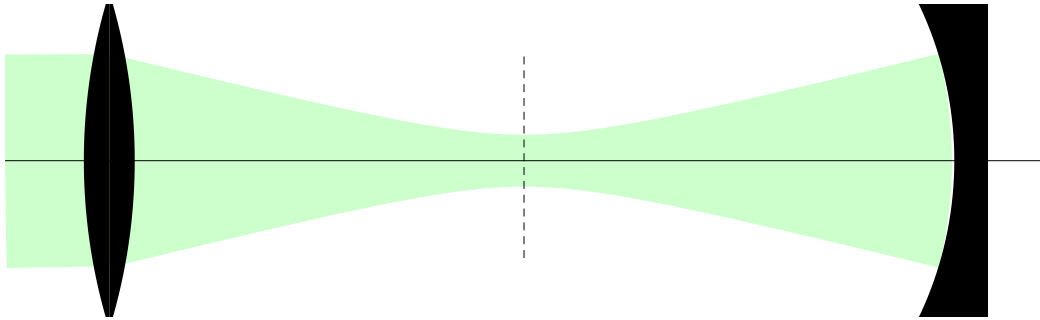
Figure 2.2: Standard retroreflected lattice configuration.

One simple way to generate an optical lattice potential is to shine counter-propagating laser beams at the location of the atom as shown in Figure 2.2. The electric fields from

the two beams add up to create an oscillating standing wave. The atom then feels the time-averaged intensity profile of this standing wave.

2.2.3   Multi-Level Optical Dipole Potential
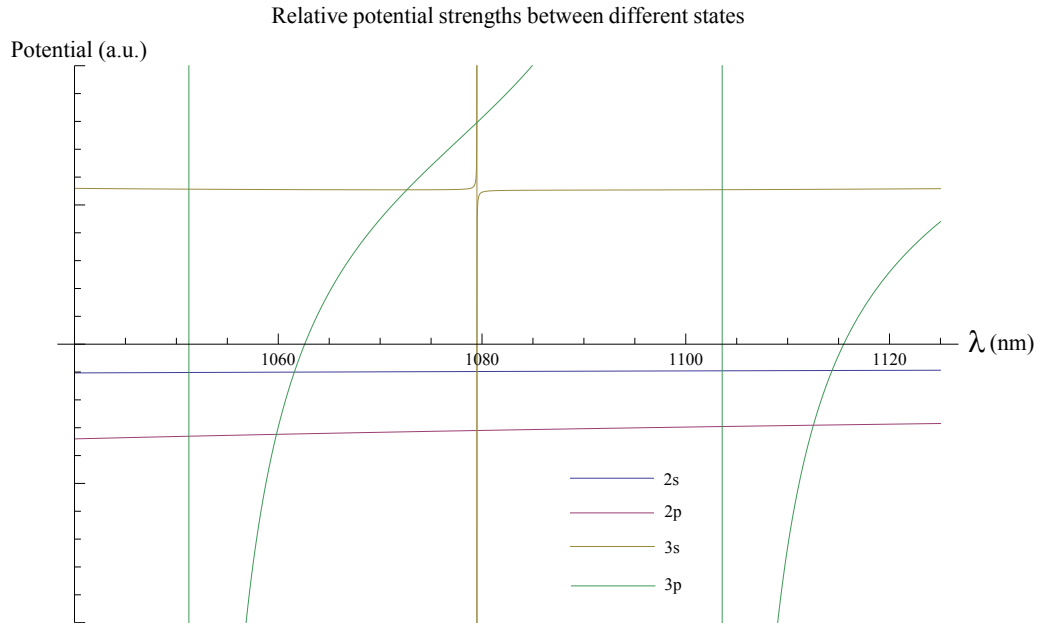
Figure 2.3:   AC Stark shift on lithium atoms in different quantum states. During the experiment, the atoms spend most of their time in the ^{2}S ground state.

In real atoms, there are multiple transitions that have to be considered for calculating the energy shift, especially for excited state energy shifts. The energy shift as a function of wavelength for ^{6}Li is shown in Figure 2.3 where the data for the figure is taken from [61].

2.2.4   Classical Stochastic Heating in Harmonic Potential

Consider a typical case when we use $1064\,\mathrm{nm}$ light to create potential on the atoms. Here, the potential shift greatly depends on the state of the atom. The ^{2}S and ^{2}P states are trapped, while the ^{3}S and ^{3}P states are anti-trapped (being pushed away by the light) as

shown in Figure 2.3. Since the light is very far detuned from any optical transition, the atom spend most of its time in the ^{2}S state. However, whenever there is an inelastic scattering event, from for example, a resonant imaging beam, the atom momentarily goes to the ^{2}P state[62]. Here the atom will feel a much larger lattice potential. I would like to calculate the effect of this time-varying potential on an atom.

Consider an atom in a classical harmonic oscillator with angular frequency ω_1. Let us set the initial energy of this atom to be E_1. At random time, the atom is excited to an excited state where the atom feel a different harmonic potential with angular frequency ω_2. This transition changes the energy of the atom to E_2. After a random time, the atom decays back to the first harmonic oscillator with angular frequency ω_1 with new energy E_3. This non-adiabatic switching back and forth between two harmonic potentials heats the atom. We can do a simple classical calculation to estimate this heating rate.

The fractional energy change over one such cycle is given by

$$\frac{E_3}{E_1} = \frac{\left(\omega_1^2 + \omega_2^2\right)^2 + \left(\omega_2^4 - \omega_1^4\right)\cos\left(2\phi_1\right) + 4\omega_1^2\left(\omega_2^2 - \omega_1^2\right)\sin^2\left(\phi_1\right)\sin\left(2\phi_2\right)}{4\omega_1^2\omega_2^2},$$

where ϕ_1 is the amount of phase accumulated before the first excitation and ϕ_2 is the amount of phase accumulated before the atom decays back to the ground state. If we integrate over the phase ϕ_1 and ϕ_2, we get:

$$\frac{E_3}{E_1} = \frac{\left(\omega_1^2 + \omega_2^2\right)^2}{4\omega_1^2\omega_2^2}$$

In ^{6}Li, the relative strength between the excited state ^{2}P trap and ground state ^{2}S trap at 1064 nm is 3.206. which implies frequency ratio of $\omega_3/\omega_1 = \sqrt{3.206} = 1.79$. Plugging this to the equation above we get $E_3/E_1 = 1.38$, . This means that every full cycle, the average energy of the atom increases by 38%. Note that the energy grows exponentially. The

energy over time is given by

$$
\begin{aligned}
E(t) &= E(t=0)\exp\left(\frac{t}{\tau}\right) \\
\tau &= \frac{t_{\text{cycle}}}{\ln\left(\frac{E_3}{E_1}\right)},
\end{aligned}
$$

where t_{cycle} is the time it takes to go one full cycle (ground state $\rightarrow$ excited state $\rightarrow$ ground state) which is inversely proportional to the inelastic scattering rate.

2.3   Periodic Potential

A common way to analyze the behavior of atoms in an optical lattice is to approximate the potential as periodic. The eigenfunctions of a periodic potential can be described using Bloch wavefunctions[44, 63]

$$
\begin{aligned}
\psi_q^n(x) &= e^{i\frac{qx}{\hbar}}u_q^n(x) \\
H\psi_q^n &= E_q^n\psi_q^n(x),
\end{aligned}
$$

Where q is the quasi-momentum which goes from $-\hbar k$ to $+\hbar k$, $u_q^n(x)$ is a function with the same periodicity as the potential, and n is the band index. E_q^n is the energy of a particle in band n with quasi-momentum q. This Bloch wavefunction is maximally delocalized across all lattice sites for a given band.

If the periodic potential is strong, the potential tends to localize the wavefunction to a single lattice site. Although the Bloch wavefunctions are still the eigenstates, using these maximally delocalized Bloch wavefunctions to describe the localized wavefunctions is not intuitive. In this case, a new (non-eigenstate) basis called the Wannier function is more appropriate. Intuitively, the Wannier wavefunction is the maximally localized

wavefunction for a given energy band which is constructed by summing up all of the Bloch wavefunctions at that band[2]. The Wannier functions for a particle located at site x_i is given by

$$w_n(x - x_i) = \sum_q e^{-i\frac{qx_i}{\hbar}} \psi_q^n(x) \tag{2.1}$$

2.3.1 Periodic Potential from Optical Lattice

In typical optical lattice experiments, the periodic potential is constructed from two interfering laser beams. The form of the period potential is

$$V(x) = -V_0 \cos^2(kx),$$

where V_0 is the amplitude of the periodic potential and $k = \frac{2\pi}{2 \text{ Lattice Spacing}}$ is the lattice wavenumber. In order to make the Hamiltonian dimensionless, it is natural to express V_0 in the unit of the lattice recoil energy

$$E_r = \frac{\hbar^2 k^2}{2m},$$

where m is the mass of the particle and k is the lattice wavenumber. Both k and E_r have intuitive meaning when the laser beams generating the lattice are counter-propagating. In this case, k is the wavenumber of the laser beam and E_r is the energy that a particle at rest gains when emitting a photon with wavenumber k. When the laser beams are not counter-propagating, E_r is not related to the recoil energy from the photons generating the optical lattice, while k is not related to the wavenumber of the laser beam.

Hubbard Model The Hubbard model assumes that the particles can only tunnel to nearest neighbour lattice sites[14]. Furthermore, the system is restricted to a single band.

---

[2]Without this band constraint, the maximally localized wavefunction is of course the delta function

Since we are trying to simulate electrons, which are spin-half particles, we also restrict our atoms to only two spin states. The Hubbard model is given by

$$H = -J \sum_{i,j} \left(c_{i,\uparrow}^\dagger c_{j,\uparrow} + c_{i,\downarrow}^\dagger c_{j,\downarrow} + \text{h.c.} \right) + U \sum_i n_{i,\uparrow} n_{i,\downarrow}$$

The J term represents the energy cost to move a particle from one lattice site to its neighboring lattice site. The sum over i, j is constrained to nearest neighbour lattice sites. The U term shows the interaction energy cost of putting a spin-up and spin-down particle in the same lattice site. Note that since we are using single-band model, Pauli exclusion forbids us to put two spin-up (or two spin-down) particles in the same lattice site. This constraint is enforced by the commutation relation of the fermionic creation and annihilation operators, c and $c^\dagger$. $n = c^\dagger c$ is the number operator. The Hubbard model is a single-band model. Therefore, the model is only valid when J and U are much smaller than the band energy spacing.

On-Site Interaction The onsite interaction U can be approximated as[14]:

$$U = g \int \mathrm{d}^3 r \, |w(r)|^4 ,$$

where $w(r)$ is the Wannier function of the atom in a single lattice site. The coupling constant, g, for s-wave scattering is given by $g = \frac{4\pi a}{m}$, which can be tuned using the Feshbach resonance. Another method of increasing U is by increasing the confinement since the U scales as w^4. One approximation to simplify the calculation of U is to replace the Wannier wavefunction with a Gaussian wavefunction corresponding to the harmonic term in the lattice potential trap. This is valid for deep trap where the ground state potential is mostly harmonic.

2.3.2   Tunneling Rate

$$H \;=\; -J \left[\begin{array}{c} \end{array} \right]$$
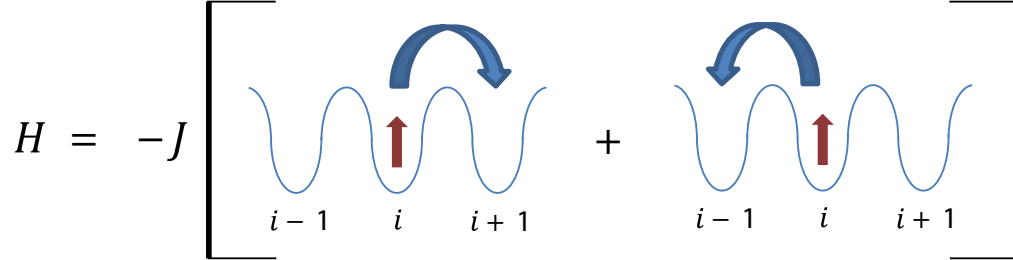
Figure 2.4:   Tunneling rate.

To calculate the tunneling, let us consider a Hamiltonian where $U = 0$ with only one spin state. Let us also assume that the system is one dimensional. The Hamiltonian, visualized in Figure 2.4, can be written as:

$$H = -J \sum_i c_{i+1}^\dagger c_i + c_{i-1}^\dagger c_i$$

We can write down a general wavefunction as a function of time as:

$$|\psi(t)\rangle = \sum_j a_j(t) \, |j\rangle \, ,$$

where $|j\rangle$ is the Wannier function of the particle on this particular band at site j.

Plugging in into the time-dependent Schrodinger equation we get:

$$
\begin{aligned}
i\hbar \frac{\mathrm{d}\,|\psi(t)\rangle}{\mathrm{d}t} &= H \, |\psi\rangle \\
&= i\hbar \sum \frac{\mathrm{d}a_j(t) \, |j\rangle}{\mathrm{d}t} \\
&= -J \sum a_j(t) \, |j+1\rangle + a_j(t) \, |j-1\rangle \\
&= -J \sum a_{j-1}(t) \, |j\rangle + a_{j+1}(t) \, |j\rangle \\
i\hbar \dot{a}_j(t) &= -J \left(a_{j-1}(t) + a_{j+1}(t) \right)
\end{aligned}
$$

A solution for this equation is

$$a_j(t) \;=\; e^{-i\omega t} e^{ij\alpha}$$

$$\hbar\omega \;=\; -J\left(e^{-i\alpha} + e^{i\alpha}\right)$$

$$=\; -2J\cos\alpha$$

$$E \;=\; -2J\cos\alpha$$

$$J \;=\; \frac{1}{2}\left(E_{\max} - E_{\min}\right)$$

This is a general result for tight binding model that says that the nearest-neighbor tunneling rate in any band is proportional to the width of that band. This tunneling rate calculation method is useful since we can calculate the tunneling rate without computing the Wannier wavefunction, simplifying the calculation substantially.

2.3.3   Full-Filling (Band Insulator)
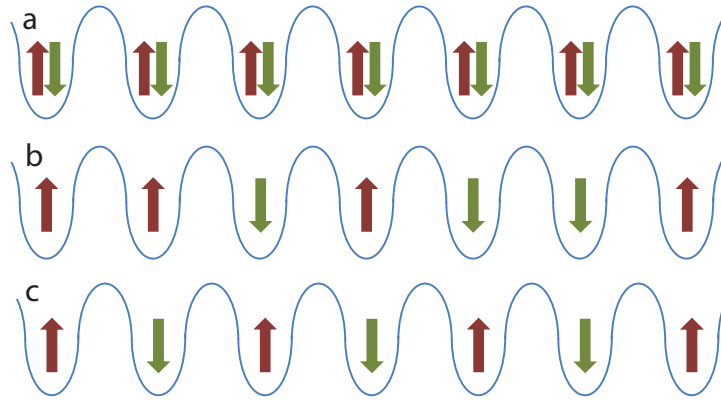
Figure 2.5:   a. Band insulator, b. Mott insulator, c. Anti-ferromagnetic ordering

Consider a single band model with two atoms per lattice site as shown in Figure 2.5a. Let us further assume that there are equal numbers of atoms in spin-down and spin-up states. In this setup, each lattice site is filled with exactly two atoms, one with spin up and the other with spin down particle. This is independent of the interaction strength. Since all lattice sites are fully-filled, there is no way for the atoms to tunnel to different lattice sites.

We call this a band-insulating state.

2.3.4   Half-Filling

Consider the case when we have exactly one atom per lattice site. This configuration is called half-filling[14]. For repulsive interactions ($U > 0$), the two spin-states do not want to be in the same lattice sites due to the positive interaction energy cost. Consider a case where $U \gg J, T$, where T is the temperature. Here, each atom will occupy a single lattice site. This is called the Mott insulating state as shown in Figure 2.5b.

If we reduce the temperature further such that $T \gg J \gg U$, we can consider a process on the order of J/U to further reduce the kinetic energy by tunneling to neighboring lattice sites. It is advantageous for two atoms in neighbouring lattice sites to have different spin-states to lower their energy. The global configuration with the lowest energy is the anti-ferromagnetic configuration as shown in Figure 2.5c, which is a special case of the Mott insulating state.

2.4   Imaging and Cooling

2.4.1   Photon Collection

In principle, in order to optically image a single atom, we only need to collect a single photon from each atom. In practice, however, there are several processes that force us to increase the number of photons collected per atom. Some of these processes are:

- **Limited photon collection aperture** - The atom tends to emit radiation over 4π

solid angle, while our imaging numerical aperture, $NA = 0.8$, can only capture around 20% of the emitted photons.

- **Limited optical resolving power** - Even though each atom can be treated as a point source, our imaging setup projects that point source into an Airy disk. The probability density function of the photon location is given by the Airy disk[64]. The more the Airy disks from neighbouring lattice sites overlaps, the more photons we need to reliably determine the original location of the atom.

- **Other losses in the imaging setup** - The imaging path consists of multiple optical elements, each of which absorbs or reflects some of the photons, preventing those photons from reaching the camera.

- **Noise in the imaging system**. Electronic readout noise on the camera and background photon noise in the room add noise on the images.

In practice, we have to collect hundreds or thousands of photons for a given atom to reliably determine its location. There are several heating sources that have to be considered.

Fluorescence Imaging   A naive approach is to shine a resonant imaging beam on the atoms. However, the calculation in Section 2.2.4 shows that the atom will gain energy exponentially with the number of photons scattered. Once an atom gains enough energy to transfer to a high enough band, the atom can then tunnel to different lattice sites, limiting the maximum number of photons that can be collected while preserving imaging fidelity. To prevent the atoms from tunneling, we can increase the lattice depth. For a $V = 1800E_{\text{rLattice}}$ lattice, which has 18 non-tunneling bands, the number of photons an

atom can scatter before starting to tunnel is on the order of ten photons.

Elastic Scattering Imaging   Another approach is to reduce the scattering rate by using weaker beam or far detuned imaging light so that the majority of the scattered photons are elastic[62]. Although the inelastic scattering is greatly suppressed, every time an atom emits a photon, the atom still receives a recoil momentum kick in random direction, which heats the atom.

It is important to realize that the atom receives a random recoil-momentum kick, and not recoil-energy kick from each scattering. The recoil energy only has physical meaning when the atom is initially at rest. When the atom moves with some velocity, the real recoil energy is much higher than the recoil energy $E_r = \hbar^2 k^2 / 2m$. Here, k is the wavevector of the imaging light, not the lattice potential. Consider an atom with momentum P. Let us assume that the atom absorbs a photon traveling in the same direction as the atom's momentum direction. The atom then emits a single photon in random direction. The net momentum change, called the recoil momentum, is $P_r = \hbar k$. The net energy change is

$$
\begin{aligned}
\Delta E &= E_{\text{new}} - E_{\text{old}} \\
&= \frac{(P + P_r)^2}{2m} - \frac{(P)^2}{2m} \\
&\approx \frac{(PP_r)^2}{m} \\
&\neq E_r = \frac{P_r^2}{2m}
\end{aligned}
$$

We can estimate the expected momentum of the atom as a function of recoil events. Let us assume that the atom receives momentum kick from random directions. The momentum / energy of the atom after N recoils is then

$$
P = \sqrt{N} P_r
$$

$$
\begin{aligned}
E &= \frac{P^2}{2m} \\
&= N\frac{P_r^2}{2m} = NE_r
\end{aligned}
$$

With $1800E_{\text{rLattice}}$ lattice, the number of photons the atom can scatter without tunneling is around 500 photons.

Cooling During Imaging   In order to collect even more photons, we can cool the atoms during the imaging process to prevent the atoms from reaching higher Bloch bands. One candidate for cooling is the optical molasses cooling and imaging method[65]. Unfortunately, the lithium excited state hyperfine splitting is not resolved as shown in Figure 2.1, forcing us to use near-detuned light. This near detuned light produces more inelastic scattering compared to far detuned light, which in turns increases the stochastic heating process described in Section 2.2.4.

Another candidate for cooling method is the Raman sideband cooling method demonstrated in [66, 67]. Cooling during imaging should allow us to collect significantly higher number of photons per atom.

2.4.2   Non-Degenerate Raman Sideband Cooling

Raman sideband cooling could be used to simultaneously cool the atoms and to generate fluorescent photons for imaging purposes. Figure 2.6 illustrates the key concepts of Raman sideband cooling. There are two relevant spin states, $m = +1/2$ and $m = +3/2$. The goal is to put the atom into the vibrational ground state of the $m = +3/2$ spin state. The Raman beam is represented by the red lines. Here, two-photon resonant light (the Raman beam) couples the $|m = +1/2, v = n\rangle$ and the $|m = +3/2, v = n+1\rangle$. The
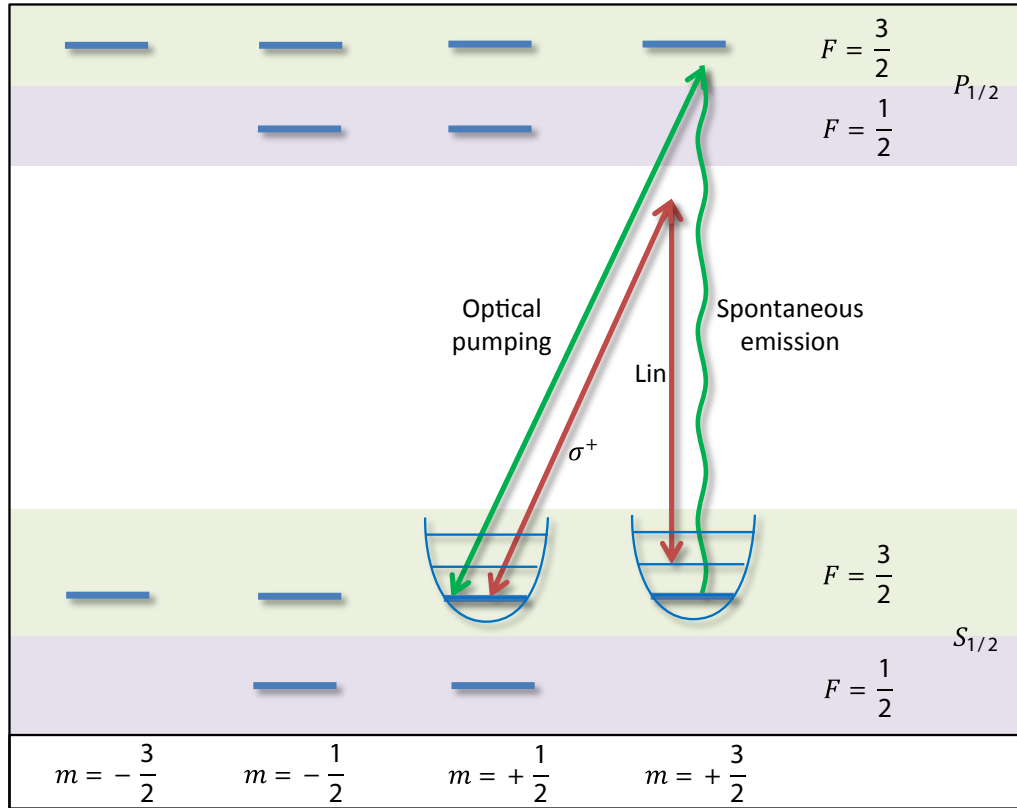
Figure 2.6: Non-degenerate Raman sideband cooling scheme.

frequency of the σ^+ light is larger than the frequency of the linearly polarized light by one vibrational energy level.

The atoms start at the $|m = +3/2, v = n + 1\rangle$ state. The atom then Rabi oscillates to the $|m = +1/2, v = n\rangle$ state. Note that the atom's vibrational energy is reduced by one vibrational quantum number during this process. Unfortunately, the atom will oscillate back to the $|m = +3/2, v = n + 1\rangle$ state after another half Rabi cycle period. In order to do cooling, we have to break this Rabi oscillation cycle. The optical pumping laser beam, represented by the green line, pumps the atom to an excited state. The atoms then spontaneously decay to the $m = +3/2$ state, completing one cooling cycle.

One key point for the optical pumping is the use of the $S_{1/2}$ to $P_{1/2}$ transition (the D1 line). The $P_{1/2}$ state does not have any $F = +5/2$ hyperfine manifold. Since the optical pumping light is σ^+ light, the $m = +3/2$ state is dark. In the best cooling scenario, once an atom reaches the ground vibrational state of the $m = +3/2$ state, the atom will be dark from all light transition. However, since we are imaging the spontaneously emitted photons, we do not want the atom to stay in the ground state, so some amount of residual heating is important.

There are a few other things that have to be taken care off. First, the Raman beam should be polarized properly such that the beam reduces the angular momentum by 1. This makes sure that we do not couple states that we do not want to couple. Second, the ground state of the $m = +3/2$ state should ideally be a dark state from both the Raman beam and the optical pumping beam. The polarization of the Raman beam makes sure that this state is a dark state. For the optical pumping beam, there are two things that have to be done. First, the beam should increase the angular momentum by $+1$. This alone would not help since the excited state has $m = +5/2$ state. The second point is to use a D1 transition instead of D2 transition so that the highest momentum state in the excited state is also $m = +3/2$. This makes sure that the entire $m = +3/2$ state is dark from the optical pumping light.

2.4.3   Degenerate Raman Sideband Cooling

Another variation of Raman sideband cooling is the degenerate Raman sideband cooling shown on Figure 2.7[66, 67]. Here, we apply a magnetic field in the negative z direction such that the Zeeman shift is exactly one vibrational energy level difference. The two Raman beams have the same frequency, simplifying the setup a little bit.
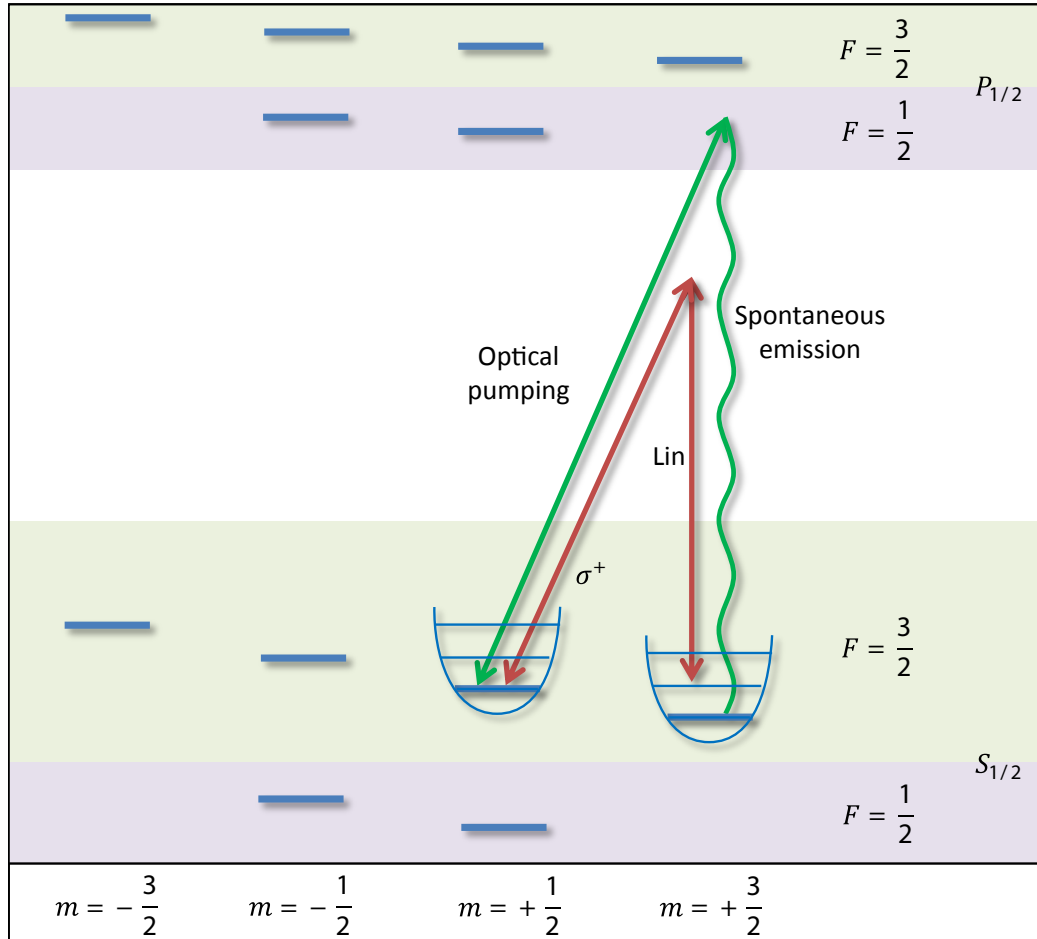
Figure 2.7: Degenerate Raman sideband cooling scheme.

2.4.4 Resolving Different Lattice Sites

Another feature that we want is the capability to resolve atoms on individual lattice sites. The Rayleigh criterion says that the resolution of an imaging optic, assuming purely incoherent scattering, is given by

$$R = \frac{1.22\lambda}{2\text{NA}},$$

where λ is the wavelength of the light and NA is the imaging numerical aperture. The intuitive meaning of this resolution is that R is the minimum distance between two point objects that can still be resolved by the imaging optic assuming infinitely good signal to noise ratio. However, in our setup there is information regarding the image that can be included in the image analysis process:

- **The atoms are identical**. Each atom emits more or less the same number of photons. In principle, from photon number counting alone, we can determine the number of atoms in the image, assuming there is no shot noise.

- **The atoms are pinned in the optical lattice**. There are only several discrete points where the atoms can reside in the image plane.

- **Binary readout**. Each lattice site can either have zero or one atom, but no more, which further reduces the possible number of configurations in the image.

If we include the above information in the image analysis, in principle we can resolve an arbitrarily small lattice spacing given infinitely high signal to noise ratio in the image[68]. In practice, however, the resolution is determined by the signal to noise ratio of the imaging system.

Imagine the case when we have exactly two lattice sites: left site at $x = 0$ and right site at $x = a$. Furthermore, let us assume that we have exactly one particle but we do not know on which site this particle is located. This particle then emits a single photon that is detected by the camera. We can calculate the likelihood that this particle is on the left lattice site.

$$f(x) = \frac{1}{\sqrt{\pi}} e^{-x^2}$$

P(x), PLeft(x)
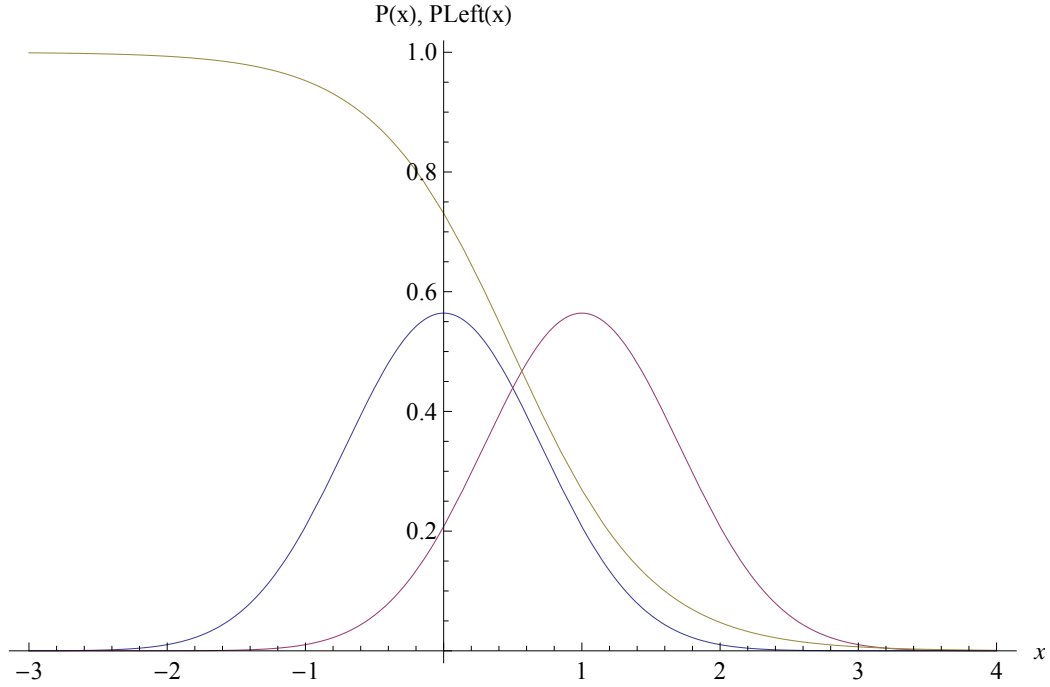
Figure 2.8: In this figure, we assume that there is exactly one point-source that can be located in either $x = 0$ (left side) or $x = 1$ (right side). The point source projects its point spread function in the camera which is represented by the two Gaussians above. We then detect a single photon from this particle. Our task is now to determine whether the particle is located on the left or on the right side. The probability that a particle is located on the left side is shown in the figure above.

$$g(x) = \frac{1}{\sqrt{\pi}} e^{-(x-a)^2}$$

$$P_{\text{left}}(x) = \frac{f(x)}{f(x) + g(x)},$$

where a is the known spacing between the two sites, $f(x)$ is the point spread function of

the left lattice site, $g(x)$ is the point spread function of the right lattice site[3], and $P_{\text{left}}(x)$

is the probability that the photon comes from the left site, given the location of the

detected photon. Without loss of generality, I also assume that the width of the gaussian

is one. Figure 2.8 shows the $a = 1$ case. If we detect the photon at the far left (right), we

---

[3]The ideal point spread function from an aperture is Airy disk[69]. Here, I use normal distribution for simplicity reasons.

are sure that the particle is on the left (right) lattice site. However, if we detect a photon that is located in between the two lattice sites, we do not gain any information regarding the source of the photon. Here, the probability is $P_{\text{left}}(x) = 1/2$ since the photon is equally likely to come from either lattice site.

The next step is to combine many photon measurements to determine the location of the source. Consider N photon measurements located at $x_1, x_2, .., x_N$. All of these photons can either all come from the left site or all come from the right site. Since each photon measurement is independent, the (unnormalized) likelihoods are just the multiplication of single photon likelihoods:

$$
\begin{aligned}
L_{\text{left}} &= \prod_{i=1}^{N} P_{\text{left}}(x_i) \\
L_{\text{right}} &= \prod_{i=1}^{N} P_{\text{left}}(x_i)
\end{aligned}
$$

We then can calculate the probabilities by normalizing the likelihoods:

$$
\begin{aligned}
P_{\text{total}}(\text{left}) &= \frac{L_{\text{left}}}{L_{\text{right}} + L_{\text{right}}} \\
P_{\text{total}}(\text{right}) &= \frac{L_{\text{right}}}{L_{\text{right}} + L_{\text{right}}} = 1 - P_{\text{total}}(\text{left})
\end{aligned}
$$

The next question is to ask how many photons do we need in order to reliably determine the location of the particle with say, 99.9% fidelity. I will describe a non-rigourous way to estimate this number. Intuitively, the more the point spread function overlaps, the more photons we need to determine the location. The overlap between two normal distribution is given by (ignoring normalization factor):

$$
\begin{aligned}
h^2(a) &= \int_{\infty}^{\infty} f(x)f(x-a)\mathrm{d}x \\
&= e^{-\frac{a^2}{2}} \\
h(a) &= e^{-\frac{a^2}{4}}
\end{aligned}
$$

We can then informally define that the amount of "information" in a photon as $1 - h(a)$,

so the estimated number of photons required should be proportional to

$$
\begin{aligned}
N_{\text{estimate}}(a) &= \frac{1}{1 - h(a)} \\
&= \frac{1}{1 - e^{-\frac{a^2}{4}}}
\end{aligned}
$$

Photons required to reach 99.9% fidelity as a function of spacing (x)
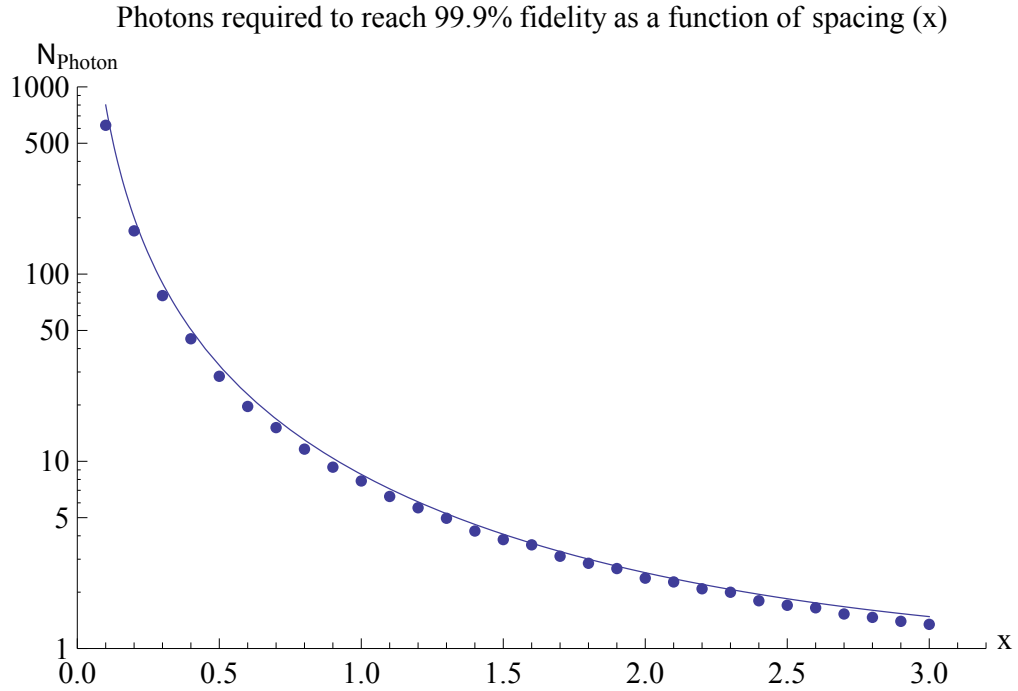
Figure 2.9:   The dots are the Monte Carlo simulation results on the number of photons required to reach 99.9% fidelity as a function of lattice spacing in the unit of the gaussian width. The solid line is $N_{es}(a) = \frac{1}{1-e^{-\frac{a^2}{8}}}$ to guide the eyes.

Another way to estimate the number of photons required is by doing Monte Carlo

simulations on x_i which is shown in Figure 2.9.

Chapter 3

Spectron Electronic Control System

Figure 3.1:   A collection of Spectron electronic boards. These are a sample of the Spectron electronic boards that we have developed.

An experiment control system is at the heart of all atomic physics experiments. In a

typical BEC setup, the user creates a sequence of electronic pulses and signals, which control various equipment in the lab. Many BEC labs use commercial solutions from companies like National Instruments or Viewpoint. Although using a commercial system is convenient, commercial systems are not usually designed for atomic physics applications. We decided from early on to design our own electronic control system from the ground up to fit exactly with atomic physics requirements.

This chapter is primarily meant to be the documentation for our lab's control system. However, I also hope that other people will find some inspiration from our control system in designing their own. The overview section describes briefly how the system works. The rest of the chapter describes each component in greater detail.

3.1   Overview

The electronic control system is called Spectron. Spectron is a distributed electronic control system that operates without a central computer. Instead, Spectron primarily uses on-board FPGAs as the core processing engine. The basic structure is shown in Figure 3.2.

The computer, which is a standard computer without any special hardware, acts as the user interface. The computer is responsible for sending commands to all Spectron boards across the lab before each experiment cycle. Each board then processes this command and generates the required real-time signal for the experiment. Note that there is no real-time requirement on the computer side.

Each command consists of multiple lists. Each row in a list represents an instruction. Each instruction is composed of the following parameters:
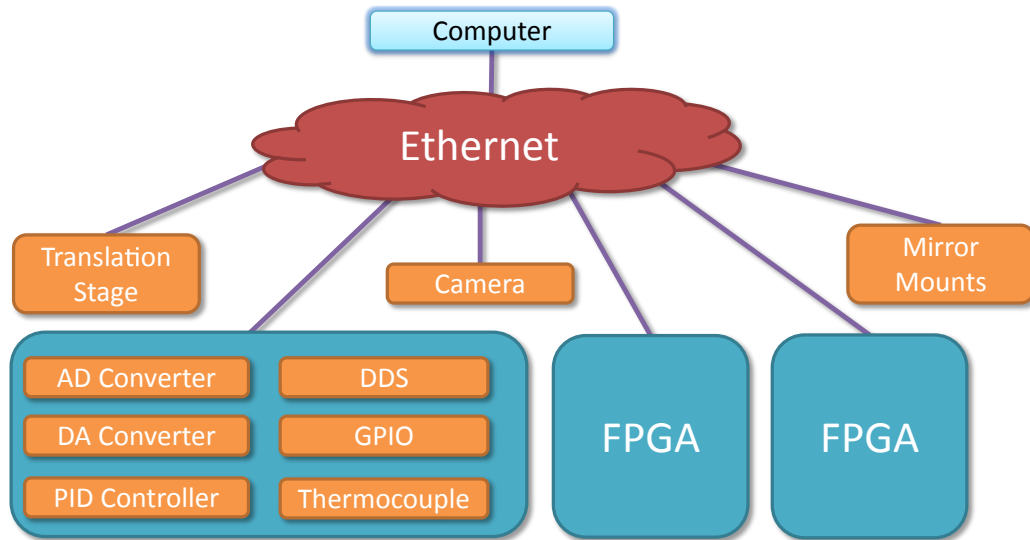
Figure 3.2: The basic structure of the Spectron system. The computer acts as the user interface only. The real-time hardware is located on the FPGAs which are distributed across the lab. All communication happens via the Ethernet.

- Timestamp for current instruction

- Initial value

- Ramp rate from the initial value

Let us consider the ramp shown in Figure 3.3 as an example. The ramp consists of 7 segments. The simplified table is shown at Table 3.1. Here, the table describes abstract waveform. We then connect these abstract waveforms into real devices to produce real signals. For example, the waveforms can represent voltage, current, frequency, amplitude, power, and phase. In this example, the waveform represents laser power. The next step is for the FPGA to turn this command into a sequence of values that can be used by other devices.

In typical commercial waveform generators, we usually have to specify the value at all times. For example, a 60- second experiment with $10\,\mu$s resolution requires 6 million data
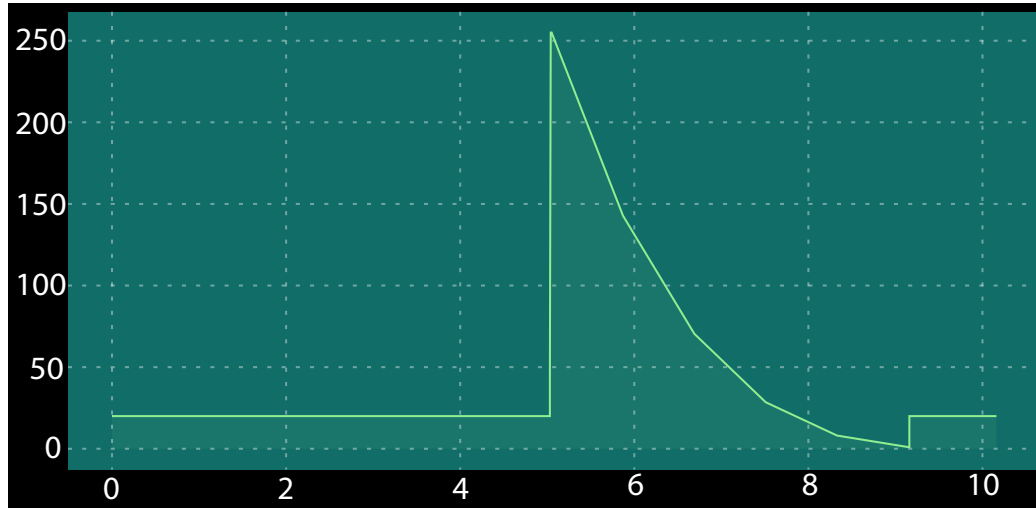
Figure 3.3: Our actual dipole trap evaporation ramp. The y axis is the laser power in Watts, while the x axis is the time in seconds. The 5 linear interpolation segments can be seen in the exponential ramp. In practice, the number of key points can be increased substantially so that the curve will look more exponential.

points. In practice the timing resolution is usually limited by the amount of available memory in the computer. In the Spectron waveform description we specify 3 things: time point, start value, and ramp rate. This greatly compresses the amount of data that has to be transmitted and stored in the system, allowing us to run the system at the maximum timing resolution of the hardware.

The next several sections describe the Spectron subsystem in greater details. The

Time	Initial	Ramp Rate
0.000 s	20.000	0.000
5.041 s	255.000	−135.518
5.869 s	142.791	−88.350
6.688 s	70.432	−51.210
7.507 s	28.491	−24.922
8.326 s	8.080	−8.520
9.157 s	20.000	0.000

Table 3.1: A sample of a ramp description. This is the actual evaporation ramp in our setup. The corresponding plot is shown in Figure 3.3.

Spectron Core Framework is the common infrastructure that is shared by all Spectron boards. The next several sections describe the Spectron Spy and Symphonia boards which are the two main implementations of the Spectron Framework.

3.2   Spectron Core Framework
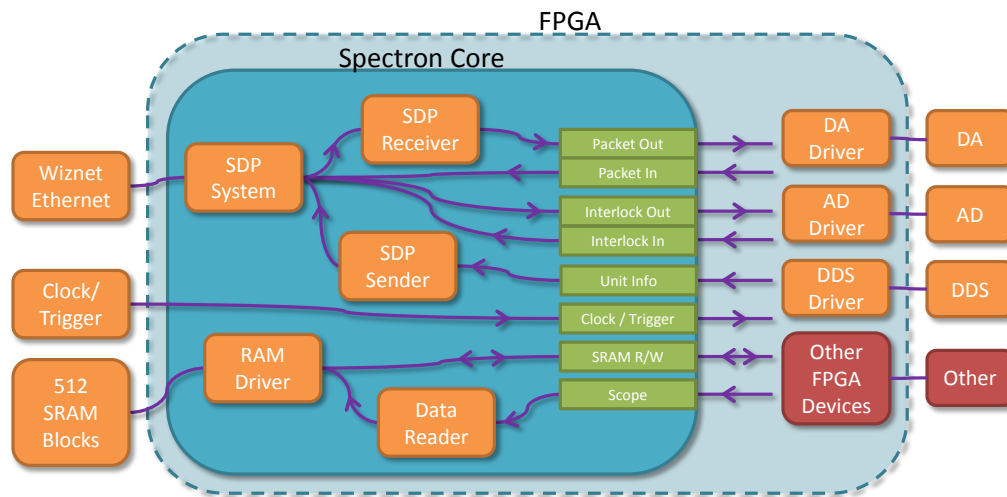
Figure 3.4:   Spectron core overview. The Spectron core is responsible in doing most of the heavy lifting such as timing, communication, memory and interlock. Every FPGA board in the lab uses identical Spectron Core as its internal engine.

The Spectron Core Framework is the core framework used by all of our Spectron boards in the lab. The framework handles the common requirements shared by all boards such as:

- Timing generation

- Communication, mac address, ip address

- Command processing and ramp generation

- Memory controller

- Oscilloscope

- Phase locked loop (PLL)

- Interlock

At the very minimum, a Spectron board that implements the framework must comprise the following:

- Ethernet connector for communications

- FPGA for real time control

- Soft Microcontroller to handle Ethernet packets

- JTAG programming port

- Clock and Trigger input for synchronization between boards

3.2.1 FPGA

There are two major vendors of FPGA, Xilinx and Altera. These two vendors also offer two different tiers of FPGA grade. From early on, we decided to use the cheaper grade FPGA since we would use many FPGAs simultaneously. At the time we chose the board, the then recently released Altera Cyclone 3 was the best cheap FPGA board available. All of our FPGAs in the lab are Altera Cyclone 3. We are using the Dallas Logic CMCS002 FPGA module with the following capabilities:

- 25k logic elements

- 594k bits internal FPGA RAM

- 512k x 8 bits external SRAM

Although this FPGA is large and fast enough for most applications, we are reaching a point where we are limited by the FPGA logic gate capacity, with the Spectron Core Framework alone taking up 11k logic elements. It might be beneficial in the near future to upgrade some of the FPGAs to higher capacity devices.

3.2.2   Communication

Communication is a critical part of the system since there are multiple Spectron boards across the lab. When the computer sends a command to a device inside a Spectron board, we have to route the command from the computer to the correct device. The first step is the communication from the computer to the FPGA. Once the information arrives at the FPGA, we have to route the information within the FPGA itself.

USB

We first tried USB as the communication protocol between the computer and the FPGA. At the time, this made a lot of sense since most computers have USB ports, not to mention the fact that some commercial products from National Instruments use also USB protocol.

We used chips from FTDI to handle most of the USB protocols which greatly simplified the USB implementation. The FTDI chips abstract out the USB protocol as a bi-directional stream of data between the computer and the FPGA. This method worked really well for our first generation FPGA board.

Unfortunately, for reasons that we still do not fully understand, USB protocol is not

scalable for a large number of devices. Computer operating system tends to crash whenever we connect multiple USB hubs with a lot of devices to a single computer. Another downside is that we need a direct connection between a computer and an FPGA at all time to run the experiment which necessitates the use of multiple computers scattered around the lab. Due to these limitations, we decided to abandon USB protocol and try Ethernet communication instead.

Ethernet

Ethernet is a much more scalable system compared to USB. In fact, most of the internet is powered by Ethernet protocol. Another advantage is that a direct computer connection to the FPGA is no longer required. Furthermore, Ethernet protocol allows for simultaneous point to point communication and also broadcast type communication.

To simplify the Ethernet implementation as much as possible, we use Ethernet controllers from Wiznet to offload most of the Ethernet stack away from the FPGA.

Once we pick Ethernet as the low-level communication protocol, we have to choose a higher-level protocol too. There are several options for higher-level protocol such as IP, TCP, and UDP. For our purposes, the communication protocol has to satisfy the following requirements:

- Lightweight due to the size/complexity constraint of the FPGA

- Reliable: make sure that all packets are received by the FPGA

- Addressable: address individual FPGAs and also individual devices inside the FPGAs

The most natural thing to do is to assign an IP address to each FPGA as is in the case of normal computers. We then use a higher level protocol on top of the IP protocol to satisfy other requirements. TCP protocol guarantees that all packets are received by the receiver in the correct order[70]. To achieve this, TCP uses an open-socket connection between the sender and the receiver. The TCP stack firmware then keeps tracks of all packets and makes sure that all packets are sent / received in the correct order. In our system, since the computer and the Spectron boards have to communicate with one another, many simultaneous open TCP sockets are required. Although TCP protocol is ideal for our reliability requirement, in practice there are not enough resources in the FPGA to handle the multiple open TCP connections. This constraint rules out TCP as our communication protocol.

Our next candidate is the UDP protocol which is very lightweight and appropriate for a small FPGA system[71]. The downside is that UDP is not reliable: packets could be lost or arrive out of order without any notification. To compensate for this, we have to ensure reliability using some other means which will be described in the next few subsections. In principle, UDP protocol also allows us to address individual device inside the FPGA by assigning a UDP port for each device. Unfortunately, due to the limitation in the Wiznet Ethernet chip on the Spectron board, we can only listen to a very limited number of UDP ports which rules out using UDP ports to address the many devices inside the FPGA.

Our solution to satisfy the addressability requirement is to create a proprietary protocol, Spectron Datagram Protocol (SDP), mentioned in the next subsection. The reliability requirement is addressed in the Ram Write Confirmation subsection.

Spectron Datagram Protocol (SDP)

In most cases, the SDP is implemented on top of an existing layer 4 protocol such as TCP or UDP[72]. The SDP itself is a datagram protocol, which means that data comes in packet format, which makes it very similar to UDP. However, the SDP is general enough that it can be implemented on top of any connection. During the early prototyping days, the SDP protocol was implemented on top of a serial connection (RS 232), USB connection, and a simple ribbon wire connection.

	Bit 31 - 24	Bit 23 - 16	Bit 15 - 8	Bit 7 - 0
0	Start Magic Number (0x3f_58_95_0d)[31:0]			
4	Source IP Address[31:0]			
8	Source IP Port[15:0]		Source SDP Address[15:0]	
12	Destination IP Address[31:0]			
16	Destination IP Port[15:0]		Destination SDP Address[15:0]	
20	Command Code[15:0]		Auxilliary Command Code[15:0]	
24	Reserved[7:0]	PacketID[7:0]	RC C RR Reserved[1:0]	Data Size[10:0]
28	Data			
n-8	Reserved [15:0]		Checksum [15:0]	
n-4	End Magic Number (0xfb_d9_36_60)[31:0]			

Figure 3.5: Spectron Datagram Protocol (SDP) specifications. The protocol is inspired by the standard UDP protocol.

In SDP, the data is transferred in a packet format, where a packet is just a finite chunk of data. The complete specification of an SDP packet is shown on Figure 3.5. A packet contains a packet header block, a packet data block, and a packet footer block which is sent in most significant bit (MSB) format. The total length of the packet is always a multiple of 32 bits which allows 32 bit processing if required. If the packet does not have the correct packet start magic number, packet end magic number, and correct checksum the packet is discarded.

Packet Header (28 bytes)   The packet header contains the information about the packet itself. The length of the packet header is fixed at 28 bytes. The format for the packet header is the following:

- 32 bit: Packet start magic number (0x3f_58_95_0d). This magic number is used to synchronize the packet.

- 64 bit: Source address. This field details the address of the sender of the packet. This is useful if the receiver has to send a reply to the sender.

- 64 bit: Destination address. This field details the address of the receiver of the packet. This is useful because there is additional routing required inside the FPGA.

- 16 bit: Packet command code. This is a multi-purpose field that can be used to describe what a device should do with the packet. Some example of possible commands are: wait for external trigger, force trigger, run and repeat, write to memory.

- 16 bit: Auxilliary command code. This serves to augment the command code.

- 8 bit: Reserved, all 0.

- 8 bit: PacketID. For identifying packet confirmation. This should be the serial number of the packet. In practice, it turns out that this field is not required for normal operation.

- 1 bit: RC (Request Confirmation). Set means that this packet requires a confirmation to be sent by the recipient. This field is not used in normal operation.

- 1 bit: C (Confirmation). Set means that this is a confirmation packet. A confirmation packet never generates another confirmation.

- 1 bit: RR (Request Routing). Set means that this packet should be routed to the proper destination by the gateway. This feature is not currently implemented.

- 2 bit: Reserved, all 0.

- 11 bit: Packet data size in bytes, NOT including the header and the footer. Can be from 0 to 1024 (Note 1024 is 2^{10}, which requires 11 bit in binary). The maximum data size is chosen so that it will fit exactly in one Altera memory block.

Packet Data (0 to 1024 bytes)   The packet data contains the data. The type of data is mentioned in the packet header. The length of the packet data can be 0. The maximum size of the packet is chosen such that it will fit in an Altera memory block (8K bits). Although the length of the packet data can be a multiple of 1 byte, in the packet itself, the packet data has to be a multiple of 32 bits (4 bytes), which is achieved by padding 0s at the end of the packet data.

Packet Footer (8 bytes)   The packet footer contains the checksum information of the packet header and packet data. The length of the packet footer is fixed. The format is the following:

- 16 bit: Reserved, all 0

- 16 bit: Checksum of the combined packet header and packet data. The check sum is computed by adding the data in 16 bit chunks, while ignoring the overflow bits. Note that this is not a good checksum. We use the simple checksum algorithm to minimize the space used inside the FPGA.

- 32 bit: Packet end magic number (0xfb_d9_36_60). This magic number is used to synchronize the packet.
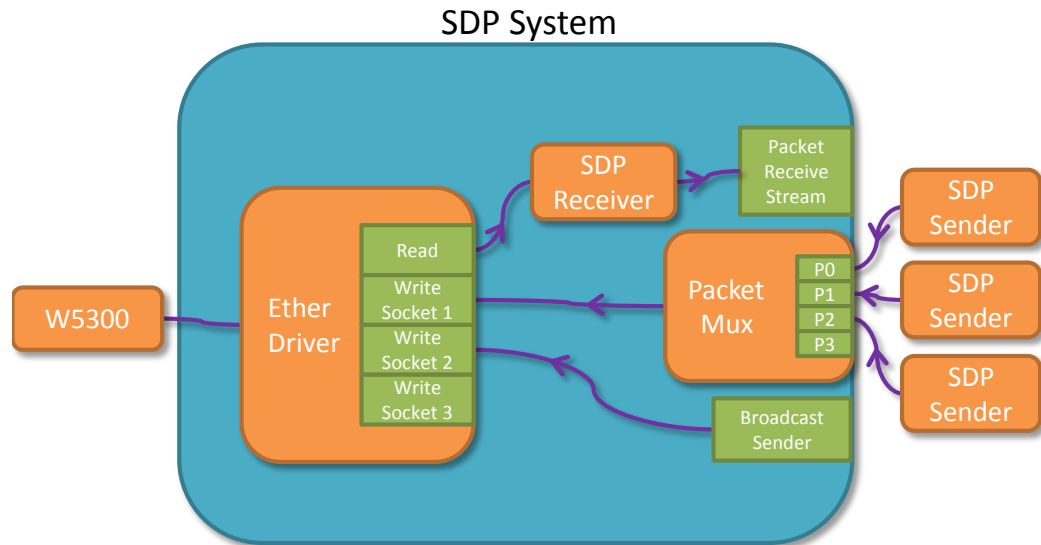
SDP System

Figure 3.6: Spectron Datagram Protocol System. The SDP system handles the communication between the FPGA and the computer.

SDP Receiver The SDP receiver is responsible for processing the incoming stream of data from the Ethernet socket and generates parsed packet data to be used by other devices. Since there is only one Ethernet connection per FPGA, only one SDP receiver is required per FPGA. The receiver uses a double buffer technique in order to be able to received new data while parsing old data at the same time.

The incoming data stream is analyzed by the SDP receiver's Input Stream Analyzer which uses the magic number to find the starting point of a packet. If a packet start magic number is found, the packet header and packet data are then stored in one of the buffers, in this case, buffer A. At the same time, the checksum counter performs a checksum
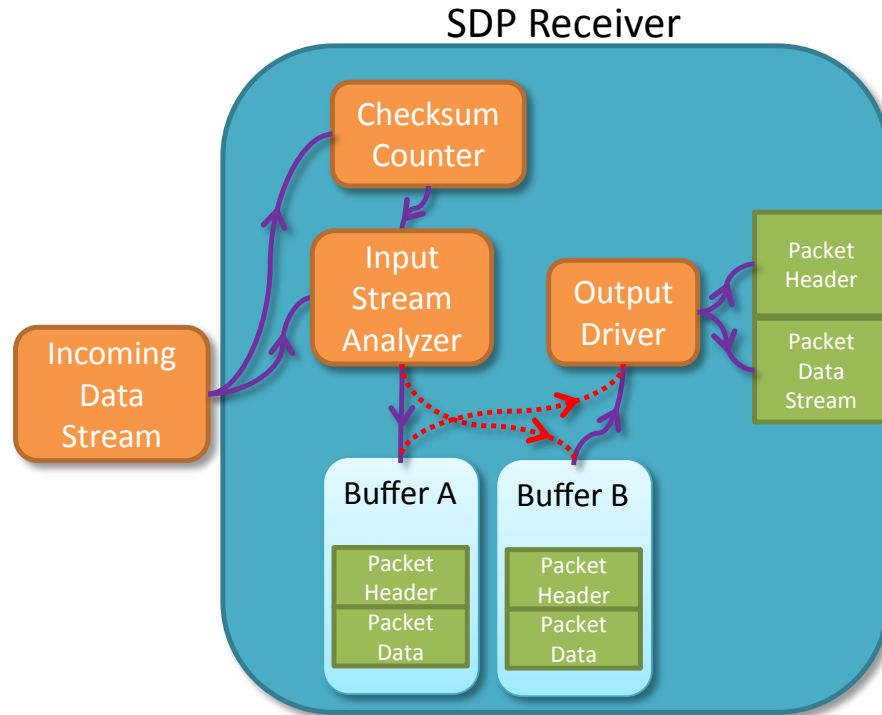
Figure 3.7: The SDP Receiver process the incoming SDP packets from the computer and also other FPGAs and distribute the parsed packets throughout the FPGA.

calculation on the packet. If the packet end magic number is correct, and the computed checksum matches with the checksum field on the packet footer, the packet is accepted and outputted by the output driver to be used by other devices in the FPGA.

While the Input Stream Analyzer is analyzing and storing the incoming packet data, the output driver simultaneously outputs the packet already stored in the other buffer to be distributed across the FPGA. Note that the packet header contains the SDP address which addresses the individual devices in the FPGA. The output driver simply broadcasts the entire packet across the FPGA. Therefore, it is the responsibility of each individual device to listen only to packets addressed to it and ignore packets addressed to other devices.
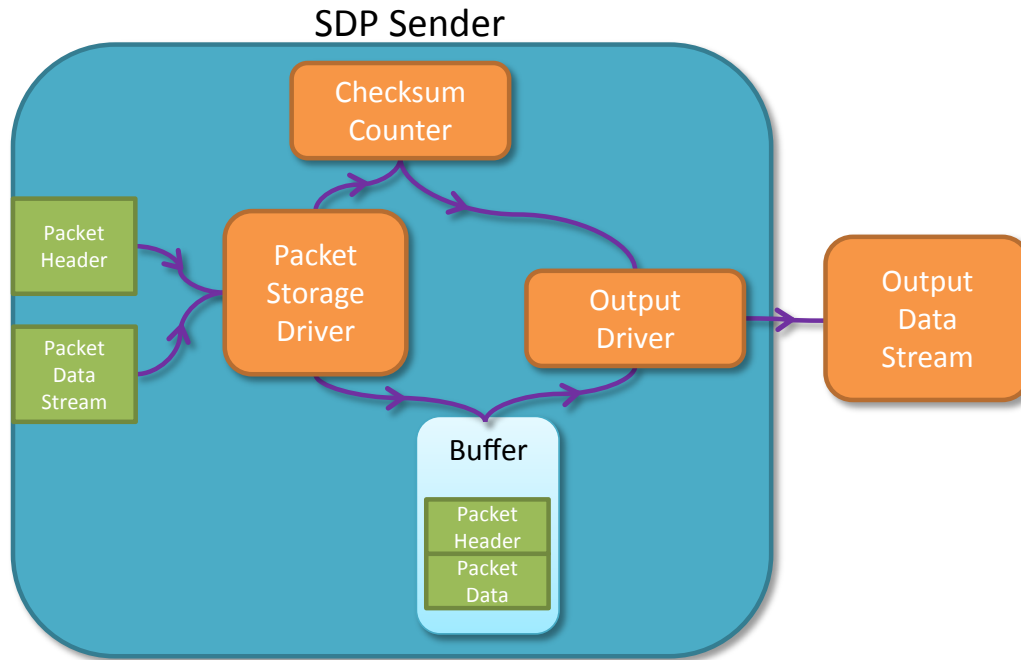
Figure 3.8: SDP Sender is used to generate an SDP packet. The SDP packet is then sent to other boards through the Ethernet.

SDP Sender The SDP Sender works the opposite way as the SDP Receiver. The default SDP Sender uses a single-buffer technique so the SDP Sender cannot store and transmit packets simultaneously. It is the responsibility of the device using the SDP Sender to wait until the SDP Sender is ready before sending the next packet.

The handshake between the SDP sender and a device is the following:

- The SDP sender tells the device whether the sender is ready or not.

- The device requests a packet to be sent to the SDP sender.

- The SDP sender then goes to busy mode

- The SDP sender stores the packet while computing the checksum at the same time.

- Once the device finishes writing the packet, the SDP sender then generates a data

stream from the packet. The stream is then used by the next communication
component, usually a packet multiplexer

- Once the sending is completed, the SDP Sender goes back to ready to send mode

Note that since there are many devices inside an FPGA, there have to be many SDP
senders too. In general, there is one SDP sender for each device that wants to send SDP
packets. Therefore, it is important that the SDP Sender does not take too much space in
the FPGA. This is the reason why the SDP sender uses single buffer technique. When a
device only needs small communication bandwidth, the device can use the Unit Info
infrastructure which is both simpler to use and consumes far less resources than the SDP
Sender. The Unit Info is discussed in a separate section.
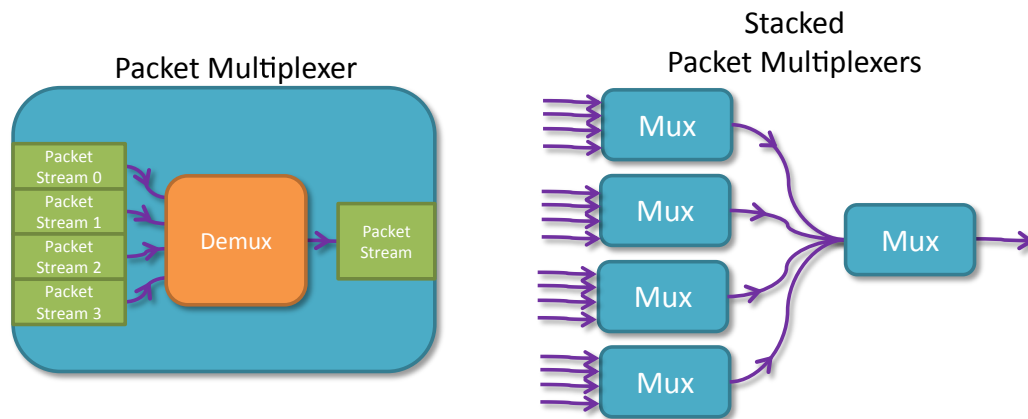
Figure 3.9:   Packet Multiplexer is used to share the single SDP sending port so that the
port can be used by all of the SDP senders in the FPGA.

Packet Multiplexer   Since there are many SDP senders while there is only one
Ethernet socket, we have to demultiplex the data streams generated by the SDP senders
into a single stream. The packet multiplexer[1] is responsible for combining this multiple

---

[1]Technically, the name should be packet demultiplexer. Packet multiplexer is used for historical reason.

data stream.

A packet multiplexer has 4 stream inputs and one stream output. If we need more than 4 inputs, we can stack packet multiplexers back to back to generate the required number of inputs. The packet multiplexer works by streaming the input stream one by one to the output stream with some flow control to prevent collisions between input streams.
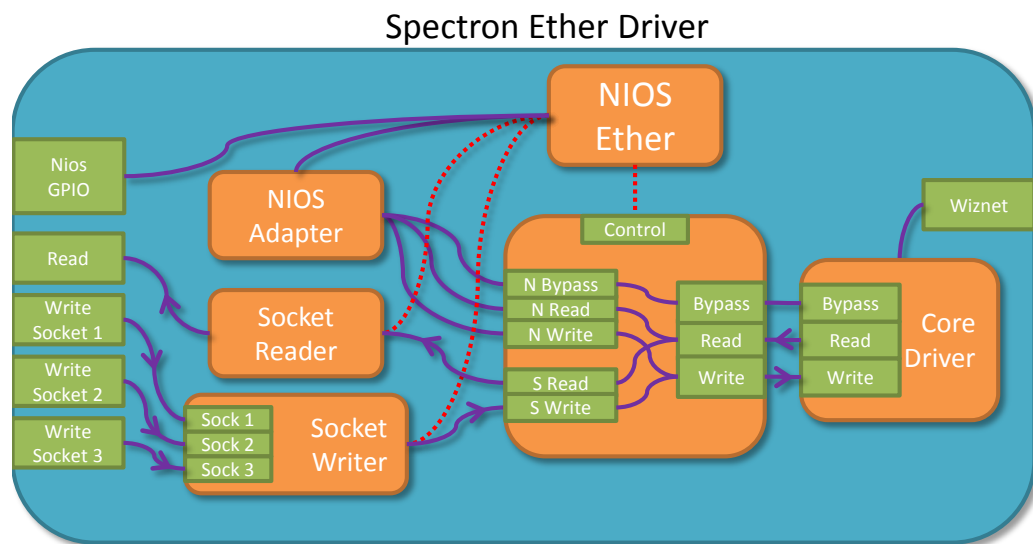
Figure 3.10: The Ether Driver handles the initialization and communication with the Wiznet Ethernet chip.

Ether Driver One of the most complex parts of the communication protocol is the Ethernet driver. The ether driver is responsible for bridging the communication from the Wiznet chip with the rest of the FPGA. Devices in the FPGA see the ether driver as a component with data input and output stream. The ether driver handles the following:

- Setting up IP Address and port allocation

- Processing incoming and outgoing UDP packets

- Communication with the Wiznet chip on when to send or receive packets

From the rest of the FPGA point of view (from the left), the Ether driver looks like a device with 1 read socket, 3 write sockets, and several GPIO handshaking pins. On the right side, the Ether driver provides the necessary wiring to the external Wiznet chip. At the center of the Ether Driver is the NIOS Ether soft microcontroller. Note that the vast majority of the data can flow from the Read and Write sockets through the Wiznet chip without going through the microcontroller. This technique greatly increases the communication bandwidth between the FPGA and the Wiznet chip at the expense of additional complexity.
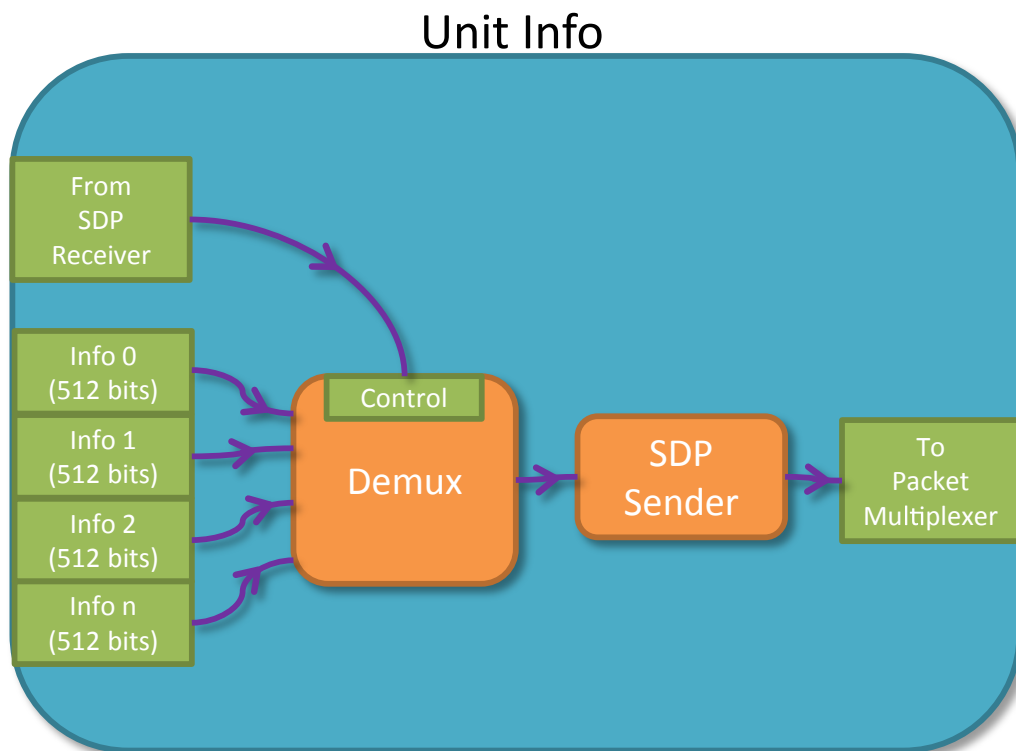
Unit Info

Figure 3.11: Unit Info is used to sent small amount of information from the FPGA to other boards in the lab. Unit Info is appropriate when the amounts of data to be sent does not warrant full SDP sender implementation.

The unit info infrastructure is created to minimize the number of SDP Senders that are required in the FPGA. The idea is to share a single SDP sender to send data from multiple devices at the expense of communication bandwidth. It turns out that the unit info is very useful for sending this slow changing data such as temperatures, state of the FPGA, flow rate, and interlock status among other things.

The basic idea is to have a device with multiple input pins, 12 x 512 bits in our setup. We then demultiplex and serialize those inputs pins so that they fit inside a single data stream. The stream is then passed to the SDP Sender to be sent to whoever requested the data. The Unit Info itself is designed only to send data when the data is requested, usually by the computer.
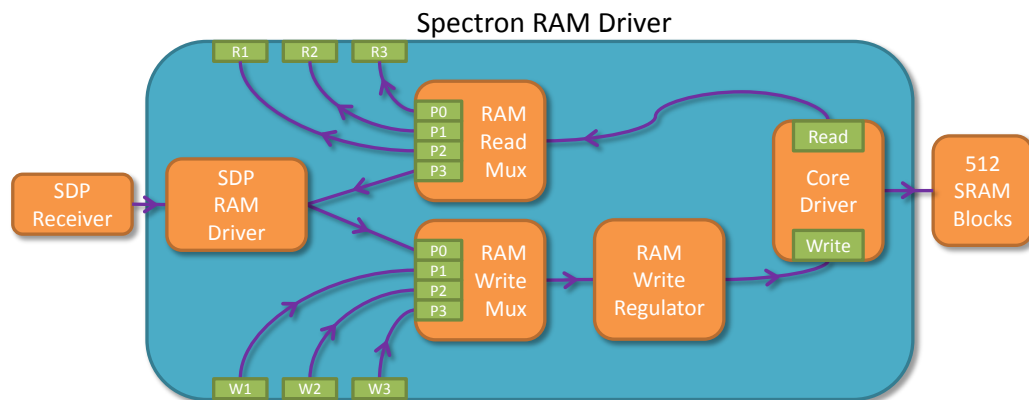
3.2.3 RAM Driver

Figure 3.12: Ram Driver handles data storage and retrieval from both the SRAM and the NVRAM.

Another integral part of the Spectron Core Framework is the RAM driver. A Spectron board uses two rams, Synchronous Random Access Memory (SRAM) and Non-Volatile Random Access Memory (NVRAM). The SRAM is used to store the real-time experiment

commands and oscilloscope data, while the NVRAM is primarily used to store the FPGA IP address. Since the SRAM and NVRAM use different communication protocols, there is a separate core ram driver for each of those RAMs. The rest of the driver code is shared between the RAMs.

The RAM driver has 4 read ports and 4 write ports. One read port and one write port are used exclusively for SDP communication. This allows the computer to read/write data from/to the RAM directly. The rest of the read/write ports are used by other devices in the FPGA for their own use. In practice, there are more than 3 devices inside the FPGA that need to use the RAM. The number of effective ports can be increased by using stackable RAM reader multiplexers and RAM writer multiplexers in a similar way as the Packet Multiplexer.

Ram Write Confirmation (RWC)

As mentioned before, the UDP protocol is not reliable. Instead of making sure that every single packet is received, we decided that only critical packets should be checked. The critical packets are the packets that write data to the SRAM to be used for real-time control. The RAM Write Confirmation is there to make sure that all write packets arrive safely.

The Ram Write Confirmation works in the following steps:

- The computer sends a start command to reset the confirmation status.

- The computer sends data to write to the ram. The RWC monitors this transmission and records which SRAM memory block is written.

Time	Output (88 bits)
0.00 seconds	000000000000000..1111
2.37 seconds	1100100110100110..1001
5.10 seconds	1100111111111110..1001

Table 3.2: Example of data writer commands.

- After all RAM blocks are written, the computer asks for the RWC to send the confirmation status.

- The computer then resends any RAM blocks that has not been written yet due to missing packets.
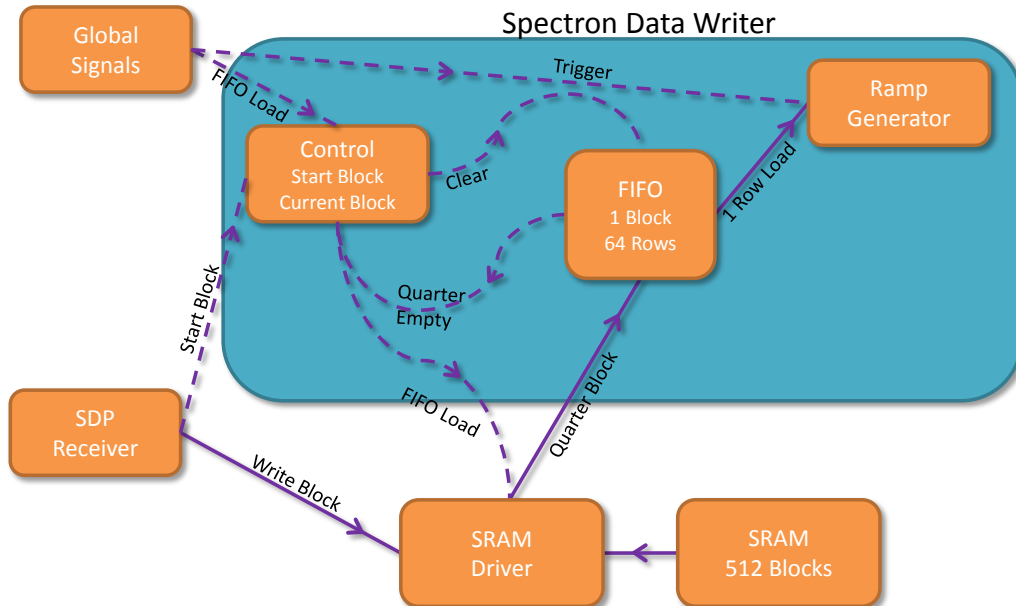
3.2.4 Data Writer

Figure 3.13: The Data Writer is used to generate real-time digital signal to be used by other devices in the FPGA.

1

The data writer is the heart of the real-time experimental sequence. Conceptually, the data writer generates an 88-bit data point at a specified time. An example of data writer commands is shown on Table 3.2. Other devices can simply use the data generated by the data writer to achieve real-time performance.

The data writer works side by side with the ram driver. The computer uploads the experiment sequence directly to the RAM. The computer then tells the data writer which ram address range is used for a particular data writer. The data writer then loads the data from the ram to its internal FPGA FIFO block. The output of the FIFO block is then fed into the output stage that outputs the desired 88-bit data at the correct time.

The data that is written to the ram contains multiple rows. Each row has the following format:

- [127 : 96]: 32-bit time stamp in 100 ns unit. This allows 429 seconds addressing.

- [95 : 88]: 16-bit bool values.

  - [95 : 91]: 5-bit reserved space.

  - Bit 90: Bypass enable.

  - Bit 89: dataChanged. If 1, dataChanged pin is pulsed. If 0, the dataChanged pin stays 0. Useful for extending the time stamp beyond 429 seconds.

  - Bit 88: isLastRow. If 0, continue read the next row. If 1, do not read the next row.

- [87 : 80]: 8-bit auxilary digital output.

- [79 : 0]: 80-bit value for this time stamp.

Real Ramp Writer

Real ramp writer takes the 80-bit output from the data writer and use the data to generate single 32-bit ramps. The 80-bit data has the following format:

- [79 : 48]: 32-bit initial value. This value is extended to 48-bit by padding the value with 16 zeros.

- [47 : 0]: 48-bit base ramp rate. The initial value is incremented by this value every 100 ns.

The 48-bit ramp rate with 100 ns is chosen to maximize the useful dynamic range of the ramp. The fastest ramp is 100 ns ramp from 0 to 1, while the slowest ramp is 0.9 years from 0 to 1.

3.2.5 Data Reader

The data reader works in similar but opposite way to the data writer. From the user points of view, the Spectron Data Reader takes a stream of data and stores the data on the SRAM for later use.

The process is the following. First, the computer tells the data reader which address block in the SRAM to be used to store the stream of data. The computer then reset the state for the data reader which involves resetting the state machine and emptying the internal buffer. The data reader then sequentially stores the 8-bit data stream to the SRAM. Since writing data to the SRAM involves significant overhead and latency, the data stream is first stored on the internal First In First Out (FIFO) block buffer. Once the FIFO is a quarter filled, the Spectron Data Reader then attempts to flush the data from the FIFO
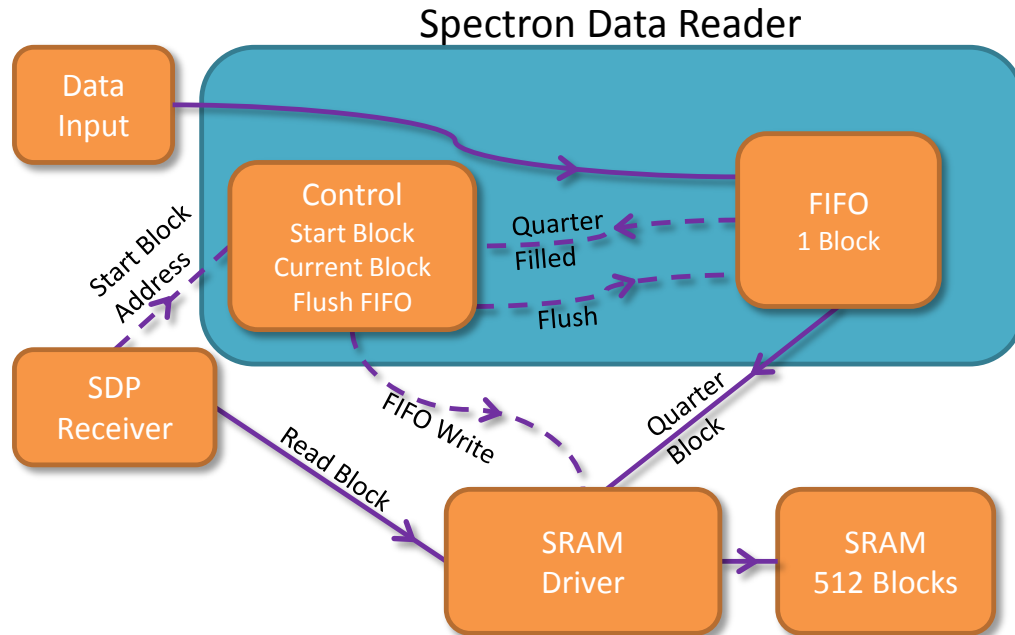
Figure 3.14: Data Reader is used to store real-time information from a device in the FPGA to the SRAM. The data can then be read from the computer for further analysis.

to the SRAM. Note that the reader has to wait until the RAM is ready before the data from the FIFO can be flushed. This process continues until the data stream is terminated. The computer then reads the data directly from the SRAM.

Spectron Scope

As of now, the only device that uses of the Spectron Data Reader is the Spectron Scope. The Spectron Scope is a general purpose 4 channel oscilloscope with the following features:

- 8x multiplexed 16-bit input

- Sample clock rate from 200 ns to 100 seconds

- Internal trigger to trigger the scope at a time synchronized with the main sequence

- External trigger

Each channel of the Spectron Scope simply serializes the 16-bit input stream into 8-bit data stream. This 8-bit data stream is then passed to the Spectron Data Reader for further storage. The computer then reads the data directly from the SRAM.

3.2.6   Clock and Trigger

An important part of the Spectron Framework is the clock and trigger distribution system. As currently implemented, the Spectron Framework shares a single clock and trigger across the entire lab. The clock source is an SRS FS725 Rubidium Frequency Standard (Atomic Clock) outputting 10 MHz. A typical Spectron FPGA board then up-converts this frequency to 100 MHz using Phase Lock Loop (PLL) circuitry for its internal clock. Since the entire lab's clock is derived from one clock source, the entire lab is phase-locked.

The trigger signal is generated by one of the Spectron Spy boards. The signal is then amplified and distributed across the lab to be used by other Spectron boards. The trigger signal itself is synchronized to the 60 Hz AC line. The idea is to reduce the noise related to the AC line phase error between experiment runs although we have not done any measurement on the usefulness of this synchronization.

Timing

One core component in the real-time sequence generation is the timing subsystem. The Spectron Core Framework uses a single timing source that is shared across the FPGA. With the current implementation, one real-time cycle is set to be 100 ns, which is 10 FPGA clock cycles. The timing system is responsible for generating, starting, stopping,

and resetting this real-time cycle clock depending on the commands from the computer and from the trigger signal.
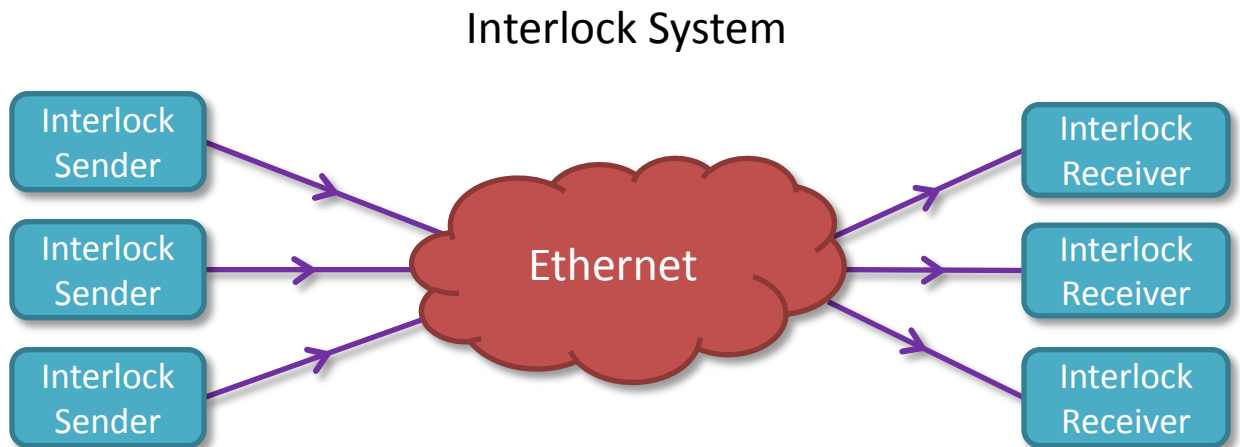
Interlock

Figure 3.15: Spectron Interlock. The interlock system is distributed. Multiple interlock senders send information about the status of each sensors via SDP packets. Multiple interlock receivers then intercept these packets and make sure that all sensors are ok before activating the interlock signal. The distributed approach reduces the chance a single failed interlock device from ruining the experiment.

We use a distributed interlock system in our lab that is integrated with the rest of the Spectron Core Framework.

The system consists of several interlock senders with several interlock receivers. Note that an FPGA board can contain both an interlock sender and an interlock receiver. Each interlock sender has one or more interlock sensors attached to it. The sensors could be temperature sensor, water flow meter, or laser power. Each interlock sender then broadcasts the status of the sensors (either within predefined range or not) to the Ethernet using UDP broadcast protocol every 200 ms.

The interlock receivers then receive these broadcasts from the multiple interlock senders for further processing. The receiver makes sure that all interlock senders are active by using a watchdog timer. The receiver then makes sure that every sensor operates within the predefined acceptable range. The moment any sensor in the lab falls outside of this range, each interlock receiver trips its own interlock status. Note that since all interlock receivers receive the same broadcast signals, when one receiver trips, every single receiver in the lab trips at the same time.

Each interlock receiver then outputs the interlock signal to the appropriate equipment using the Interlock watchdog board. The watchdog board makes sure that the interlock receiver is working properly at all times for it to output the interlock signal.

The distributed approach has the following advantages:

- There is no single point of failure. A single failed interlock sender will trip the entire interlock system. A single failed watchdogged interlock receiver, however unlikely that is, will not affect other interlock receivers. For example, the magnetic coil is interlocked at multiple locations by several different interlock receivers to make sure that there is always one working receiver to shut down the coil.

- There is no dedicated cable required between the sensors and the receivers since we are using the existing Ethernet connection for all communication.

3.2.7 Spectron Advanced Loop

Spectron Advanced Loop is a general purpose feedback controller module that can be implemented inside the FPGA. Some uses are laser frequency locking, laser power stabilization, and current stabilization.

Some of the features are:

- Real time P and I gain change. Sometimes, it is useful to change the gain in real-time during experiment sequence when the plant to be stabilized has non-linear response depending on the stabilization level.

- 4 measurement input ports. These 4 measurements can be combined using the built-in fuzzy circuitry (or any other custom function) to generate the final measurement signal. The relative weighting of the 4 measurements can be adjusted in real-time. This is useful for example when we have to transfer smoothly from one measurement device to another one due to limited dynamic range of the measurement device. One such application is laser power stabilization with two photo diodes.

- Real-time set-point change using Spectron Real Ramp Writer.

- Build in function generator. We can use the function generator to both scan the output value directly and also to scan the set-point.

- Lock/unlock pin to switch between feedback mode and direct output drive mode. During changing from unlock to lock, the initial value of the integrator is set to be the previous output value. Since we can ramp the I gain in real-time, this locking process can be made smooth.

3.2.8 Summary

At the end of the day, the Spectron Core Framework provides most of the heavy lifting required in order to create a real-time device. Here, a device could be a digital to analog
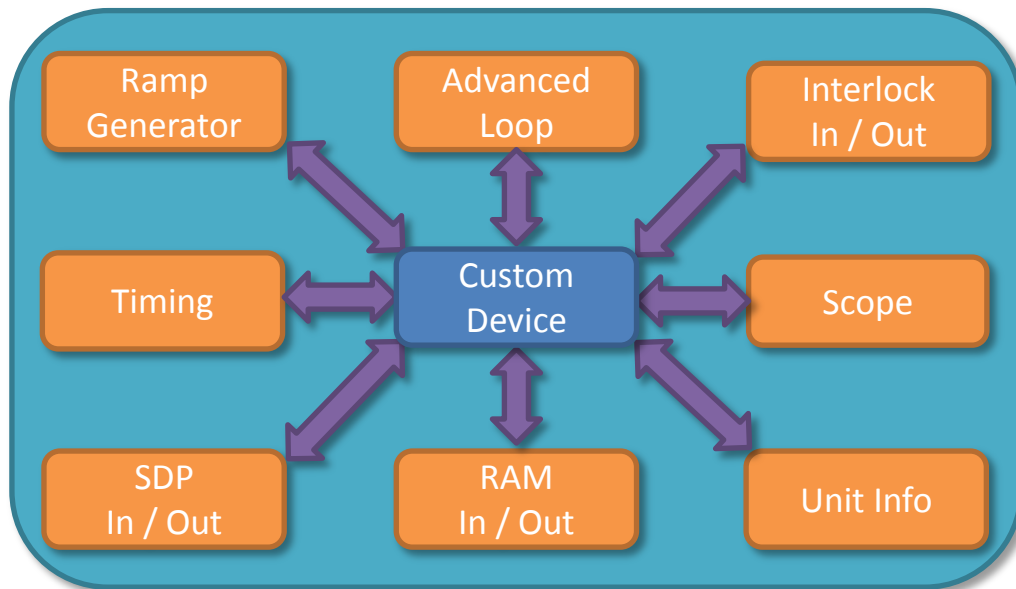
Figure 3.16:   From the custom device point of view, the Spectron Core Framework provides the implementor with numerous commonly used tools so that the implementor can focus on implementing the custom device.

converter, DDS, or analog to digital converter. From the device programmer's point of view, the Spectron Core Framework provides the majority of the commonly used features which greatly simplifies the custom device programming.

3.3   Spectron Master Spy

The most versatile and simplest Spectron Board implementation in our lab is the Spectron Master Spy. The board has 8 standard 8P8C Ethernet connectors on the front panel. We then connect modular Spectron Spy daughter boards on one (or more) of the Spectron Master Spy connector. In this section, I will describe how the Spectron Spy communication protocol works while the next section describes some Spectron Spy daughter boards that we have in the system.
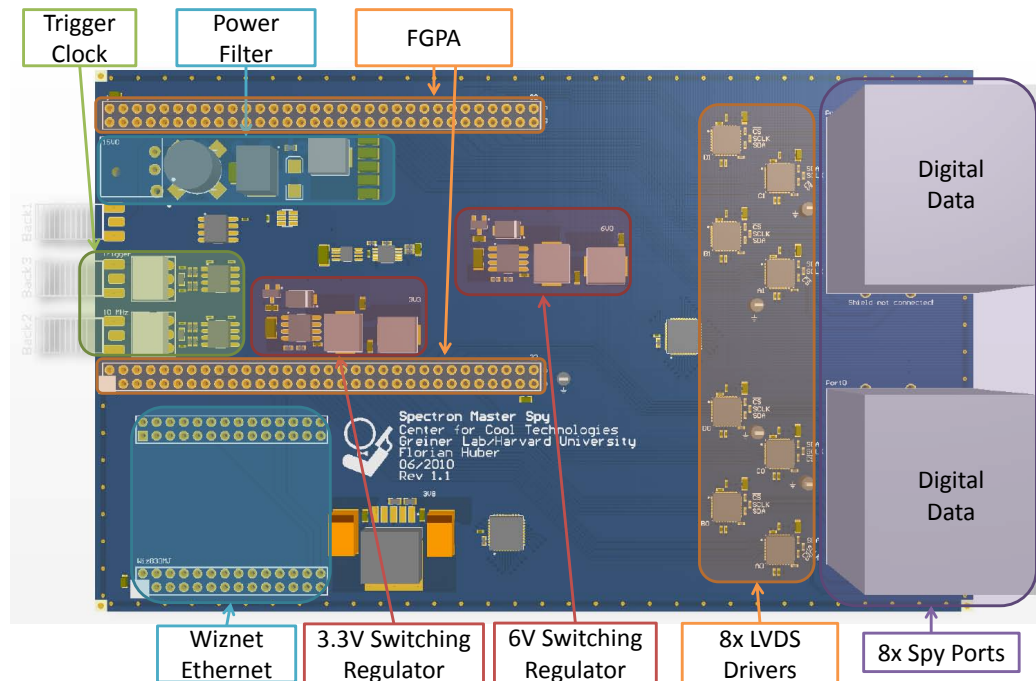
Figure 3.17: Spectron Master Spy, the most versatile FPGA board in our lab. The Spectron Master Spy can drive up to 8 devices that support the Spectron Spy protocol. There are more than 20 different Spectron Spy devices (and counting) that can be connected to the Spectron Master Spy.

3.3.1 Spectron Spy Protocol

The Spectron Spy Protocol is our own proprietary communication protocol to connect Spectron Master Spy board to Spectron Spy daughter board. The protocol pin configuration for the master board, following EIA-568-B is shown on Table 3.3.

Each Spectron Spy connector has three bi-directional pairs for general purpose data connection with an additional pair reserved for 6 Volt power transmission from the master board. Depending on the daughter board, each of the bidirectional data connection can be configured as either input or output. The power connection allows us to power the daughter board using the same cable which greatly reduces the need of a separate external power supply for each daughter board. The voltage for the power transmission is chosen

Pin	Color	Function	Pair	SPI
1	Orange Stripe	Bi-directional data negative pin	2-	$\overline{CS}$
2	Orange	Bi-directional data positive pin	2+	$\overline{CS}$
3	Green Stripe	6 Volt	3-	
4	Blue	Bi-directional data positive pin	1+	SCLK
5	Blue Stripe	Bi-directional data negative pin	1-	SCLK
6	Green	Ground	3+	
7	Brown Stripe	Bi-directional data negative pin	4-	Data
8	Brown	Bi-directional data positive pin	4+	Data

Table 3.3: Spectron Spy Protocol

to be 6 Volt so that we can use 5 Volt linear regulators (or lower) at the daughter board side to power 5 Volt circuits. The SPI pinout on the protocol is given only as recommended pinout since from the FPGA point of view, the three bi-directional pins are fully symmetric.

Why Digital Signal is Awesome   The Spectron Spy protocol is designed to keep all signal digital as much as possible, especially in the long cable section. The reason is that digital signal is exponentially more immune against noise compared to analog signal. This advantage is best illustrated by an example.

Consider an extreme example where we want to send a 1000 dBc signal over an optical fiber. As a comparison, an iPod has 83 dBc dynamic range. With digital signal, 1000 dBc is equivalent to 166 bits signal, which can be transmitted through the fiber with a few times 166 photons. With analog signal, 1000 dBc is equivalent to 10^{50} dynamic range. To overcome shot noise, we need to send 10^{100} photons over the optical fiber, which is a lot of photon. As a comparison, if we convert the entire 10^{53} kg mass / energy of the universe into photons, we will only get 10^{88} photons. Even if we use squeezed light, we still need 10^{50} photons, which is still far higher than the digital requirement.

3.3.2   Low-Voltage Differential Signaling (LVDS)

The Spectron Spy relies heavily on the Low-Voltage Differential Signaling (LVDS) standard. LVDS protocol was first introduced in 1994 in order to increase communication speed between devices. Some advantage of LVDS compared to single ended signaling:

- **Differential**. In the PCB layout, differential signal trace pairs are routed together, as close to each other as possible, while in a long transmission line, the wires are typically twisted continuously along the line. This configuration greatly reduces the loop area exposed by the wire pair, which in turns reduces the amount of RF noise picked up by the cable pair. Since the signal has much lower noise than typical single ended signals, the overall differential voltage can be significantly reduced.

- **Low-Voltage**. The difference in voltage between the pairs is small (350 mV). This is much smaller than the TTL signal level of 5 Volt and the CMOS signal level of 3.3 Volt. The significantly reduces differential termination power requirement in typical $120\,\Omega$ transmission line from 200 mW in TTL to 1 mW in LVDS. Note that low voltage also implies low current and low power. The low common mode voltage of the differential pair of 1.2 V also allows direct communication with newer low voltage ICs.

- **Less radiation**. This advantage derives directly from the previous two advantages. First, the twisted pair nature of the cable tend to cancels any RF radiation emitted by the cable. Second, since LVDS uses significantly less current than TTL signal, the amount of radiation emitted is also reduces.

- **Ground loop free**. A typical differential signal cable has a separate ground shielding around them. The differential signal pair cables themselves are never

connected to ground. This eliminates any possibility of ground loop since there is no ground involved in the signaling.

- **Symmetric**. The cable pair in LVDS is symmetric. There is nothing inherently different about the positive and negative wire. We only have to swap these wire in order to convert a binary 0 into a binary 1. This is in contrast with a single ended signal where one wire is ground while the other is the signal. The overall signal tends to change towards binary 0 if there is significant signal loss along the cable.

- **Widely used**. Not surprisingly, LVDS is a very popular protocol used by many commonly known standards such as SCSI, Firewire, SATA, PCI Express, and Hyper Transport. This means that there are plenty of ICs that support LVDS protocol.

3.3.3   Ethernet Cable

We decided to use Ethernet cable for the following reasons:

- Ethernet cable is widely used so it is very cheap.

- Compared to other widely used cables such as USB and Firewire cables, Ethernet cable has more wires which increase the flexibility of the daughter board.

- There are 4 twisted pair wires inside an Ethernet cable which makes it ideal for our LVDS signals. Ethernet cable is available at all possible length up to 100 meters which is far longer than alternative cable (USB or Firewire).

- Ethernet cable is available with all possible shielding including individual pair shielding with outer double shielding. Theoretically, this should greatly reduce cross talk between pairs while also reducing RF noise transmission from the cable.

In our test run, we can reliably use 50 MHz SPI clock rate for most cable length, and 25 MHz SPI clock rate for 50 meters and longer.

3.3.4   Serial Peripheral Interface Bus (SPI)

One popular serial communication protocol is the Serial Peripheral Interface Bus. The protocol has one master with multiple slaves. The SPI bus specifies the following signals

- SCLK. Serial clock, generated by the master.

- CS. Chip select. Generated by the master to select a particular slave device while at the same time define the timing of the communication.

- MOSI. Master Output Slave Input. Used to send data from the master to the slave.

- MISO. Master Input Slave Output. Used to send data from the slave to the master.

The Spectron Spy protocol has 3 bi-directional general purpose signals. This means that the Spectron Spy can support unidirectional SPI protocol but not bi-directional SPI protocol. This is actually good enough for most purposes since most slave devices are either an input only device (TTL input, analog input) or an output only device (TTL output, analog output).

3.4   Spectron Spy (Daughter Boards)

Each Spectron Spy daughter board is connected to a Spectron Master Spy board via an Ethernet cable.

A typical Spectron Spy board is shown on the figure above with the following signal path:

- Signal comes from the left from the Ethernet cable

- The LVDS signal is converted to 3.3V CMOS signal using a three channel LVDS to CMOS converter chip. This chip is chosen with the signal direction in mind. For example, an Analog Output board uses 3 output converters, while an Analog Input board uses 2 output converters and 1 input converter.

- The 6V power is converted to a 5V power using a linear regulator

- Both the data and the power signal pass through the ADUM chips. One of the ADUM chip floats the power transmission while another one floats the data from the left part of the board to the right part of the board.

- Once the power and signal is on the right side of the board, we can put whatever circuit we need for this particular board. This part of the board is fully floated and ground loop free due to the ADUM chip.

Some remarks about the awesome Analog Devices ADUM chip:

- The ADUM chip is a single chip solution to isolates the power with minimal footprint. This greatly reduces the need of external power supplies for isolation purposes. A different chip from the same ADUM series isolates the digital data line.

- The ADUM chip isolates the power using an on chip transformer running at 180 MHz. This high frequency is required in order to transfer significant power through the small on chip transformer. This frequency is significantly higher than typical switching power supply frequencies which are usually below 1 MHz. This is good because theoretically, the output switching ripple noise should be greatly reduced. Most circuits also operate at much lower frequency making this high frequency noise

irrelevant to those circuits. Furthermore, in principle, filtering this high frequency noise with low pass filter should be relatively easy. Note however that in practice 180 MHz noise tends to radiate easily. Therefore, it is probably a good idea to follow the board layout recommendation from Analog Device to minimize this radiation.

- The chip uses pulse width modulation at 635 kHz in order to control the output voltage. This frequency range is far more relevant for both the circuits that we want to isolate and also for atomic physics experiment in general. For sensitive analog application, it is important that this noise is filtered properly. We found that ferrite beads at both the input and output pins significantly reduces noise at this frequency level.

The next subsections describe several notable Spectron Spy daughter boards that we have in our experiment.

3.4.1 TTL Input - Output

This is one of the simplest Spectron Spy boards. The board simply converts the three LVDS signals into TTL input or output. Since we digitized most of the electronics, we only use TTL signal for legacy applications where we have to control equipment that takes TTL signal.

3.4.2 Analog Output

As the name suggests, this board generates a floating 18-bit analog output signal. The board has built in 1 ppm / $°C$ voltage reference and output buffer with output voltage range of 0 to 5 volts. We primarily use the board for controlling our high current power
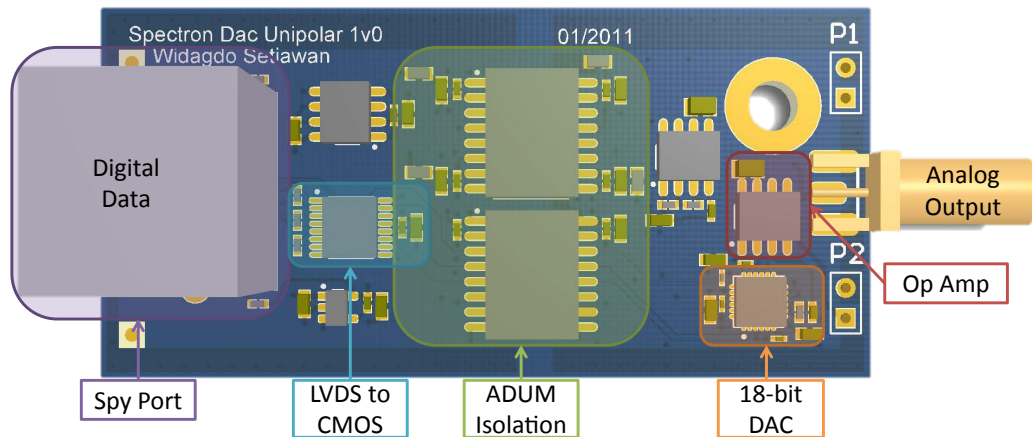
Figure 3.18: Spectron Dac. This Spectron Spy board generates 0 - 5 V signal with 18-bit resolution at 2 MSPS.

supplies and the piezos in our laser diodes. In order to maintain the 1 ppm / $°C$

performance, it is important to design the circuit without any resistive voltage divider

since the temperature coefficients of standard resistors are rather high, 20 ppm / $°C$. [2]

The important parts for this boards are:

- DAC: DAC 9881

- Voltage reference: MAX 6350

- Op amp buffer: Opa 211

3.4.3 Analog Input

This board converts analog signal from -10 Volt to 10 Volt into 14-bit digital signals. We

use this board to read the lithium laser spectroscopy signal in order to lock the laser to

the reference. Another use is to read the current probe signal for stabilizing the current in

---

[2]Note that we do not need to use the analog output board as PID loop setpoint or as laser power control setpoint since our setpoint system is purely digital.
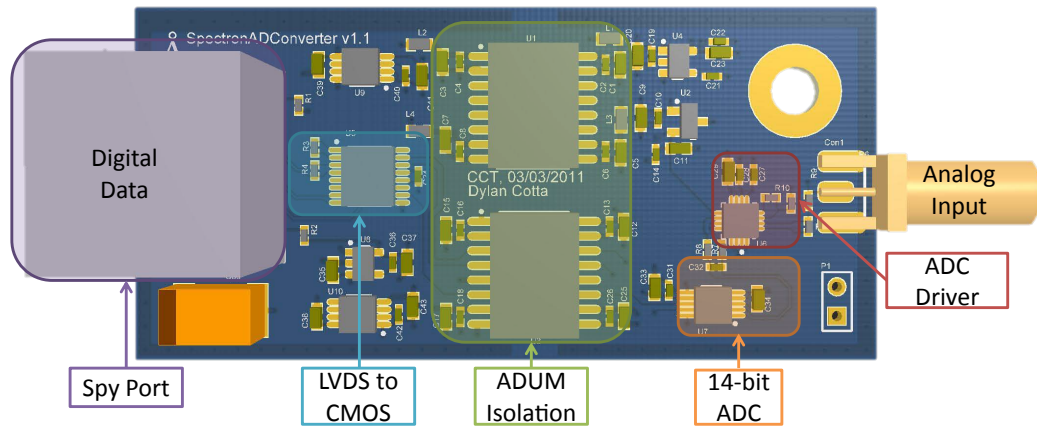
Figure 3.19: Spectron Adc. This Spectron Spy board digitize -10 Volt to + 10 Volt analog signal into 14-bit digital signal at 2 MSPS.

the coils. One upgrade that can be made with this analog input board is to use higher resolution 16-bit AD converter instead of the 14-bit one.

3.4.4 Thermocouple Reader

This board has 8 thermocouple input ports for simultaneous reading 8 temperature signals. This is primary used for temperature interlock of the magnetic coils. We also use this board to monitor the temperature during the vacuum chamber baking and also for interlocking the baking system from over temperature.

3.4.5 Interlock Watchdog / Mini Fuse

This board is used to drive other equipment that has to be interlocked, for example IPG laser, RF amplifiers, and coil power supplies. The board has a watchdog timer that makes sure that the FPGA is still running. The FPGA can stop working for example when the clock input stops or is running at a wrong frequency. The watchdog chip then safely shuts
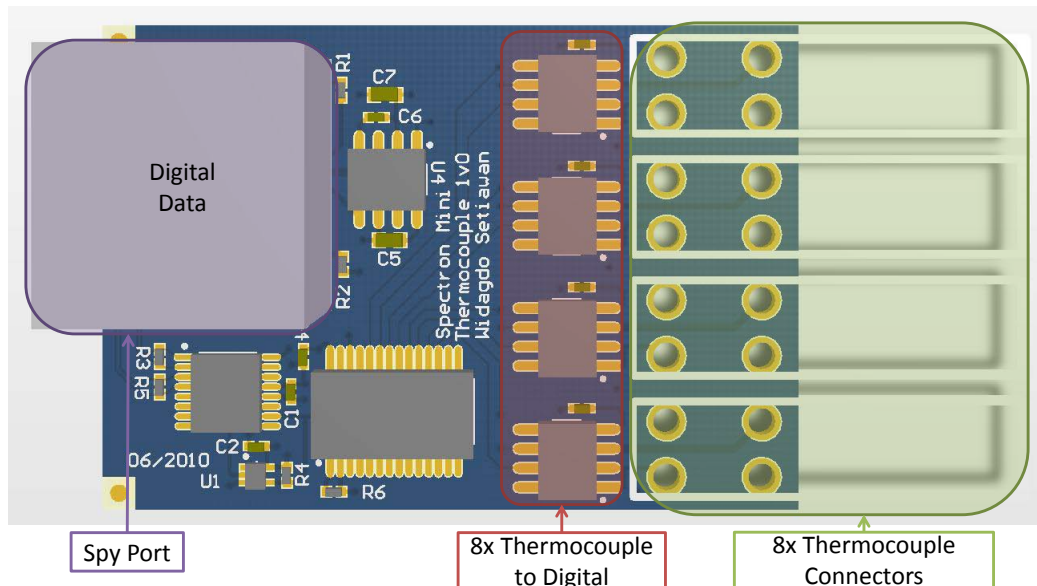
Figure 3.20: Spectron mini thermocouple accepts 8 mini thermocouple connectors. This board is used to monitor coil temperature and also to monitor the temperature zones during vacuum chamber baking process.

down the equipment that is interlocked.

3.4.6 Frequency Counter

This board is used to count frequencies from 1 MHz to 3 GHz. A typical application is to count the frequency for beat-locking two laser beams. This technique allows us to continuously change the beat-lock frequency over multiple gigahertz range.

3.5 Symphonia

By far the most complicated Spectron board is the Symphonia board. The project consists of two sub-boards, the motherboard and the DDS board, that are stacked together to form the complete circuit.
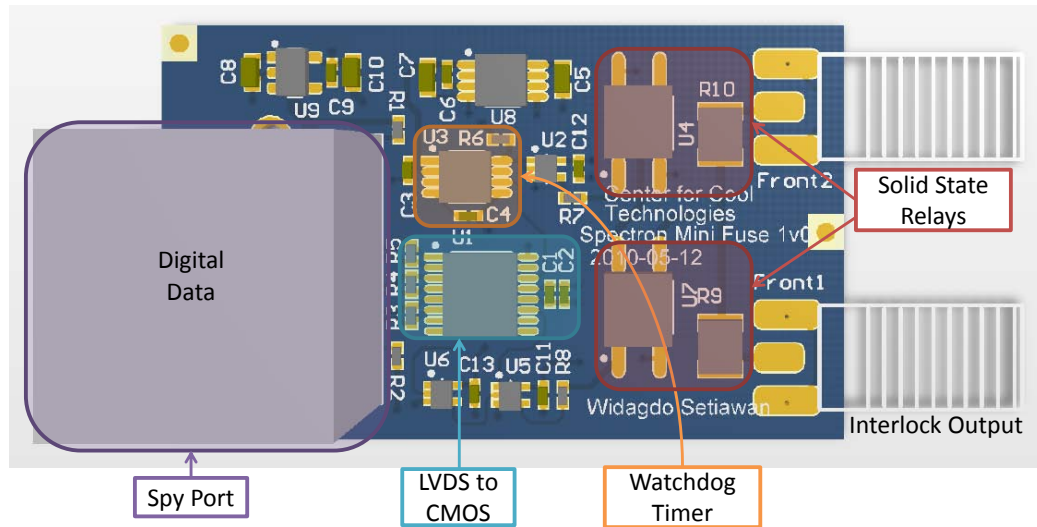
Figure 3.21: Spectron mini fuse is an integral part of the interlock system. With a complicated system like our FPGA, we do not want to solely rely on the FPGA to produce the correct output at all times. One failure mode that comes to mind is the failure of the clock signal to the FPGA, making the FPGA stuck in "Interlock OK" mode. This boards make sure will independently shut down the interlocked devices when this happens.

The motherboard has 2 analog outputs and 6 analog inputs at the front panel. The analog input, in combination with the DDS board, can be used to stabilize the power of a laser. The motherboard also has several switching power supplies which generates the voltages needed by the DDS board.

The DDS board, designed by Florian Huber, has 4 phase-locked RF outputs that are generated by 4 AD9910 DDS chips. The DDS board has the following features:

- 4 RF outputs, each with the following components:

  - AD9910 DDS chip, 1 GSPS, 14-bit, 2 us parameter update rate [73]

  - Exponential amplifier to attenuate/amplify the RF signal

  - High order low pass filter to reduce higher harmonics and nyquist mirror image

- Synchronization chips to phase synchronize the state machine inside the 4 DDS
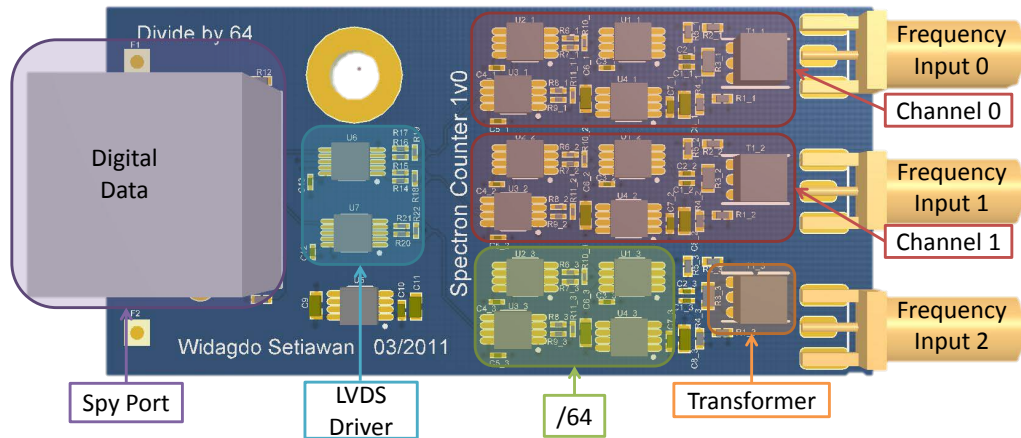
Figure 3.22:   The Spectron Counter is used to count frequencies from 1 MHz to 3 GHz which is very useful for offset locking lasers.

chips.

- Crystal oscillator at 100 MHz as a stable frequency reference source for the 2 GSPS DDS clock signal. This crystal oscillator is then locked to the lab-wide atomic clock signal.

Combined with the motherboard AD converters, the Symphonia board can be used directly to stabilize laser power without additional components.

3.6   Blastia

Another critical component of the control system is the software, called Blastia. Although in the grand scheme of things the software only acts as a user interface to the Spectron framework, a good user interface makes a big difference in terms of usability and productivity of the user. Blastia is designed with enough abstraction layers that the top layer, which is the interaction layer with the end user, is completely decoupled from the Spectron system.
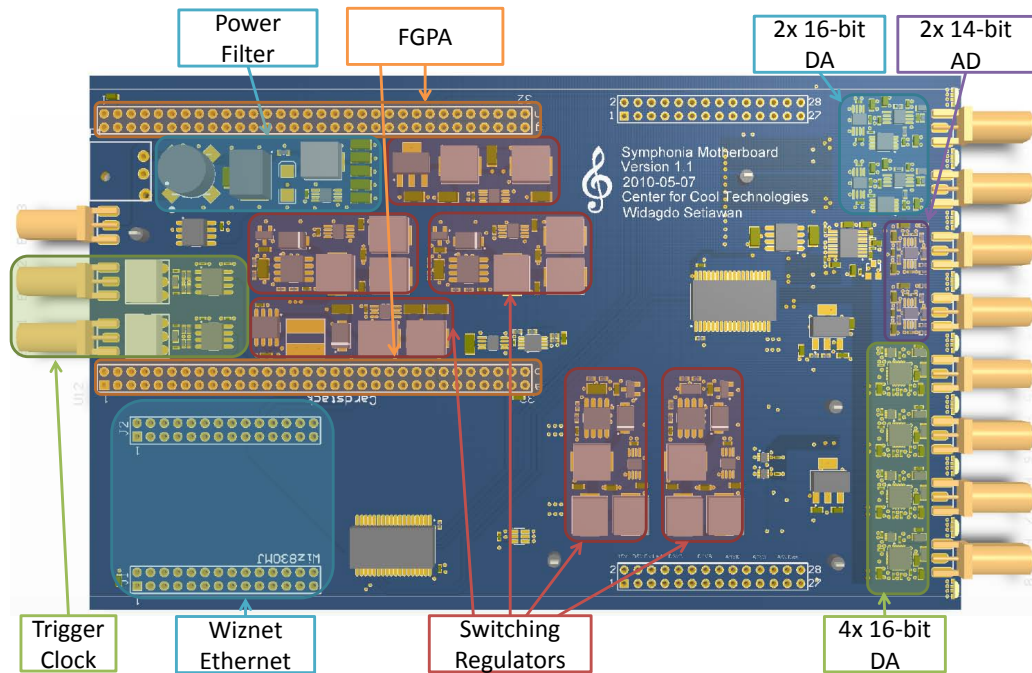
Figure 3.23: Symphonia Motherboard is equipped with 2 DA converters and 6 AD converters. The motherboard also generates the necessary voltages to be used by the Symphonia DDS board. The voltages are generated using 5 switching regulators running in spread spectrum mode. Switching regulators and spread spectrum are explained in Appendix A.7.

3.6.1 Experiment Sequence Code

The experiment sequence code is a source code that describes how each channel should behave as a function of time. The sequence code itself is a full Turing complete programming language (C#) which gives the user a lot of power in describing the experiment sequence.

The following is an actual snippet of our experiment code.

```
1  if (gdtv.IsFeshbachAdiabaticRampDown)
2  {
```
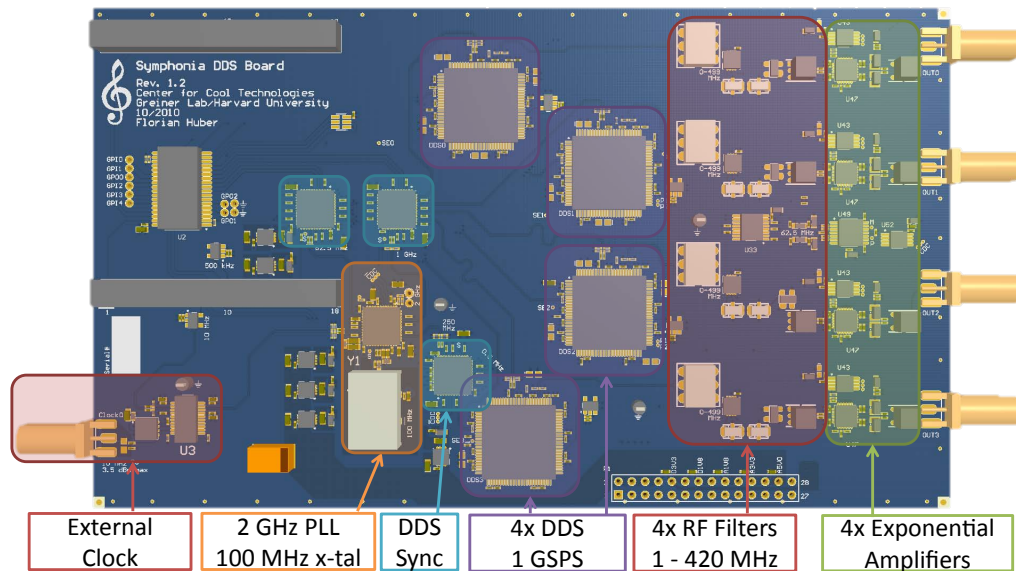
Figure 3.24: Symphonia DDS generates 4x 1 - 420 MHz phase locked signals. The board accept a 10 MHz reference clock signal that is used to lock an internal crystal oscillator running at 100 MHz. This 100 MHz signal is then PLL up to 2 GHz. The 2 GHz signal is then distributed to the 4 DDS chips. Two other synchronization chips provide IO Update and Sync signals to synchronization the 4 DDS chips. There is an exponential amplifier at the output of each DDS chip to provide the output with high dynamic range.

```
3    c.MainChamberFeshbachCoil.Feedback.ClosedLoopSetPoint.AddRamp(

4        currentTime, gdtv.FeshbachAdiabaticTimeScale, "0 A");

5    c.AdResync.AddPulse(currentTime, "1 ms", true);

6    currentTime += gdtv.FeshbachAdiabaticTimeScale;

7    c.MainChamberFeshbachCoil.Fet.AddEdge(currentTime, false);

8    c.AdResync.AddPulse(currentTime, "1 ms", true);

9  }

10 else //Non adiabatic ramp down

11 {

12   c.MainChamberFeshbachCoil.Feedback.ClosedLoopSetPoint.AddRamp(

13       currentTime, gdtv.FeshbachRampDownTime, "0 G");
```
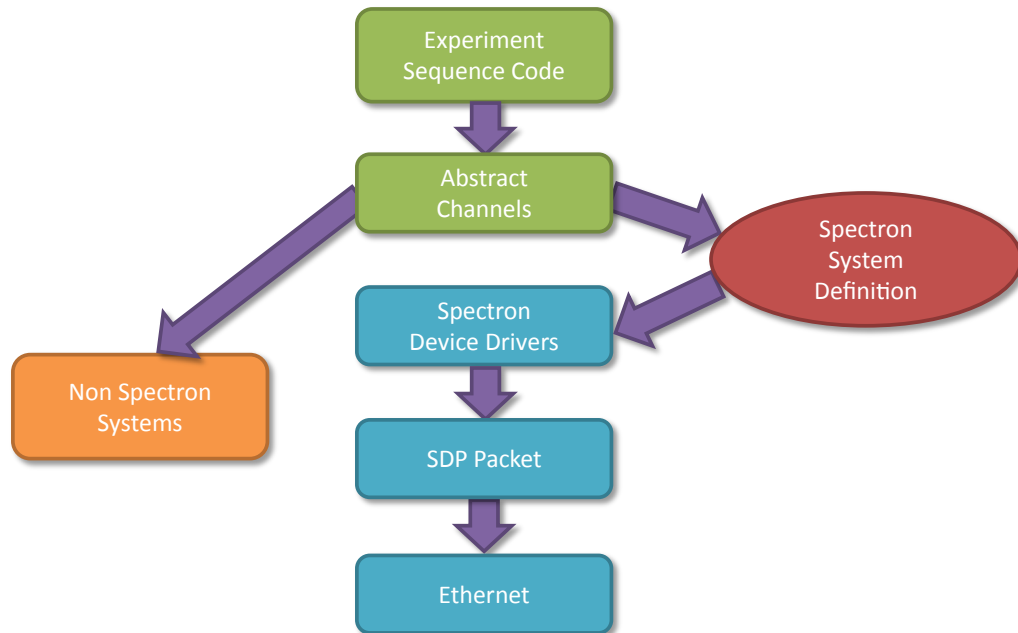
Figure 3.25: Blastia abstraction layers. The user writes the experiment sequence in C#. The code is then compiled in to abstract channels. Spectron device drivers then grab the appropriate abstract channel as assigned by Spectron System Definition, and compile the abstract channel into appropriate SDP packets. These SDP packets are then sent to the FPGA to be stored in the SRAM in preparation for the experiment sequence. Other non Spectron system can also take these abstract channels for their own use. Some non Spectron system in the lab are: camera, motorized mirror mounts, and air bearing translation stage.

```
14    c.MainChamberFeshbachCoil.Fet.AddEdge(

15        currentTime + gdtv.FeshbachRampDownTime, false);

16    c.AdResync.AddPulse(

17        currentTime + 0.9 * gdtv.FeshbachRampDownTime,

18        2 * gdtv.FeshbachRampDownTime + "100 us", true);

19 }
```

I want to draw your attention to several aspects of the code:

- The system fully supports variables. Some variables in the code:

IsFeshbachAdiabaticRampDown, FeshbachAdiabaticTimeScale, FeshbachRampDownTimem, currentTime. The variables can be changed easily in the user interface so that the behavior of the experiment can be change easily without touching the experiment code. Another option is to scan one or more of the variables automatically.

- Conditional code. The if statement on line 1. Here, the code branches depending if we want to ramp down the Feshbach coil adiabatically, or as quickly as possible.

- Full dimension and unit support. Expression such as "100 us", "0 G", "0 A" are fully supported. The software will tell the user when a number with a wrong dimension is used, such as attempting to set laser power to "0 Gauss". Furthermore, there is automatic unit conversion as illustrated in line 4 and line 13. In line 4, we set the current to "0 A", while in line 13, we set the magnetic field to "0 G". The software knows how to convert magnetic field into current for that particular coil.

3.6.2   Abstract Channel

A channel is an abstract object that can have values that vary in time. Example of channels include: Feshbach magnetic field (in unit of Gauss), DDS frequency (in Hertz), laser power (in Watts), and wave plate rotation angle (in degrees or radians). At this stage, a channel is abstract in a sense that it does directly map to anything in the experiment. From the channel's point of view, it does not matter how the magnetic field is generated or how the laser power is stabilized. The channel only describes what the end result should be.

This strict abstraction layer is useful since over the course of the experiment, the way we

generate magnetic field / laser power / frequency could change over time while we still want the same magnetic field regardless of implementation detail. Consider the following scenario. Initially, we generate the magnetic field by setting the current directly on the power supply. Later on, it turns out that build in the current stabilization circuitry in the power supply produces significant noise in the magnetic field. To solve this problem, the current is now stabilized by using some external feedback circuitry. Due to the new feedback circuit, controlling the current on the power supply directly is no longer possible. Without this abstraction layer, all of the experiment sequence code that drives the current of the power supply directly has to be rewritten. However, since the abstract channel only describes the magnetic field, all of the experiment sequence code can stay. The only thing that we have to change is the mapping between abstract magnetic field and the way to generate that magnetic field.

The abstract channels, in turn, are used by other subsystems such as the Spectron System Definition which maps the abstract channels to Spectron device drivers or other subsystems such as camera drivers and Aerotech airbearing stage drivers.

3.6.3 Spectron Device Drivers

Each Spectron device in the FPGA is represented by a Spectron Device Driver on the computer side. This abstraction layer hides the implementation detail on the FPGA side from the rest of the software. The Spectron Device Drivers take the abstract channels and convert the abstract channel into proper data format to be used by the FPGA. The device drivers then generate SDP Packets and sends those packets through the Ethernet to the Spectron boards.

3.7   Possible Upgrades

There are two notable systems that can be significantly upgraded. First, the FPGA can be upgraded to a much larger and faster version. Second, the Symphonia board can be upgraded increase the frequency range significantly.

3.7.1   FPGA

Since we are using Dallas Logic modules, upgrading the FPGA should be as simple as replacing the modules with newer modules. At the time of writing, the state of the art economical FPGA is the Cyclone 5 which is significantly better than the Cyclone 3 in terms of speed and capacity, while the best Dallas Logic module is the CMC1003 Cyclone 4 module with the following capabilities:

- 115k logic elements

- 3888k bits internal FPGA RAM

- 64M x 16 DDR2 SDRAM

With large FPGA, we can utilize the full capability of the DDS chip including the real-time phase control. We can also create a more general purpose Spectron Spy board source code that can be reconfigured without reprogramming instead of using different code for different Spectron board. Finally, the large DDR2 SDRAM can be used to store long waveforms both for sequence generation and oscilloscope.

3.7.2   Symphonia Microwave

Currently, the Symphonia frequency output is limited to 400 MHz. Although this is high enough for most application, there are times when we need to generate higher RF and microwave frequencies. It is possible to increase the frequency range significantly while maintain most of the advantages of DDS system by using single sideband generation[74]. Some possibly useful components are: ADL 5386 (quadrature modulator), ADF 4350 (PLL synthesizer with integrated VCO), and LTC 6946 (low noise integer-N synthesizer).

3.7.3   Noise Reduction

Most of our electronics described in the previous sections can be considered as "prototype" grade as we have not optimized the performance of any of the boards. If in the future it turns out that the noise generated by any of the boards limits the experiment in any way, I am confident that there will be clear and easy path to improve the noise performance of the circuits considerably.

One example is that the Phase-Locked Loop (PLL) on the Symphonia board was never optimized. With proper optimization, the servo bump could possibly reduced significantly. The Spectron Spy Analog Output board was optimized to minimize PCB board size and not for minimal noise. We can put more power filter on the analog side to reduce the power supply noise if required. Another method is to disable the clock / data transmission when there is no value change instead of sending the same data over and over again. This method would completely eliminate the digital clock noise.

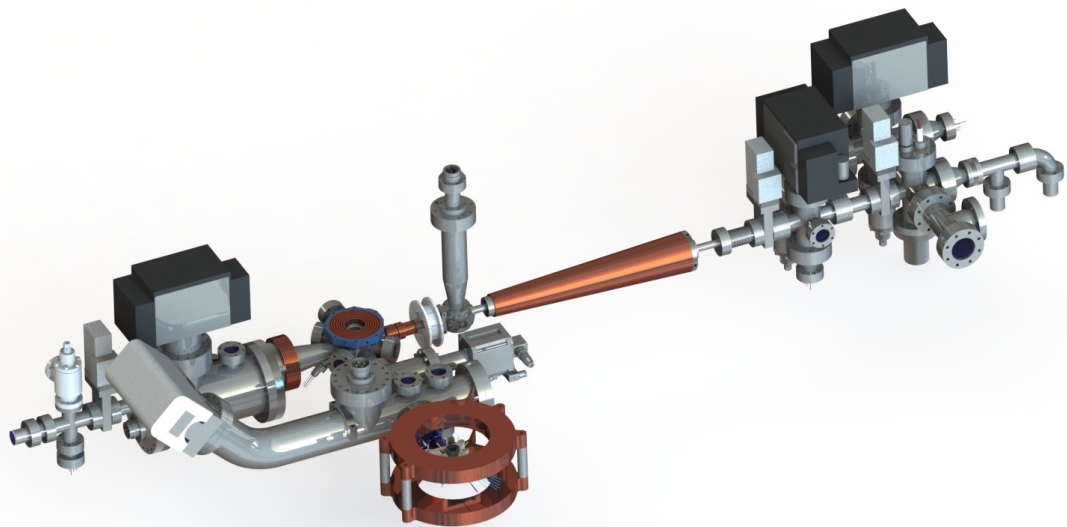# Chapter 4

Experiment Setup

Figure 4.1:   A perspective view of the experimental setup. This chapter explains the experimental setup in greater detail.

This experiment was originally designed to be a sodium - lithium experiment[75, 76]. The idea was to sympathetically cool the fermionic lithium atoms with the bosonic sodium atoms in order to obtain a large degenerate Fermi gas sample[52]. However, due to the success of the Rubidium Quantum Gas Microscope experiment[77], we decided to switch direction and focus on obtaining similar single-site resolution imaging capability but with fermionic lithium atoms.

Florian Huber's diploma thesis[76] describes the early vacuum chamber, Zeeman slower, and optical design from the atomic beam source up to the MOT Chamber. I will only briefly review those descriptions here while mentioning some changes that we have made since that time. Instead, I will focus more on new additions that we have in the experiment since that time.

For the rest of this chapter, I will adopt the following convention:

- The atoms in the atomic beam are moving from the oven towards the positive x-axis

- The atoms are transported from the main chamber towards the glass cell. This transport direction is the positive y-axis

- The Feshbach magnetic field points at the positive z-axis, which is the vertical axis.

4.1   Vacuum Chamber and Magnetic Coils

The vacuum chamber was optimized for a sodium - lithium experiment and thus not optimized for lithium only experiment that we are currently pursuing[1]. However, since we

---

[1]For an overview of an optimized lithium only experiment, see for example, Selim Jochim's lithium experiments[78, 79].
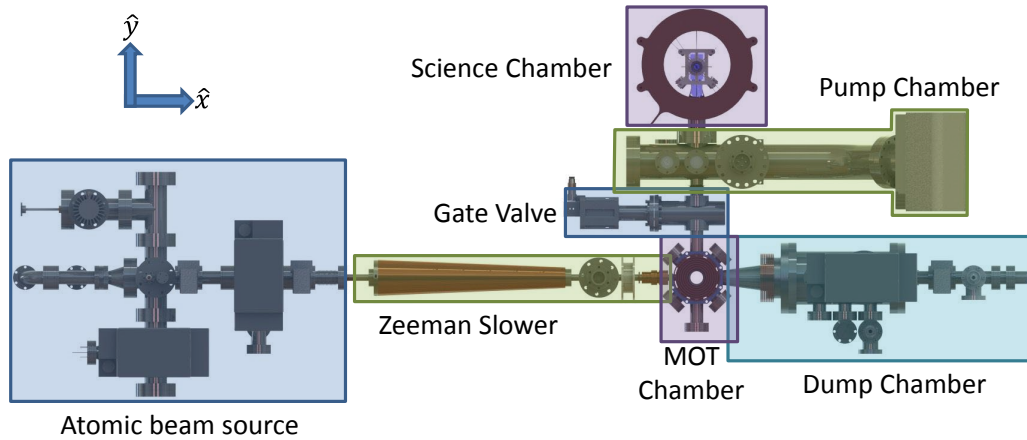
Figure 4.2: The experiment setup is logically separated into different sections: atomic beam source, Zeeman slower, MOT chamber, dump chamber, pump chamber, and science chamber (glass cell).

are interested in creating 2D gasses instead of 3D gasses, we do not need to cool too many atoms per experiment run, so efficiency should not be a big problem.

The logical divisions of the vacuum chamber is shown in Figure 4.2. The atoms start at the atomic beam source at $450°\,$C. These atoms are then slowed down to $1\,$K by the Zeeman slower and captured and further cooled by the MOT to $1\,$mK. The atoms are then moved towards the science chamber / glass cell where they will be further cooled before the physics experiment.

4.1.1 Atomic Beam Source

The atomic beam source setup closely follows the setup designed by Claudiu Stan from Wolfgang Ketterle's group at MIT[80, 81] as shown in Figure 4.3. The lithium atoms start at $450°\,$Celsius on the oven in gaseous state. Most of the lithium atoms are filtered through a series of apertures so that the lithium atomic beam is mostly traveling straight towards the Zeeman slower and the MOT chamber. The lithium beam can be blocked
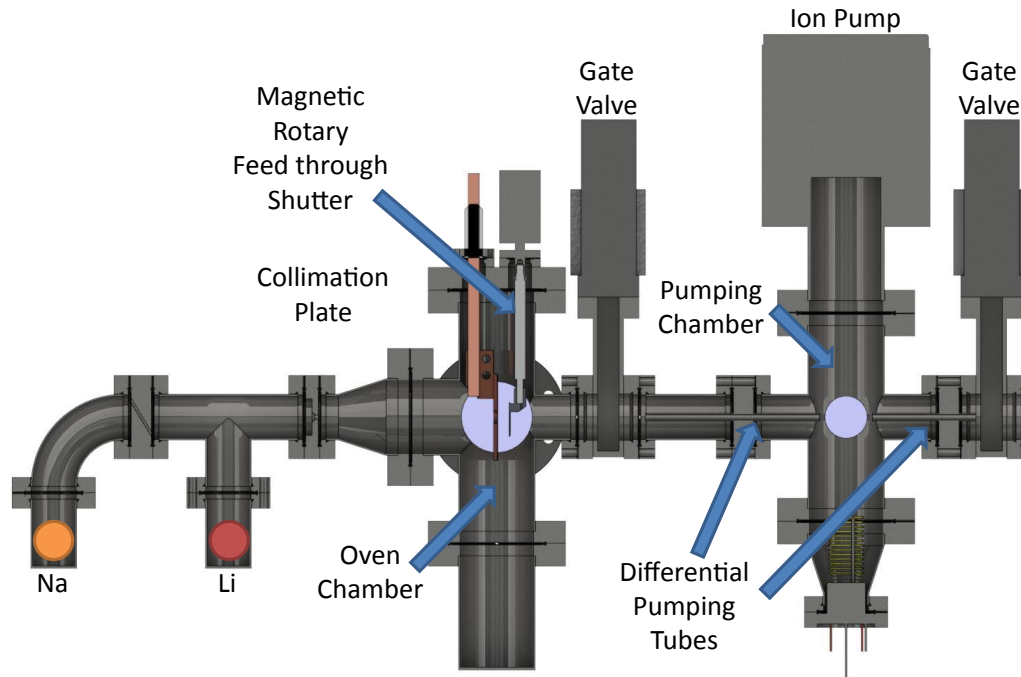
Figure 4.3: The atomic beam source is an almost identical copy to Claudiu Stan's design[80, 81]. The only difference is that we use magnetic rotary feed through atomic beam shutter instead of linear actuator.

using mechanical shutter once the beam is no longer required.

One change that we made from Stan's design is that we are using a magnetic rotary feed through for the atomic beam shutter instead of bellow based linear feed through. Stan mentions in his thesis that the bellow on the pneumatic linear stage is not too reliable for repetitive use[80]. The magnetic rotary shutter is chosen so that we do not have to rely on any bellow for the mechanical motion[2].

From the oven chamber, the lithium beam enters the differential pumping chamber. This chamber acts as an intermediate chamber between the higher pressure oven chamber and the low pressure MOT chamber. This additional chamber allows us to have 4 orders of

---

[2]Another option is to use magnetic linear feed through that also does not rely on any bellow.

magnitude difference in pressure between the oven chamber and the MOT chamber.
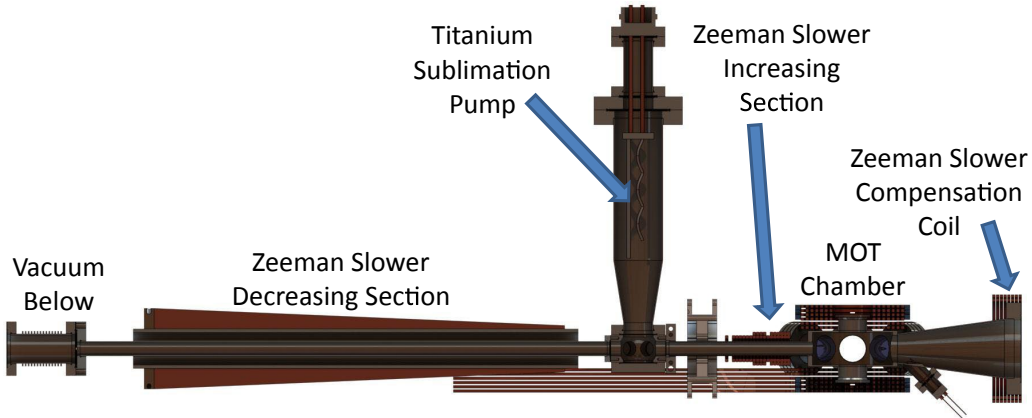
4.1.2 Zeeman Slower

Figure 4.4: The Zeeman slower is a spin-flip Zeeman slower similar to [82].

After the differential pumping chamber, the atoms enter the Zeeman slower, shown in Figure 4.4, where the atoms are slowed and cooled simultaneously. The Zeeman slower that we have is a standard spin-flip Zeeman slower which was optimized for maximum sodium slowing performance[82]. It turns out that lithium atoms can sustain higher acceleration rate compared to sodium. Since we are using lithium instead of sodium, the electrical current on the Zeeman slower coils are modified to increase the lithium loading rate while sacrificing the sodium loading performance.

4.1.3 MOT Chamber

From the Zeeman slower, the slowed and cooled atoms arrive at the MOT chamber shown in Figure 4.5. Here, a Magneto Optical Trap (MOT) cools and captures the atoms from the atomic beam. The MOT has finite capture velocity range so it is important that the

Figure 4.5: The MOT chamber is a simple chamber with 10 ports. The chamber was initially designed for magnetic transport[83] which greatly constrained the vertical dimensions of the chamber.

Zeeman slower slows down the atoms enough so that their velocity is below the MOT capture velocity.

Coil

The magnetic wire for the coil is a 4x4 mm square cross section wire with 2.5 mm circular diameter hole inside. The coil dissipates a few kilowatts of heat during normal operation. Therefore, it is important that we cool the coil effectively. The cooling water flows through this inner bore. This way, the coil is in good thermal contact with the cooling water at all times.

There are two coils on the MOT chamber, the MOT coil and the Feshbach coil as shown in Figure 4.6. The MOT coil, the outermost pancake pair, is wired in quadrupole
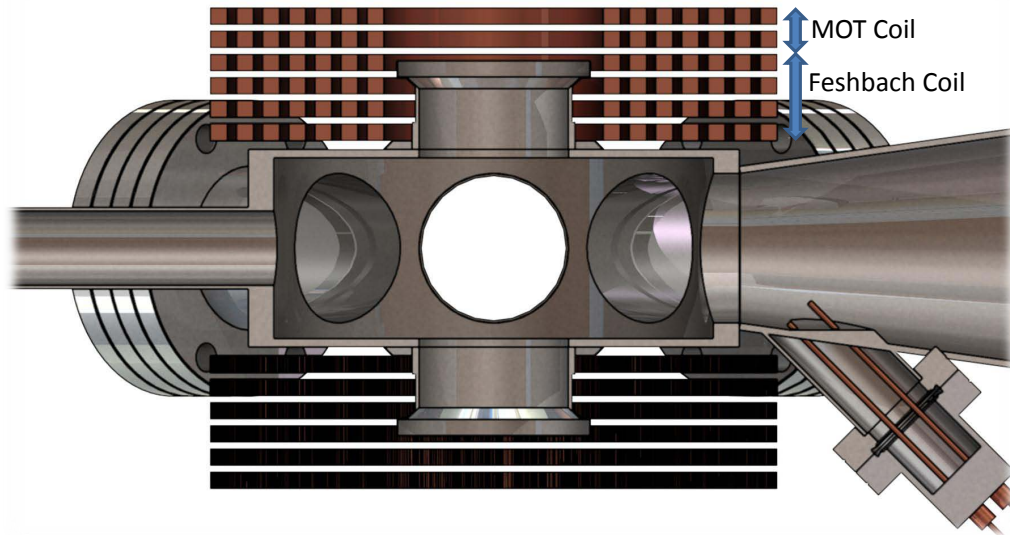
Figure 4.6: The MOT coil and Feshbach coil are mounted very close to the chamber. Since initially the coils were design to generate strong quadrupole field for magnetic trapping[84], the Feshbach field has large residual curvature. Coincidentally, this residual curvature is very useful for adding extra confinement on the atoms.

configuration in order to generate the necessary magnetic field for creating the MOT. The remaining two inner coil pairs are wired in Helmholtz configuration in order to create the required magnetic field for the lithium Feshbach resonance at $834\,$Gauss. The inductance of the MOT Coil is $13.5\,\mu$H and the inductance of the Feshbach coil is $35.2\,\mu$H.

4.1.4 Dump Chamber

The dump chamber is located on the opposite side of the Zeeman slower as shown in Figure 4.2. The purpose of the dump chamber is to pump the residual gas load from the unslowed atomic beam while pumping the MOT chamber at the same time. There are two pumps at the dump chamber: a Titanium sublimation pump and a $75\,L/s$ ion pump. These to pumps in combination reduce the pressure on the MOT chamber to less than $10^{-12}\,$Torr level. This low pressure reduces the amount of background gas so that the lifetime of

trapped atoms limited by the background gas collision is in the order of a few minutes.

4.1.5   Pump Chamber

The pump chamber is the chamber between the main chamber where we load the atoms and the science chamber / glass cell where we do our experiment. As the name suggests, the pump chamber is responsible for pumping the gas load from the glass cell. Similar to the dump chamber, the pump chamber also uses both Titanium sublimation pump and an ion pump. At the time of writing, the pump chamber is still not assembled.

4.1.6   Science Chamber / Glass Cell
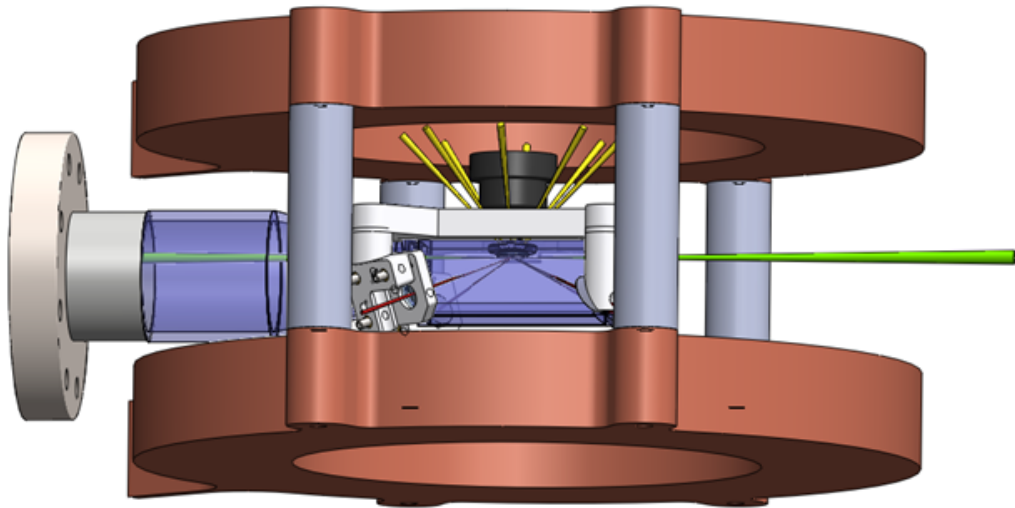
Figure 4.7: The glass chamber is sandwiched between two Feshbach coils. The objective is mounted directly on the glass cell to minimize relative movement between the objective and the glass cell. The retroreflectors, which reflects the red beams, are also mounted directly on the glass cell to minimize lattice misalignment (see Section 4.3.11). The green beam is the transport beam while the yellow beams are the Raman sideband cooling beams.

The science experiment will be done exclusively in the glass cell. A glass cell designed has

Figure 4.8:   The side view of the glass cell.

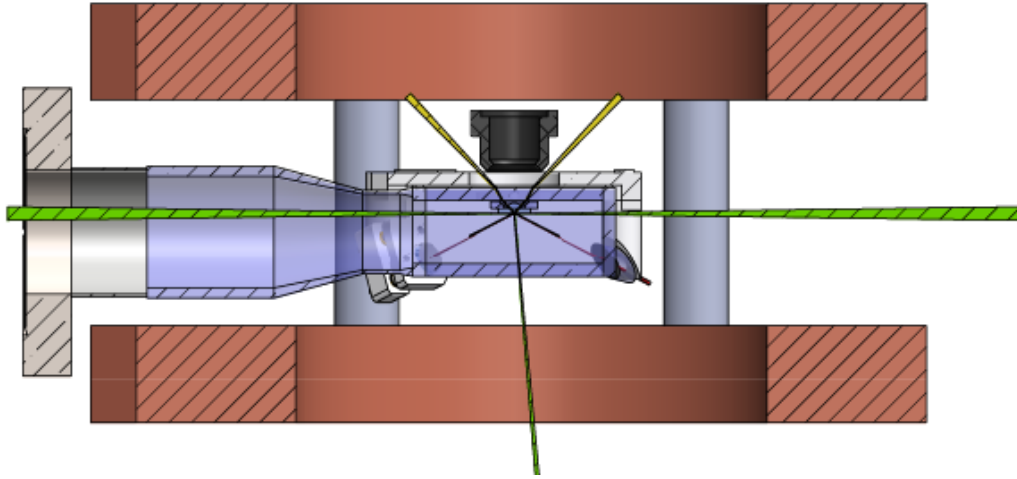several advantages over steel chamber design:

- More optical access. Since the entire cell is made out of glass, we can shine laser beam from a significant fraction of the solid angle from the atom

- No eddy current. Since glass is an insulating material, there is no eddy current induced while we are switching the magnetic field on off. This is important since there are several instances where fast switching of magnetic field is required.

Glass cell chamber also have several disadvantages. First, glass cell is far more fragile than stainless steel chamber. It is usually very hard to mount a glass cell at more than two points without using stress relief bellow. Second, the glass cell has to be mounted to the rest of the chamber which is mostly made out of stainless steel. The glass to metal transition on the glass cell usually cannot take too much differential thermal expansion between the glass and the stainless steel. However, we feel that the advantages outweigh these two disadvantages. The in-vacuum optical components in the glass cell are described

later in the experiment sequence section. At the time of writing, the glass cell has not yet been assembled on the experiment.

Part of the glass cell setup is shown in the figure. The glass cell is mounted from the left to the stainless steel chamber. The glass cell is asymmetrically sandwiched around the Feshbach coil. Note that the atoms are located exactly in between the two coils. We can clearly see that the glass cell give us optical access to the atom from a significant part of the solid angle.

Coil
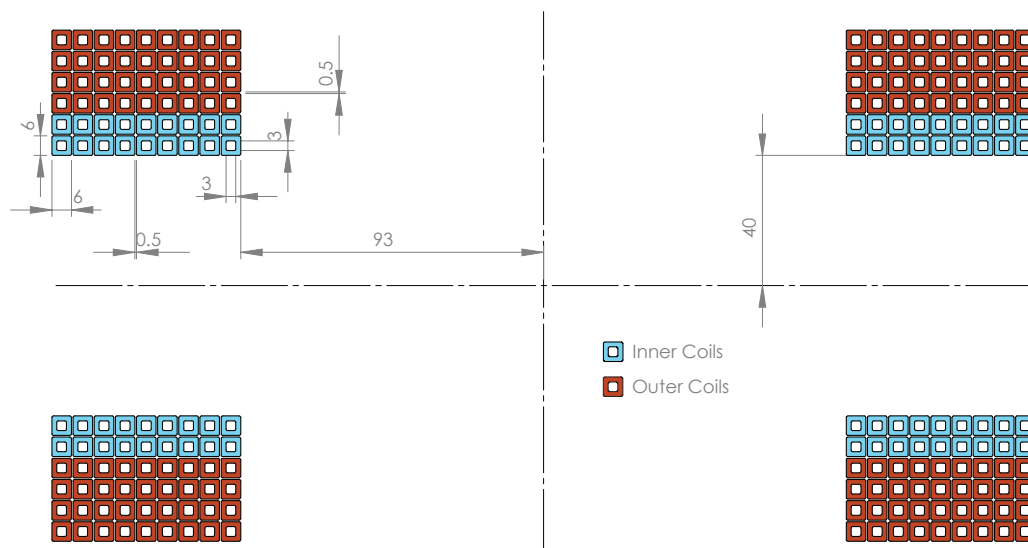
Figure 4.9: The cross section of the science chamber wires. All numbers are in millimeter. The atoms are located exactly in between the two coils.

The coils in the glass cell are made using similar technique as the coil in the MOT chamber. However, the coil is significantly larger for the following reasons:

- The larger the coil, the more uniform the generated magnetic field which is somewhat important for the experiment. In reality, the coil is significantly larger

Property	Value
Outer cross section	$6\,\mathrm{mm} \times 6\,\mathrm{mm}$ square
Inner cross section	$3\,\mathrm{mm} \times 3\,\mathrm{mm}$ square for water cooling
Insulation	Two layers of overlapping Kapton tape
Total insulation thickness	$0.06\,\mathrm{mm}$
Annealing	Soft temper for easier winding

Table 4.1: The wire specifications for the science chamber coil.

than it has to be if we only care about magnetic field homogeneity.

- Large coil allows us to have larger optical access to the atoms in the glass cell. The atoms in the glass cell are not located symmetrically on the z-axis of the glass cell because the location of the atoms is constrained by the optical component inside the glass cell. This constraints the vertical dimension of the coil significantly.

- The coil is designed in Helmholtz configuration so that there is no residual confinement at the atom location. Helmholtz configuration dictates diameter of the coil to be twice the distance between the coil. This makes the horizontal dimension of the coil very large.

We purchased the wire for the science chamber coil from S&W Wire Company. The specification for the wire is shown on Table 4.1. The dimensions and the scaled coil cross section of the coil are shown in Figure 4.9. As currently designed, the expected specification of the science chamber coil is shown in Table 4.2

4.1.7   Coil Winding

We decided to wind our own coil for the science chamber. The double pancake coil configuration is tricky to wind since we need to wind the two layers simultaneously. We

Property	Value
Magnetic field	$4\,{}^{\text{Gauss}}\!/_{\text{Ampere}}$
Water flow rate	$5\,{}^{\text{liters}}\!/_{\text{minute}}$
Total wire length	$83\,\text{meters}$
Total wire mass	$20\,\text{kg}$
At Feshbach resonance:	
Current	$210\,\text{Amperes}$
Voltage	$11\,\text{Volt}$
Magnetic field	$840\,\text{Gauss}$
Power dissipation	$2{,}300\,\text{Watts}$
Wire heating	$< 4°\,\text{Celcius}$
At maximum:	
Current	$500\,\text{Amperes}$
Voltage	$26\,\text{Volt}$
Magnetic field	$2000\,\text{Gauss}$
Power dissipation	$13{,}200\,\text{Watts}$
Wire heating	$< 20°\,\text{Celcius}$

Table 4.2: Science chamber coil specification

constructed a coil winding tool shown on Figure 4.10 to help us wind the science chamber coil.

Before we wind the coil, we have to unwind enough wire from the wire spool to find the wire spot where that would be the center of the coil. One tricky part about handling square copper wire is that the wire tends to twist on its own axis. The Teflon holder is responsible for fixing the orientation of the copper wire while providing enough friction so that the tension of the wire in the coil section is reasonably high. The two winding arms are moved simultaneously in opposite directions so that the radial dimensions of the two layers are about the same at all time. It is important to apply superglue during the winding process itself to make sure that coil is properly bonded to itself. We also have to make sure to Teflon coat every single part that touches the wire itself because we do not want to glue the wire to the aluminum parts. After the super glue cures, we separate the
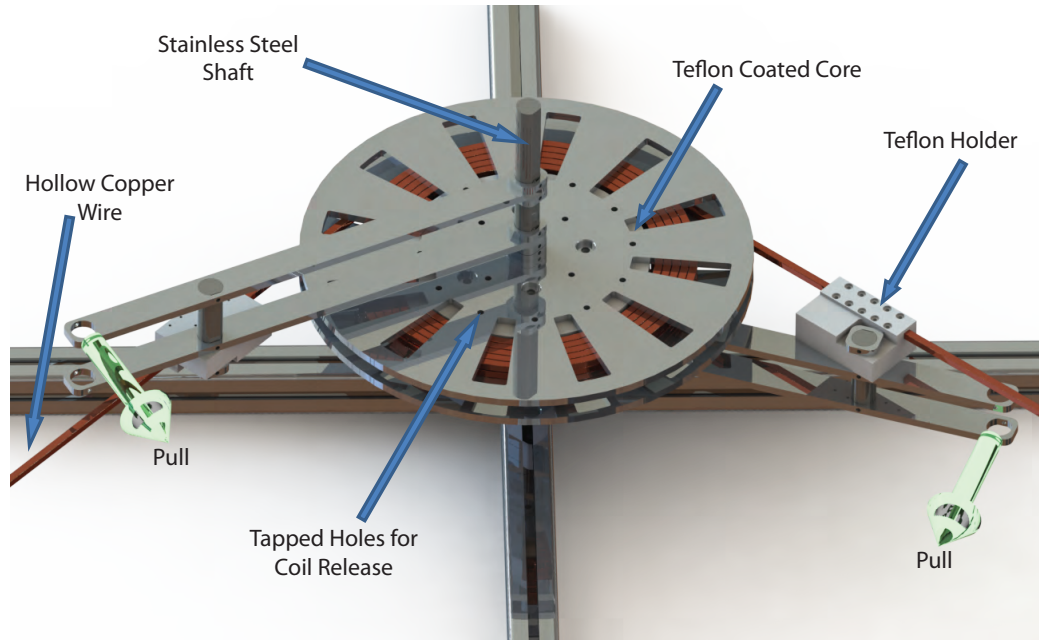
Figure 4.10: The two layers of the pancake are wound simultaneously in opposite direction using this coil winding tool. The teflon holders are responsible in orienting the square wire before the wire reach the coil.

coil from the winding tool by screwing screws on the tapped holes.

Coil Winding Handedness   There are two possible way to wind double pancake coil: left handed spiral and right handed spiral. Figure 4.11 shows the simulations result for the two possible case. Since we want uniform magnetic field at the atom location, it would make sense to arrange the top and bottom coils to have the same handedness.

4.2   Power Supplies and Coils

This section describes the power supplies that are required to power the magnetic coils. All of the high current power supplies that we use are the Genesys series power supplies from TDK-Lambda. We are using 8 x 1.5 kW power supplies with 3 x 15 kW power
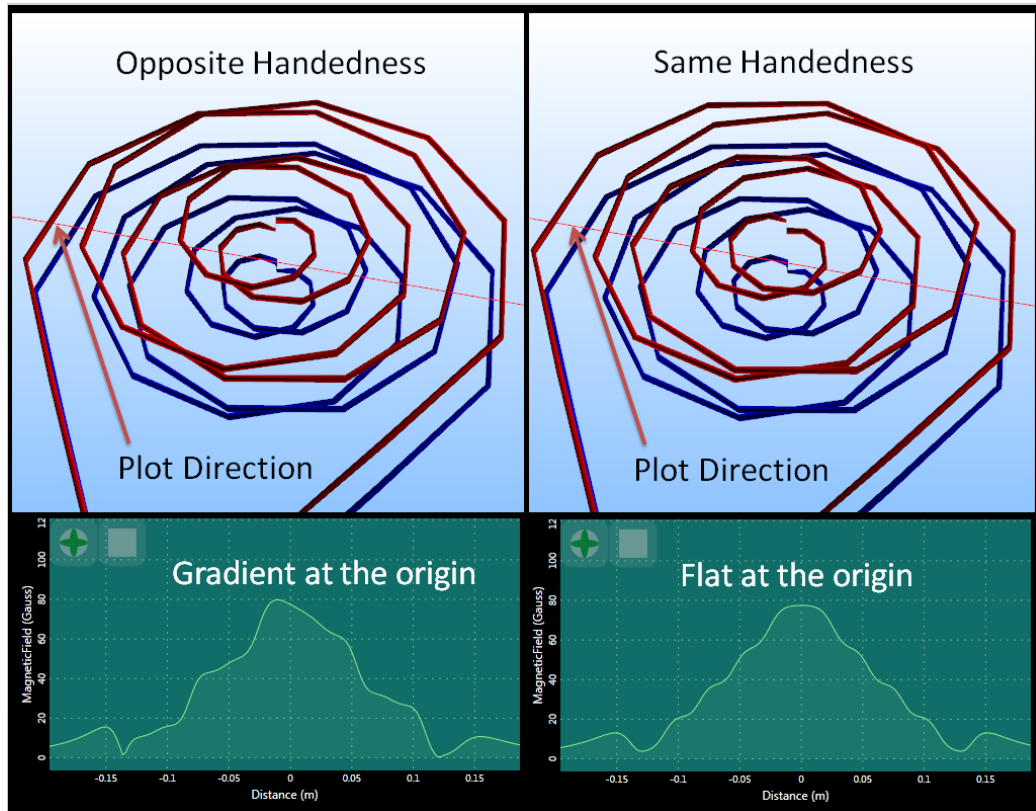
Figure 4.11: The image shows two simulations of highly exaggerated spiral coil winding configurations for Feshbach field. On the left, the top(red) and bottom(blue) coils have different handedness, left handed and right handed respectively. With this configuration, there is significant gradient at the origin. The right figure shows the two coils with the same right handedness. The magnetic field at the origin has no gradient.

supplies. The power supplies are located outside our lab for the following reasons:

- Minimize acoustic noise in the lab since acoustic noise tends to vibrate sensitive optical elements.

- Minimize heat dissipation source in the lab. Even at 88% power conversion efficiency, the power supplies would still dissipate 6kW of heat on the air at full power. Large heat source in the lab could cause unwanted temperature fluctuations which in turn would cause optical misalignments.

Part Number	Current (A)	Voltage (V)	Power (kW)	Quantity	Names
Gen 12-120	120	12	1.5	6	PS[0]..PS[5]
Gen 20-76	76	20	1.5	2	PS[6]..PS[7]
Gen 60-250	250	60	15	3	PSB[0]..PSB[3]

Table 4.3: We use these 11 power supplies to power our magnetic coils. The power supplies are shared between the MOT chamber and the science chamber. All power supplies are from TDK-Lambda High Power.

- Of course the real reasons are far more practical. We do not have space in the lab to house these huge power supplies and our electrical wiring in the lab is also not equipped to power these current hungry power supplies.

The power supplies are well decoupled from the coil both physically and electrically. Since the power supplies are located at some other location, we have to run 35 meters long 4/0 gauge wire from the power supply to the lab where each wire acts as a resistor. Overall, the resistance is actually a good thing since resistance generally increases the stability of a feedback system with minimal power dissipation. Note that for some coils, we are using two power supplies in parallel to power a single coil, so the increased feedback stability is very welcomed. We also use several Mosfets and IGBTs in the lab in order to change the wiring configuration between the power supplies and the coil in real-time. This way, we can share some power supplies between the MOT chamber and the glass cell. The power supplies specifications is shown on Table 4.3.

4.2.1   MOT Chamber Coil Wiring

The MOT chamber coil wiring is shown on Figure 4.12. I will describe the purpose of each component on the diagram. The resistors represent the 35 meters cable connecting the power supplies from outside the lab to the lab. The value of each resistor is $5.6\,\text{m}\Omega$. By
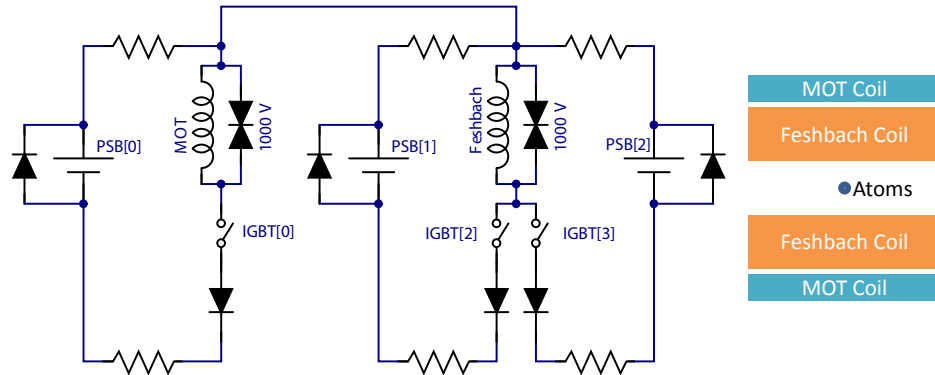
Figure 4.12: MOT Chamber Coil Wiring. The MOT coil is powered by one power supply while the Feshbach coil is powered by two power supplies connected in parallel for extra current.

arbitrary convention, the positive sides of the power supplies are connected together in the lab. We can treat this common connection as the ground, and assume that the power supplies produce negative voltage to power the magnetic coil. Note that this connection happens after the long wire or resistor. Theoretically, these resistors should help stabilize the feedback circuit in the power supply.

Each power supply is connected in parallel with a diode. This diode was installed as a safety precaution to protect the power supply against reverse voltage. This could happen on the following scenario. First, we turn on both PSB[1] and IGBT[2] at some high current. We then turn IGBT[2] off. This will induce high negative potential on the coil due to the intrinsic inductance of the coil. If we then turn IGBT[3] switch on, the PSB[2] power supply will be subjected to this high voltage. The diode is there so that the negative voltage build up on the power supply is limited by the bias voltage of the diode.

The diodes connected in series with the IGBTs are there to protect the power supply from high voltage. The diodes prevent current from flowing the opposite direction through the power supplies.

One critical component in the setup is the varistor. Each inductive coil is connected in parallel with a 1000 V varistor. The varistor is there to dissipate the energy or current quickly during fast magnetic field shut down. Consider the following scenario. First, we run 500 A through the Feshbach coil. At $t = 0$, both IGBT[2] and IGBT[3] are switched off simultaneously. Since the coil is really inductive ($35\,\mu$H), for a short moment, the coil acts as an ideal current source, which means that the coil will produce / take whatever voltage it needs to maintain the 500 A current. Without the varistor, the voltage will build up to some high value until that voltage exceeds the breakdown voltage of the IGBT. Although I have not verified this theory, it seems likely that the IGBT cannot take this kind of abuse too many times. With the varistor in place however, once the voltage across the coil reach the turn on voltage of the varistor, the current will start flowing through the varistor, protecting the IGBT in the process. The relation of the voltage of the coil is given by the inductor relation:

$$V = L\frac{\mathrm{d}I}{\mathrm{d}t}$$

For a fast current ramp down, it is important to have a high varistor turn on voltage. The 1000 V varistor is chosen to maximize this voltage while still safely below the 1200 V breakdown voltage of the IGBT. With the $35\,\mu$H inductance, we expect total ramp down time of $17\,\mu$s.

Note that the varistor has a finite lifetime. During fast current turn off, the varistor dissipates $P = VI = (1000\,\mathrm{V})(500\,\mathrm{A}) = 0.5\,\mathrm{MW}$, which is a very high power. At maximum energy dissipation scenario, the varistor could only survive around $100,000$ shots. In practice however, we rarely operate at maximum energy dissipation mode, so the lifetime of the varistor is expected to be significantly higher.
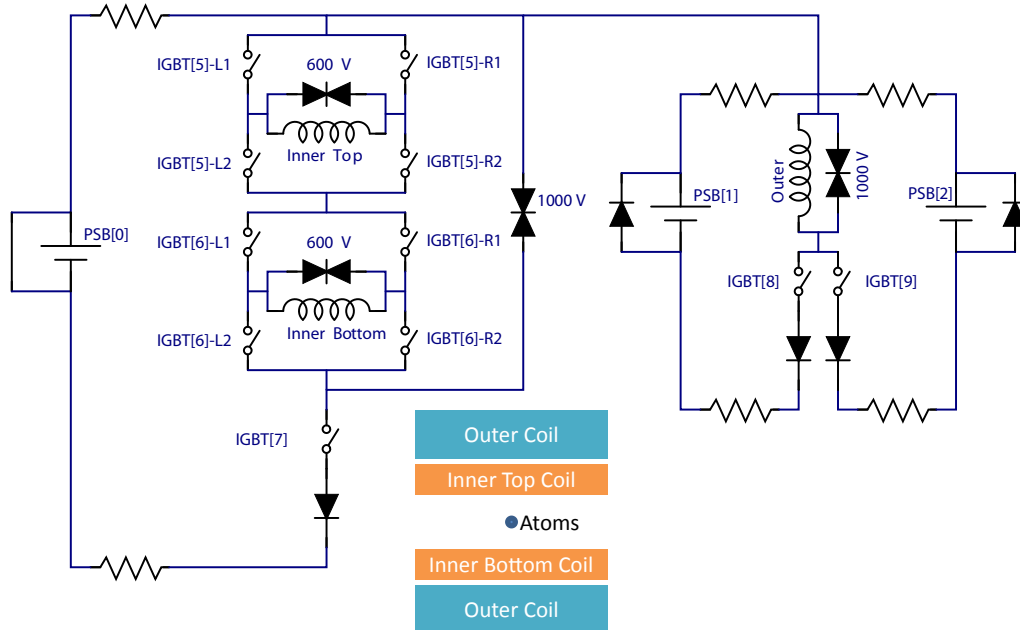
4.2.2   Glass Cell Coil Wiring

Figure 4.13:   Glass cell coil wiring. The glass cell coil is carefully designed such that when there is equal current in all coils, the magnetic field generated has zero residual curvature. The inner coil can also reconfigured independently in all possible permutations of current directions in order to create curvature or gradient on the atoms.

The glass cell coil wiring diagram is shown on Figure 4.13 above. The power supplies and the resistors are shared with the MOT Chamber coil wiring. The inner and outer coils are driven independently. The outer coils can only be turned on and off, while each of the inner coils is connected using H-bridge configuration in order to flip the direction of the magnetic field in real-time. The coil as a whole has the following modes of operations:

- **Helmholtz** mode. All coils generate magnetic field in the positive z direction. The dimensions of the coils are designed such that when the current on all wire are equal, the generated magnetic field profile has zero curvature.

- **Curvature** mode. The inner coils' field points in the negative z direction while the

outer coils' field points in the positive z direction. In this configuration, we can tune the curvature at the location of the atom with maximum radial trap confinement of 8.5 Hz while the axial direction is anti-trapped by 12 Hz. This configuration is useful if additional confinement is required either during the trap transfer process or during the physics experiments. Note that Maxwell equations forbid trapping configuration in all 3 dimensions for high-field seeking states.

- **Quadrupole** mode. Each of the inner coils produce magnetic field in the opposite direction. This creates a maximum linear gradient of 65 $^{\text{Gauss}}$/$_{\text{cm}}$ in the strong (axial) axis and 33 $^{\text{Gauss}}$/$_{\text{cm}}$ in the radial direction. In pure quadrupole configuration, the magnetic field at the location of the atom is zero in this configuration which is probably not so useful for physics experiment. We can simply add additional bias magnetic field using other coils as quantization axis or Feshbach resonance. Note that with strong bias field, the gradient of the Zeeman shift only exists in the direction of the bias magnetic field.

- **Top/bottom coil only** mode. In this configuration, we only turn on either the inner top coil or the inner bottom coil. This is actually similar to the quadrupole mode that it produces magnetic field and magnetic field gradient in z axis.

There are 9 IGBTs in the double H-bridge configuration. As currently planned, only IGBT[7] should be used for fast ramp down of magnetic field. When IGBT[7] is suddenly turned-off, the current will flow through the 1000 V varistor, dissipating the energy stored in the coils at the same time. The two 600V varistors are there to protect the IGBTs if there is accidental turn off in any of IGBT[5] and IGBT[6] during high current operations.

4.3 Experiment Sequence

This section describes the generic experiment sequence starting from the oven to the final imaging at the glass cell. At the time of writing, the glass cell is still yet to be assembled. All of the sequence up to an including the transport has been completed while everything after the transport would be implemented later.

The proposed experiment sequence is roughly the following:

- **Oven**. Atoms start at the oven chamber at $450°$ C. The nozzle allows the atom to propagate towards the main chamber direction. The two differential pumping tubes further select the radial velocity distribution of the atoms[80].

- **Zeeman Slower**. The axial velocity and temperature of the atoms are then reduced by the Zeeman slower. The final velocity of the atom at the end of Zeeman slower is in the order of $50\,m/s$, just slow enough to be reliably captured by the MOT.

- **The Magneto Optical Trap (MOT)** traps and further cool this atoms down to 1 mK.

- **MOT Compression**. The MOT is then compressed by reducing both the detuning and laser power in order to further cool and spatially overlap the MOT with the optical dipole trap (ODT). The temperature at this point is around $500\,\mu$K[78]. The atoms are then optically pump to the lower hyperfine state by shutting down the repumping light while increasing the power of the MOT light[52].

- **First evaporation**. The magnetic field is increase close to a Feshbach resonance to increase the collision rate between the two spin states which evaporates and reduces the temperature of the atoms.

- **Transport**. The atoms are then transported to the glass cell by moving one of the transport lens using a linear translation stage[85, 86].

- **Accordion lattice**. On the glass cell, $150\,\mu$m from the glass surface, the atoms are then transferred to the accordion lattice trap. This compresses the atoms along the transport axis.

- **Accordion transport**. The accordion lattice angle is then changed to compress the atom in the z direction while moving the atoms closer to the hemisphere surface at the same time[87]. An alternative is to use a slip-stick like transport scheme to bring the atoms further to the hemisphere surface.

- **Vertical lattice**. The vertical lattice beam is used to increase the confinement in the z direction, transforming the atomic sample into a 2D gas.

- **Dimple evaporation**. The atoms are further cooled using dimple evaporation before the science experiment.

- **Horizontal lattice**. We then load the atoms into horizontal optical lattice for the science experiment.

- **Pinning**. Right before imaging, we increase the trap depth of all lattice beams to the maximum non-adiabatically in order to suppress tunneling and pin each atom into a single lattice site.

- **Imaging**. A Raman sideband cooling beam is then used to cool the atoms [67] while at the same time, a camera collects the florescent light from the atom through both the hemisphere and the objective lens[65].
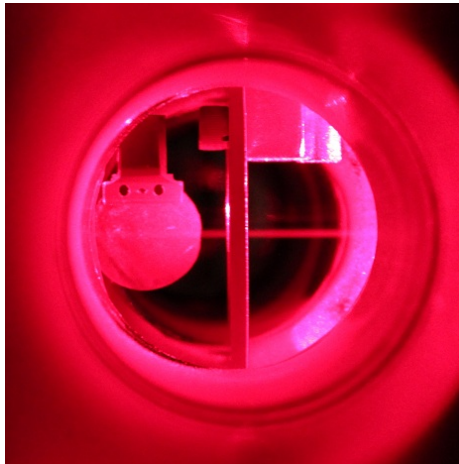
Figure 4.14:   Image of the atoms fluorescing the Zeeman slower laser beam. The rotary shutter on the left and the collimating plate on the center can also be seen.

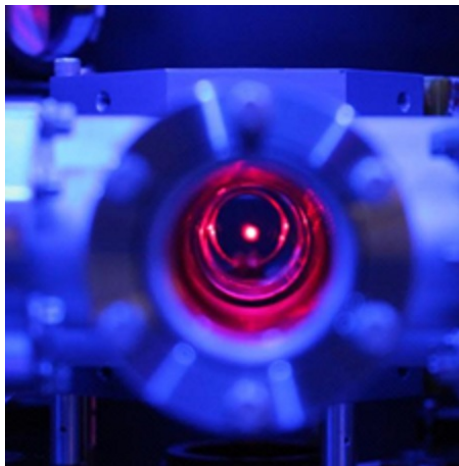### 4.3.1   Magneto Optical Trap (MOT)

Figure 4.15:   Image of the atoms in the Magneto Optical Trap (MOT).

The Magneto Optical Trap for lithium uses the following parameters:

- 3 independent laser beams that are retroreflected onto themselves

- Beam waist is 10 mm

- Each of the 3 beams consists of:

  - MOT light to trap and cool the atoms

  - Repump light to pump the atom back to the $3/2$ state

- Quadrupole field with $12\,{}^{Gauss}\!/_{\text{cm}}$ on the strong axis. The quadrupole field is increased during MOT compression stage to $20\,{}^{Gauss}\!/_{\text{cm}}$ over 20 ms.

- MOT loading $\approx 10^8$ atoms over $5\,\text{seconds}$

Lithium Laser   The lithium laser light for the Zeeman slower and the MOT is generated by three tapered amplifiers seeded by a grating stabilized diode laser. A detailed description of the setup can be found in Tobias Schuster's diploma thesis[75].
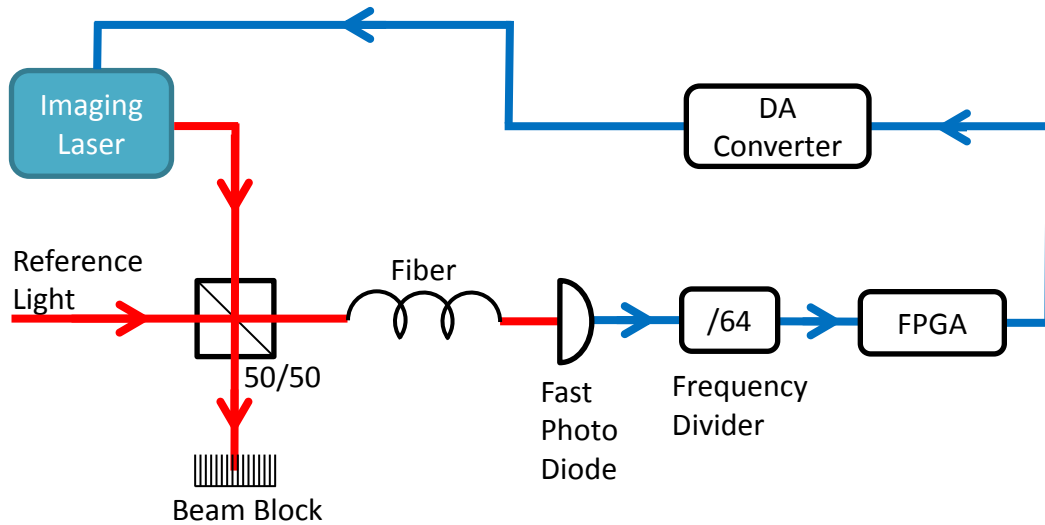

Figure 4.16:   Frequency offset lock schematic. This simple setup is capable of locking two lasers with $\pm 3\,\text{GHz}$ offset.

Offset Locked Imaging Laser   We use an offset locked laser diode for the absorption imaging light. This imaging light is only used for diagnostic purposes since the main single
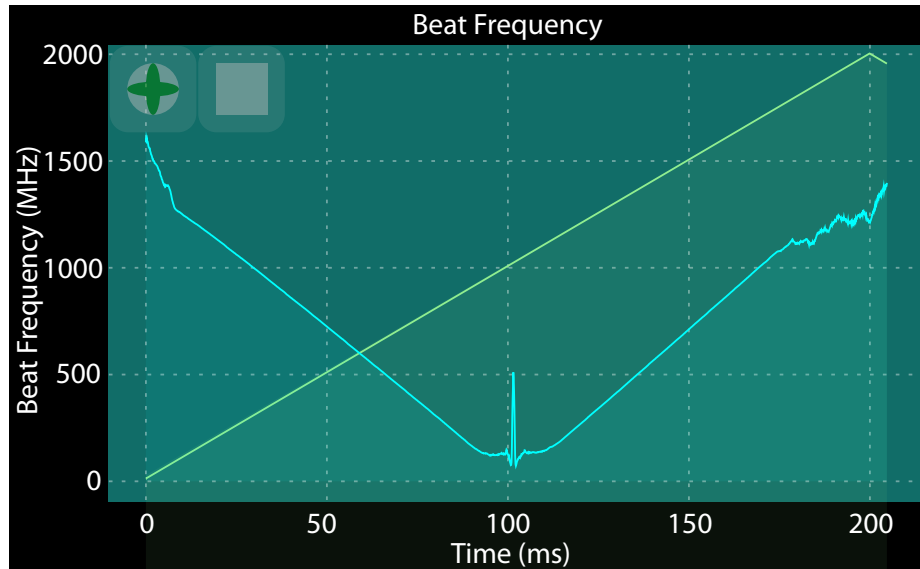
Figure 4.17: Typical beat lock signal. The green line is the piezo scan voltage of the laser diode grating. The blue line is the frequency readout spanning from negative frequency to positive frequency range. We can lock the laser at any point on the blue line except for the center region with low frequency.

lattice site imaging system uses different imaging technique. The offset locking is crucial since we want the capability of imaging the lithium at different magnetic fields. Our offset locking scheme is slightly different from the typical homodyne phase lag method.

The offset lock scheme is conceptually simple. First, the imaging laser is mixed with the reference light from the MOT laser setup on a 50/50 beam splitter. The mixed beam is then coupled to a single mode fiber. The single mode fiber is there for practical reason since we want the full spatial mode overlap between the two laser beams. The light coming out of the fibers is modulated by the difference in frequency between the imaging laser and the reference light. This modulation frequency is then measured by the fast photo diode. Up to this point, the setup is similar to standard offset lock setup.

The electrical signal from the photo diode is then amplified and treated as digital / binary signal. The signal goes to the divide by 64 frequency divider, which is basically just a

counter. The output of the divider is connected to the FPGA using Spectron Spy protocol. The reason for the divider is that the FPGA is running at 100 MHz clock, so the FPGA can only measure frequency up to 50 MHz. This division by 64 makes the effective upper frequency range to be higher than 3 GHz. Overall, the FPGA can then read beat lock frequency ranging from −3 GHz to +3 GHz.

Once the FPGA measures this beat frequency, the FPGA then uses a digital feedback system to feedback on the imaging laser to achieve the desired beat frequency. Theoretically, this method allows (almost) continuous tuning of frequency over 6 GHz with the exception of a few MHz region where the beat frequency is close to zero. A typical frequency readout signal is shown in Figure 4.17.
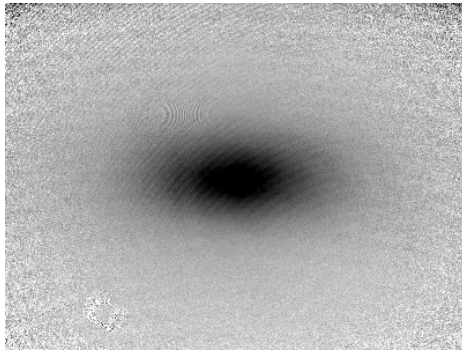
4.3.2   Checkpoint: Magnetic Trap

Figure 4.18:   Absorption image of the atoms in the magnetic trap.

Although not part of the main experiment sequence, we can trap the atoms from the MOT on the magnetic trap simply by turning the MOT beam offs and ramping up the magnetic field, creating a quadrupole magnetic trap. The atom absorption image is shown in Figure 4.18.

4.3.3   Optical Dipole Trap (ODT)

Figure 4.19:   To scale contour plot of the ODT / transport trap.

From the MOT, the atoms are then loaded to the optical dipole trap. The dipole trap is created by a single $54\,\mu$m waist, $250\,$W laser beam which is turned on during the compression stage of the MOT. The transfer efficiency from the MOT to the dipole trap is approximately 2%. The ODT creates a trap with frequency of $11\,$kHz $\times\ 50\,$Hz $\times\ 11\,$kHz , corresponding to trap aspect ratio of 220, with a trap depth of $2.5\,$mK. The shape of the trap can be visualized as a needle or a hair, but it should not be visualized as a blob as shown on Figure 4.19. The extreme trap aspect ratio is the side effect the the problem requirements to trap the spatially large and energetically hot MOT.
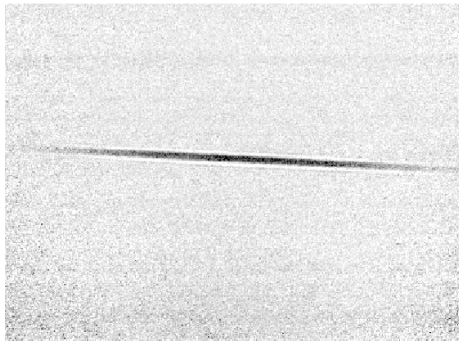

Figure 4.20:   Absorption image of atoms in the hybrid optical and magnetic trap.

At this point, the atoms are still rather hot (1 mK) and the ODT can only barely hold the atoms. These atoms would probably spill out during the acceleration stage of the transport process. Therefore, before the transport, it would be a good idea to evaporate some of the atoms by ramping up the magnetic field close to a Feshbach resonance [52]. The magnetic field also confines the atoms along the trap axis, further increasing the

collision rates. This increases the collision rate between the atoms and thermalizes the

atoms to a lower temperature. Figure 4.20 shows the absorption image the atoms in the

hybrid optical and magnetic trap.
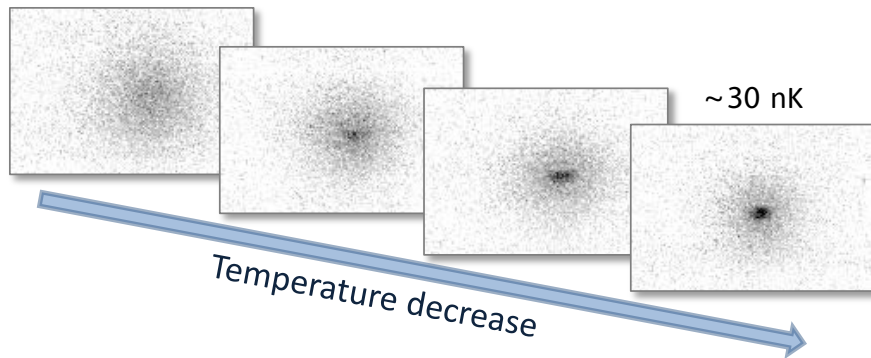
4.3.4 Checkpoint: BEC

Figure 4.21:   Evaporation of two spin states of lithium on Feshbach resonance. At the end
of the evaporation, a molecular lithium BEC is formed.

Once the atoms are in the ODT, we can use Feshbach resonance to do perform

evaporative cooling down to degeneracy. An absorption image of molecular lithium

BEC[88] is shown in Figure 4.21. We can also apply RF pulse to transfer the atoms into

different m_F levels as shown in Figure 4.22.

4.3.5 Transport

Once we load some atoms in the ODT, we have to move the atoms to the glass cell which

is achieved by translating the focus of the ODT to the science chamber[89, 90, 91, 86].

The optical transport setup is shown on Figure 4.23. The three lenses, f_1, f_2, f_3 forms a 6f

Fourier imaging system. A collimated beam with waist of around 5 mm comes from the
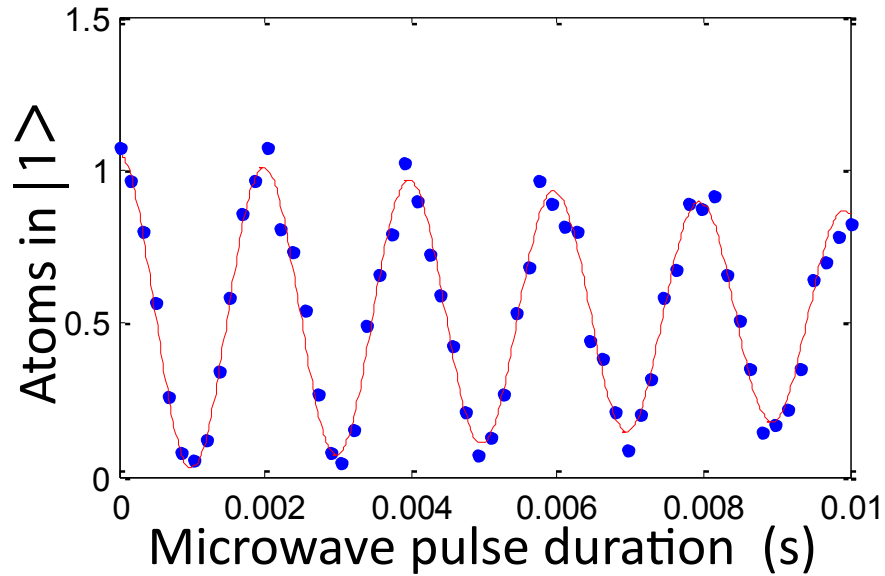
Figure 4.22:   Rabi oscillation of atoms between two spin states driven by RF / microwave pulse.

left. Initially, the atom is located at the focus of f_3. We then translate f_1 to the right in order to shift the final focus location as shown in Figure 4.24. Note that the transport beam is the same as the ODT beam.

The transport setup has the following characteristic:

- The last 2 lenses forms a $4f$, $2.5\times$ magnifying imaging system

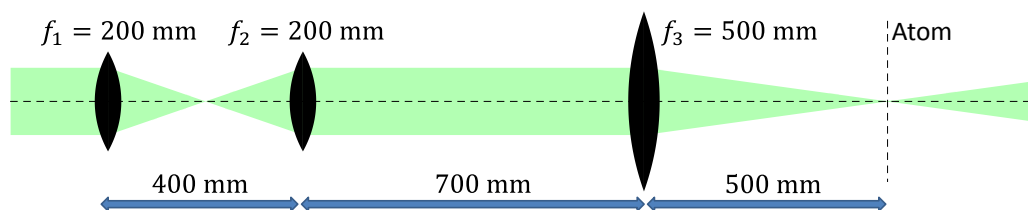Figure 4.23:   Transport lenses. The first lens is moved by a translation stage over $100\,\text{mm}$. This movement is amplified by 6.25 by the next two lenses which act as a telescope. The focus location at the atom can move by $625\,\text{mm}$.
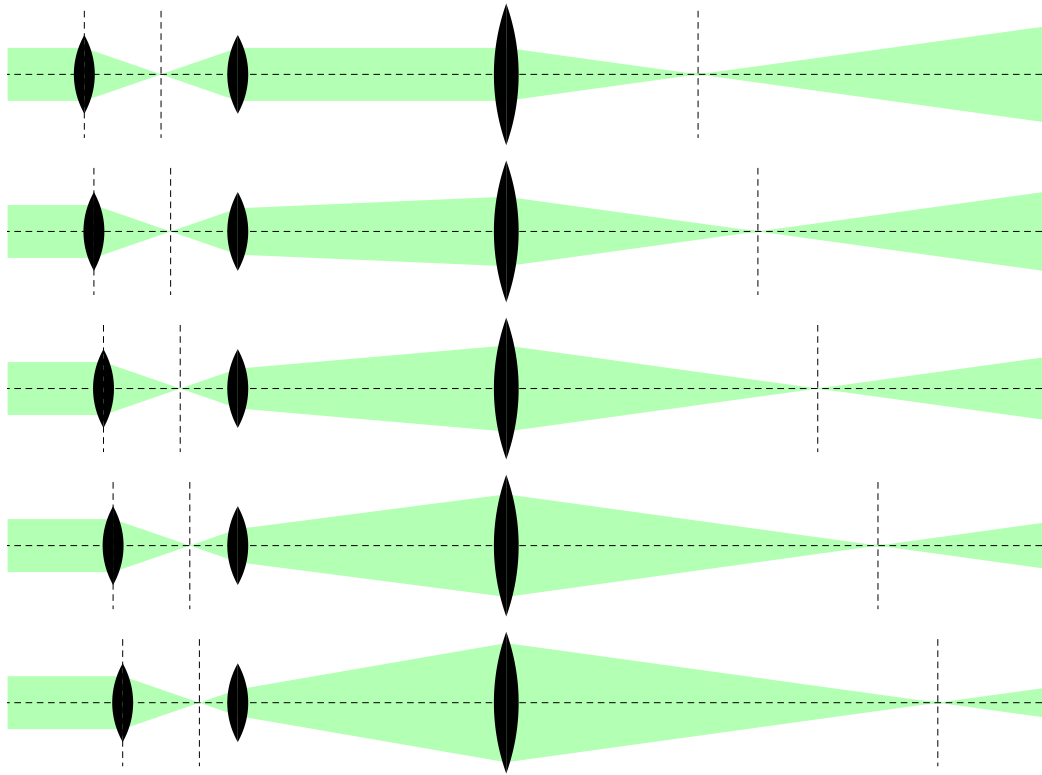
Figure 4.24: The Gaussian beam trace with different positions of the first lens. The movement of the final focus is far larger then the movement of the first lens. Although not visible in the diagrams, the waist of the beam at the atom is constant over the entire process.

- By moving the first lens to the right, the focus location after the last lens is moved by $2.5^2 = 6.25$ times due to the magnification setup. This can be understood intuitively using phase space conservation of the laser beam. Since the waist is reduced by 2.5 times, the Rayleigh range must increase by 2.5^2 time to conserve phase space.

- The waist after the final lens is constant during the movement of the first lens. This is a special property of a $4f$ imaging system. The first lens creates a constant waist regardless of the lens position located between f_1 and f_2. The $4f$ imaging then

images the curvature of this beam. Since a 4f imaging preserves curvatures, the curvature of the beam after the final lens is also preserved, thus keeping the final waist constant.

Note that although the current setup preserves the waist of the beam during transport, there is no fundamental reason why this has to be the case. By changing the distance between f_2 and f_3, we can for example compress the atomic cloud during the transport to better match the final transport configuration with the final trap while at the same time matching the trap size with the MOT for initial loading.

Transport Optics   Lens f_1 is mounted on an Aerotech ABL10100 air bearing linear translation stage with maximum travel range of 100 mm. The air bearing is used to minimize the amount of vibration during the transport process. The stage is driven by a linear servomotor to minimize vibration on the transport axis.

All three lenses f_1, f_2, f_3 are standard Thorlabs plano-convex singlet lenses. For reasons that I do not fully understand, OSLO simulation, confirmed by actual trial, indicates that these lenses perform better than fancier achromatic doublets for transport purposes.

Transport Ramp   We can think of the transport process as changing the trap location. Ideally, we want to change the trap adiabatically in order minimize the amount of heating in the trap[56]. There are several considerations:

- **Velocity**. According to the special theory of relativity[92], the atoms should not care about the velocity of the trap as long as the velocity is constant. Therefore, the velocity of the trap is just limited by maximum velocity of the translation stage.
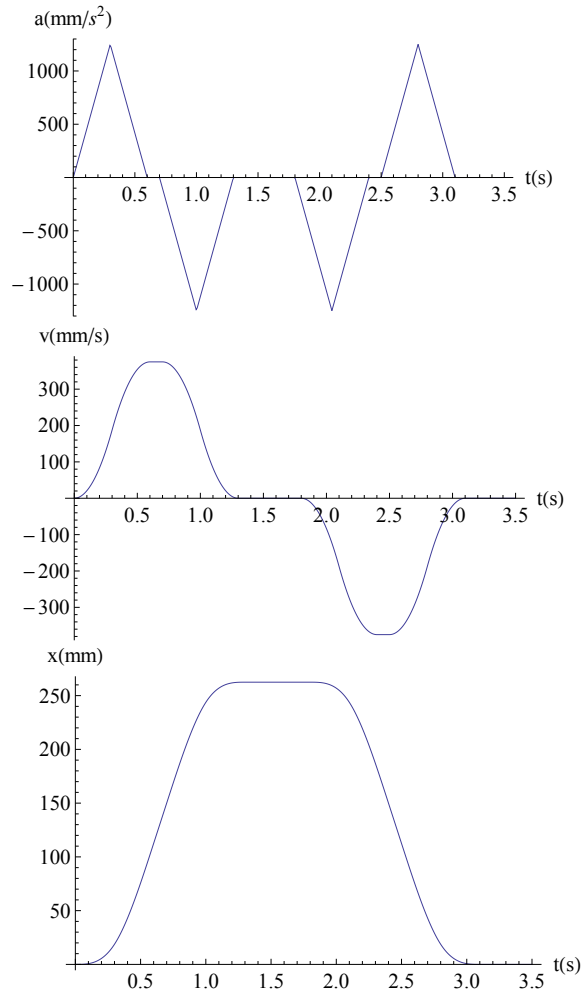
Figure 4.25: The ramp profile of the location of the atoms during the transport step. The profile is designed such that the acceleration is always continuous over the entire process.

- **Acceleration**. If we assume that the trapping potential is harmonic[3], constant acceleration is equivalent a shift in the harmonic potential. Therefore, the atom should not care about constant acceleration either.

- **Jerk**, the rate of change of acceleration. There is no physical argument that I can think of that says that the atoms should not care about the jerk. Intuitively, the jerk produces shock wave in the atomic cloud that contributes to heating on the

---

[3]This is a really bad assumption. I am using this assumption for the sake of argument only.

sample. The transport profile is designed to minimize the amount of jerk in a given

total transport time.

The shape of the transport profile is shown in Figure 4.25.

ODT / Transport Laser   The transport laser is a 300 W IPG multi frequency modes,

single spatial mode fiber laser. Although the laser produces rather significant amount of

intensity noise, the coupling of the intensity noise to the atoms is limited due to the

relatively small trap frequency and the high temperature of the atoms. The high power is

required since we are loading the atom from the MOT directly to the dipole trap. The

coherence length of the fiber laser is measured to be $270\,\mu$m.

Transport Data   Figure 4.26 shows the absorption image during the transport process.

4.3.6   Trap Transfer

Once the atoms arrive at the glass cell, we have to transfer the atoms to a 2D final trap

required for the experiment. Unfortunately, the needle shape of the transport trap

($24\,\mu$m $\times 6000\,\mu$m $\times 24\,\mu$m) is significantly different from the required final trap shape

($15\,\mu$m $\times 15\,\mu$m $\times 0.05\,\mu$m) so that direct trap transfer would be inefficient. The second

problem is that it is really difficult to transport the atoms directly to the final location,

$10\,\mu$m from the substrate, since the waist of the transport beam is $54\,\mu$m. Therefore, an

intermediate trap is most likely required to put the atom in the final trap.

The figure above shows the major trapping beams that we need for the atoms. The green

beam is the transport beam. The blue beam is the accordion beam represented at 5
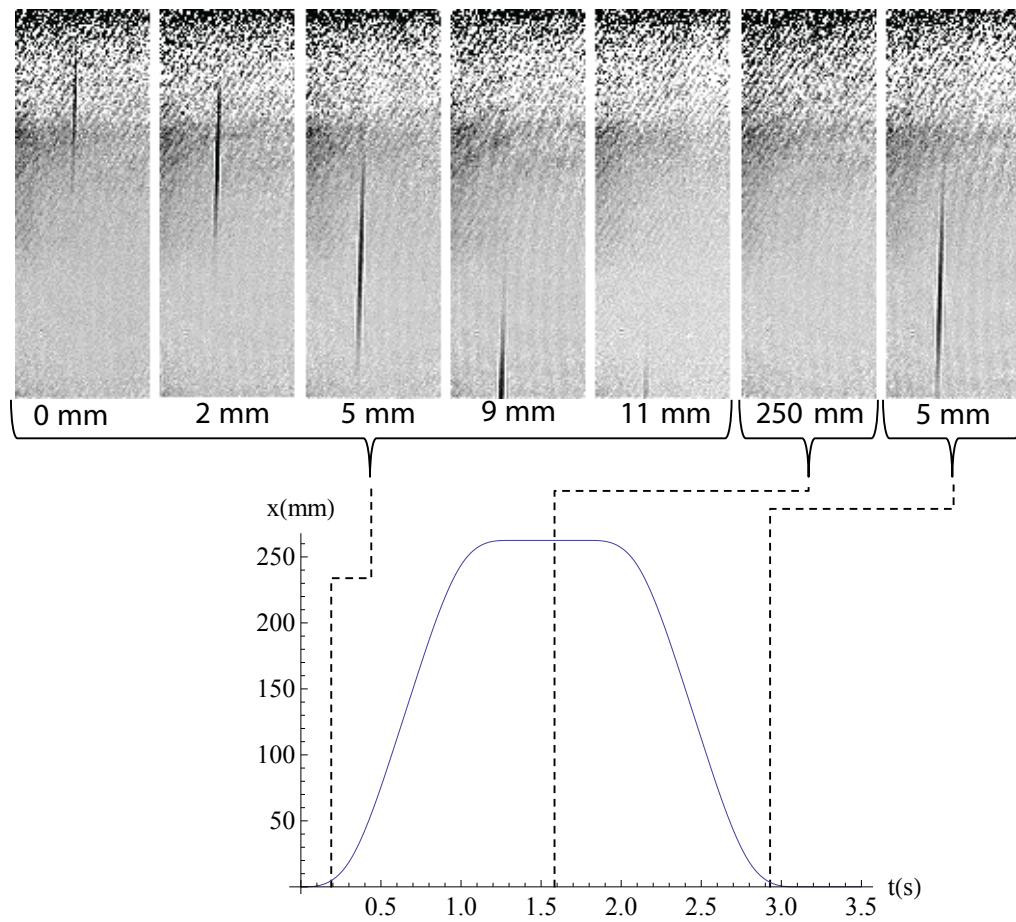
Figure 4.26: Absorption images of the atoms during the transport process. Here, the atoms are transported from $x = 0\,\text{mm}$ to $x = 250\,\text{mm}$ and back to $x = 0\,\text{mm}$. The first 5 images shows the atoms early in the transport process while the last image shows the atoms close to the end of the transport step.

different angles. The purple beam is the vertical dipole trap. The dimple beam is now

shown, but it will be projected from the objective. Finally, the red beams are the

horizontal lattice beams.
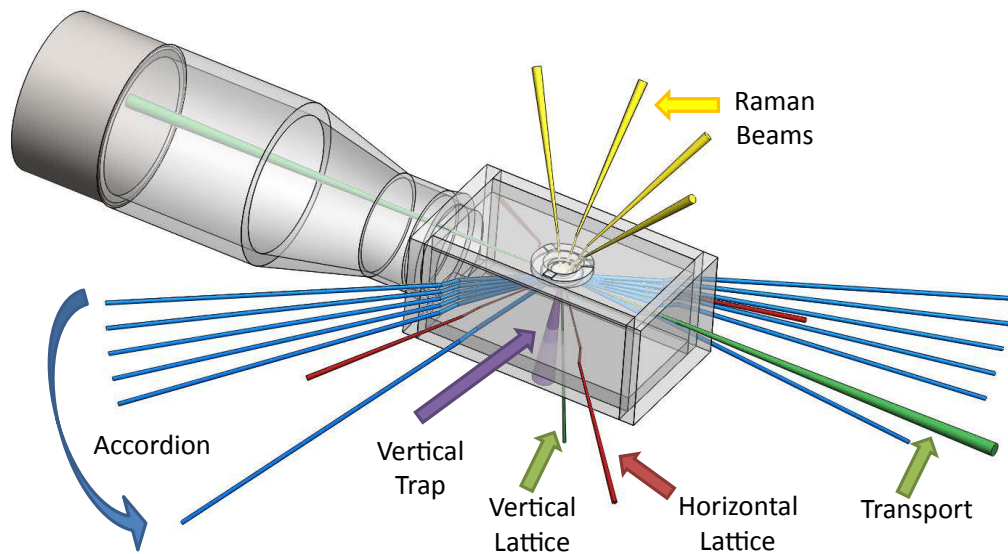
Figure 4.27: All of the trapping beams in the glass cell. The image also shows some of the Raman beams.

4.3.7 Transfer form Transport Trap to Accordion Trap

At the time of writing, the glass cell is not ready to be assembled yet. Instead, I will discuss our current plant for the experiments on the remaining sections.

The first transfer trap is the accordion lattice trap. This 25 W, 289 μm waist lattice beam

Figure 4.28: Transfer form transport trap to accordion trap. The trap aspect ratio is reduced from 60 : 1 to 12 : 1 during this process.

is aimed at $35°$ from the $-\hat{y}$ direction, and tilted $2°$ up towards the hemisphere. This creates a lattice in the vertical direction, while producing some confinement in the transport axis. The goal is to load all of the atoms into a single lattice plane on the accordion lattice. We then ramp down the transport laser power, down to 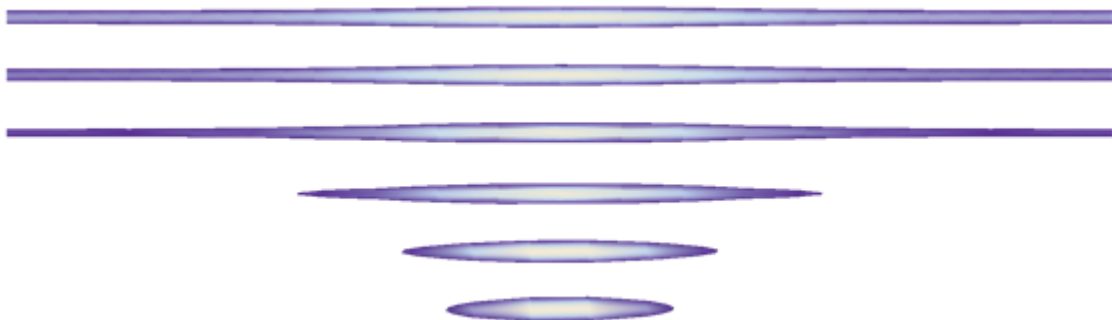5 W. The cloud now has an elongated pancake shape instead of needle shape. Note that there is only 3% spatial overlap between the transport trap and the accordion lattice trap. Intuitively, at least 3% of the atoms will be trapped by the accordion lattice trap. This number could potentially be increased substantially by evaporating the atoms in the transport trap to the accordion lattice trap. At the time of writing, it is not obvious what the total efficiency of the transfer process will be.

Figure 4.28 shows the calculation on how the potential changes as we ramp down the transport potential from 50 W to 5 W. The aspect ratio changes from $60 : 1$ to $12 : 1$ during this process.

The next step is to shine a radial dipole trap from bottom towards the hemisphere. This $100 \, \mu$m waist dipole trap provides radial (x-y) confinement while the accordion lattice primarily provides z confinement. The transport laser beam can then be switched off adiabatically.

4.3.8 Accordion Vertical Transport

We then compress the accordion lattice by changing the tilt from $2°$ to $20°$ as shown in Figure 4.27 and Figure 4.29. This step both compresses the atom in the vertical direction while transporting the atoms upwards at the same time. At the end of this step, the atoms would be located in a single pancake, $15 \, \mu$m from the surface. Since the lattice
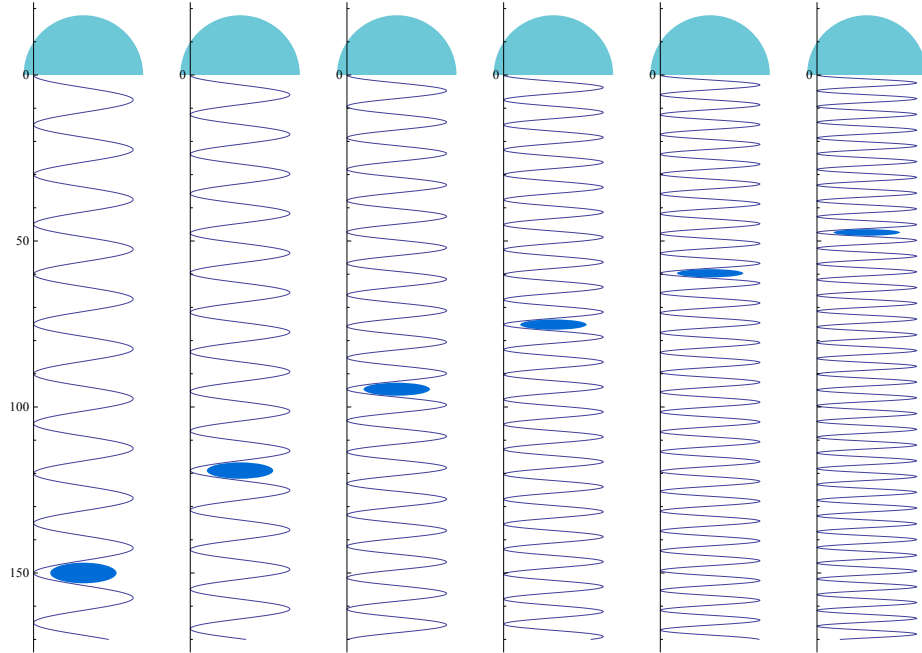
Figure 4.29: Accordion vertical transport. The accordion lattice beam angle can be changed from 2° to 20° in order to both compress the atoms and transport the atoms closer to the hemisphere.

angle is only 20°, the confinement on the axial direction is rather weak.

4.3.9   Vertical Lattice

Depending on the experiment, it might be beneficial to shine a vertical lattice beam from the bottom upwards towards the hemisphere. This lattice beam increases the axial confinement, making the system closer to an ideal 2D system. To remove the possibility of etaloning, the vertical lattice beam will be tilted by 5° from vertical axis.

4.3.10   Final Cooling
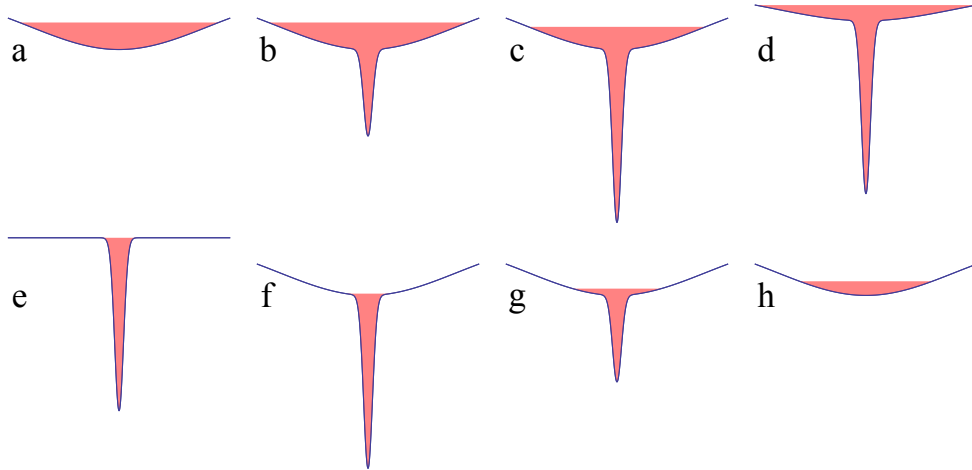
Figure 4.30:   The dimple cooling stages. First, the atoms sit in a wide, low frequency trap. A much smaller but stronger dimple trap is raised at the center of the initial trap. The weak trap is then switched off quickly to eliminate most of the hot atoms, while the dimple holds the colder atoms. The colder atoms from the dimple is then released back to the initial trap adiabatically.

Once the atoms are in the 2D trap, we are now ready for the final cooling / evaporation stage. Before reaching this stage, the atoms went through a series of trap transfer and evaporation steps so the atom number fluctuation is expected to be very large. We have to reduce both the number of atoms and the atom number fluctuation before we can perform reproducible physics experiment.

One method that works really well in the Rubidium experiment is the dimple evaporation method shown in Figure 4.30. Here, a small laser beam with small waist ($5\,\mu$m) is shined on the atoms with $1\,$W of power, creating a $1.2\,$mK trap with $52\,$kHz $\times$ $52\,$kHz $\times$ $1.1\,$kHz confinement. Using simple 2D harmonic potential approximation, the number of states is

given by:

$$N \approx \left(\frac{\text{Trap Depth}}{\text{Trap Frequency}} \right)^2 \propto \text{Laser Power}$$

Ideally, the number of atoms loaded to the dimple trap is exclusively determined by the number of available states on the dimple trap which can be tuned linearly by tuning the laser power. The remaining atoms can then be discarded simply by switching the vertical dipole trap off. Once we have the required number of atoms in the dimple trap, we can then transfer the atoms back to the vertical dipole trap adiabatically[4].

4.3.11 Horizontal Lattice

The horizontal lattices are generated by two independent non-mutually interfering lattice beams, each located at 45° angle from the x-y axis. Each lattice beam is shined at 20° angle upwards pointed at the hemisphere. The reflection from the hemisphere is then retro-reflected back to the hemisphere. The generated one dimensional lattice comes from 4 interfering lattice beams. This is very beneficial since the trap depth is 4 times higher than typical 2 interfering beams lattice setup. In addition of creating lattice in the horizontal direction, this beam configuration also produces a weak lattice in the vertical direction.

The summary of all trapping laser beams is on Table 4.4

---

[4]An alternative entropy explanation is given in [93]

Beam	Waist μm	Power Watt	λ nm	θ degree	ϕ degree
Transport	54	250	1070	0	180
Accordion	290	20	1064	2-20	35
Vertical Trap	100	1	780	85	0
Dimple(from above)	5	1	780	-90	0
Vertical Lattice	80	20	1064	20	45 (4x)
Horizontal Lattice	80	20	1064	85	90

Table 4.4: Summary of all trapping laser beams

Lattice Mirror Configuration

There are two common way of generating the reflected light for the lattice beam: the double imaging approach and single curve mirror approach.
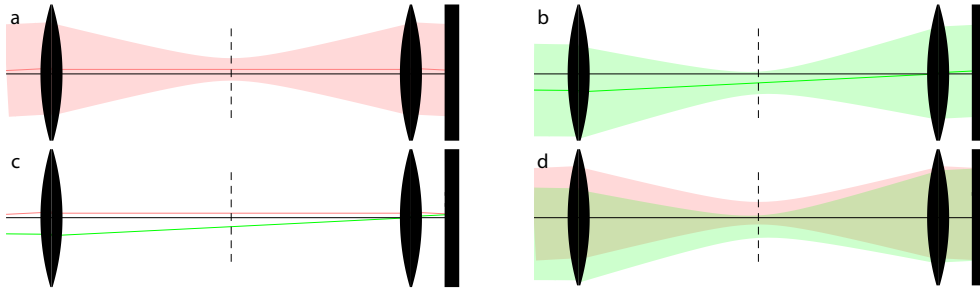
Figure 4.31: Lattice with lens and plane mirror retroreflector. This figure shows how the setup performs under vertical misalignment of the input beam. The center of the combined trap is shifted due to this misalignment.

Double Imaging In this approach, the mirror is just a plain mirror. The lens is responsible for refocusing the light back on the atom[26]. Figure 4.31 shows what happens when we shift the incoming beam by some angle. The reflected beam will shift even more in the opposite direction such that the center of mass of the two beams are not at the location of the atoms. This configuration is not so stable against a shift in input beam.

One solution is to move the mirror further away such that we form a 4f imaging system to balance out the shift on both the incoming and reflected beam. This configuration is insensitive against a tilt of the incoming beam.

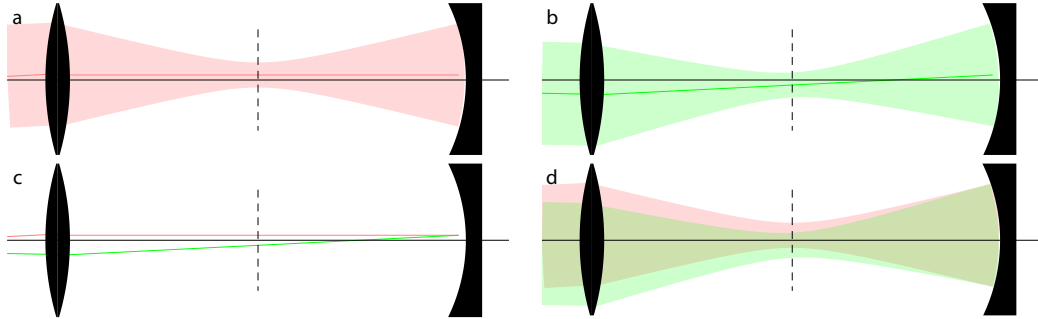Figure 4.32: Single curved mirror retroreflector under vertical misalignment of the input beam. The retroreflected beam shifts down by the same amount as the vertical misalignment so that the center of the trap stays the same independent of the vertical misalignment of the input beam.

Single Curve Mirror Retroreflector Figure 4.32 shows the single curve mirror approach under vertical displacement of the input beam. In the final configuration, the atom is located at the center of mass of the two beams which means that the atom is still located at the center of the overall trap. This means that this lattice configuration is insensitive to vertical misalignment of the input beam.

Figure 4.33 shows what happens when we tilt the input beam. The reflected beam exactly overlaps with the incoming beam. This configuration is also insensitive to a tilt misalignment in the input beam.

Lattice Laser Since the optical lattice axis is produced by retro-reflecting a single beam, it is important that the coherence length to be significantly larger than the distance from the atom to the mirror. The IPG laser with a coherence length of $unit[270]\mu m$
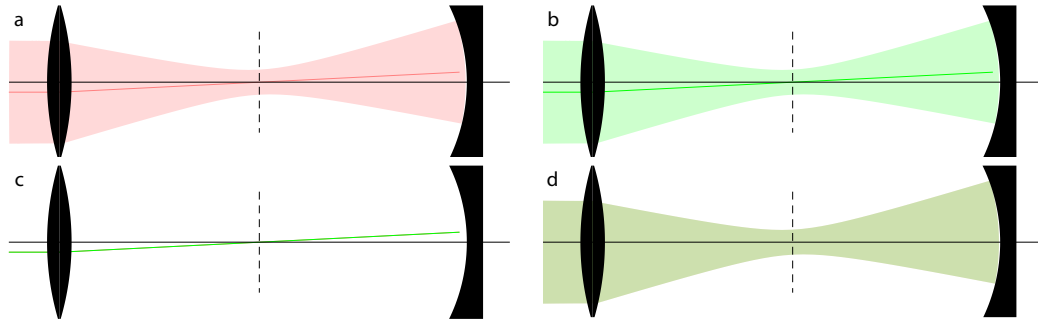
Figure 4.33: Single curved mirror retroreflector under tilt misalignment of the input beam. The retroreflector reflects the beam directly on the input beam so that the center of the trap remains the same independent of the tilt misalignment of the input beam.

cannot be used for this purpose.

The lattice laser beam is generated in two stages. The seed laser is a Mephisto laser which has a coherence length greater than 1 km. After some frequency shifting using AOM, the laser beam is then amplified using three 50 W Nufern fiber amplifiers, one for each lattice axis.

The performance of the lattice laser is important for the following reasons:

- Coherence length. When the coherence length is too short, the location of the lattice interference pattern shift at fast length scale.

- Intensity noise. The parametric heating grows really fast with the trap frequency. First, lithium has very light mass which increases the trap frequency. Second, since we are using optical lattice, the trap frequency in each lattice site is usually very high. This two compounded effects make it very important that the intensity noise is minimal.

The Mephisto seed laser is probably good enough for our applications. However, there are

several potential problems in our Nufern fiber aplifiers:

- SBS. This problem is fundamental to a fiber system and there is not much we can do about this other than reducing the laser power inside the fiber.

- Switching power supplies. The Nufern fiber laser uses several switching power supplies to drive the pump diodes both on the per-amplifier stage and the main amplifier stage. Although there is some switching frequency filtering on those power supplies, we found that there is not enough filtering and some fraction of the switching noise can be measured on the laser power intensity noise.

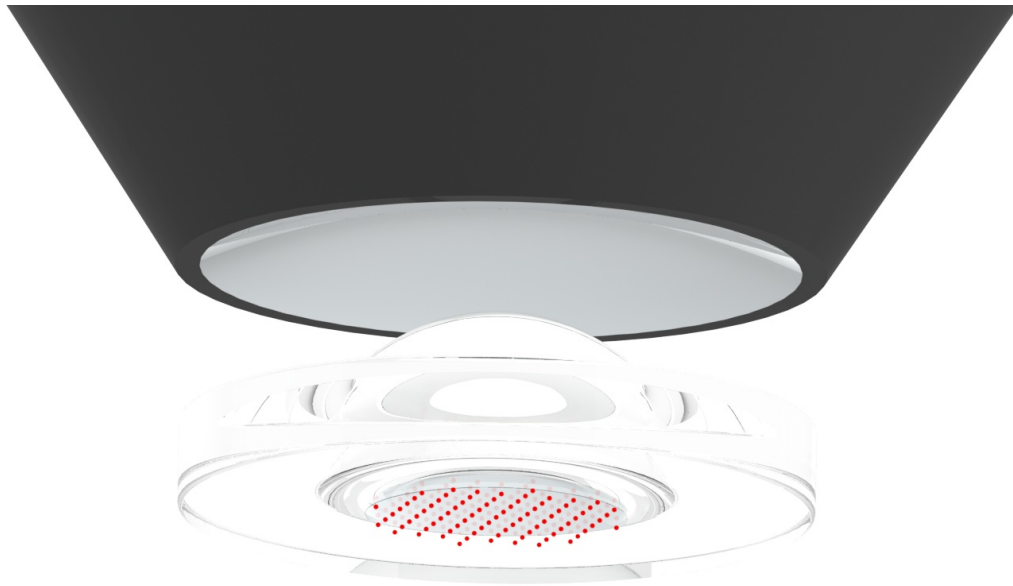### 4.3.12 Single Lattice Site Resolution Imaging

Figure 4.34: An artistic drawing of the atoms under the objective and hemisphere. The image is not to scale.

There are two challenges that have to be solved in order to image single atom with single site resolution. First, we have to have high numerical aperture imaging system. This is

achieved using the hemisphere and a relatively high numerical aperture objective. Second, we need to collect enough photons from each atom. The number of photons required greatly depends on the imaging resolution. In our case, we probably need at least around 1000 photons per atom in order to reliably resolve atoms from different lattice sites. This is achieved using Raman sideband cooling described later.

Hemisphere
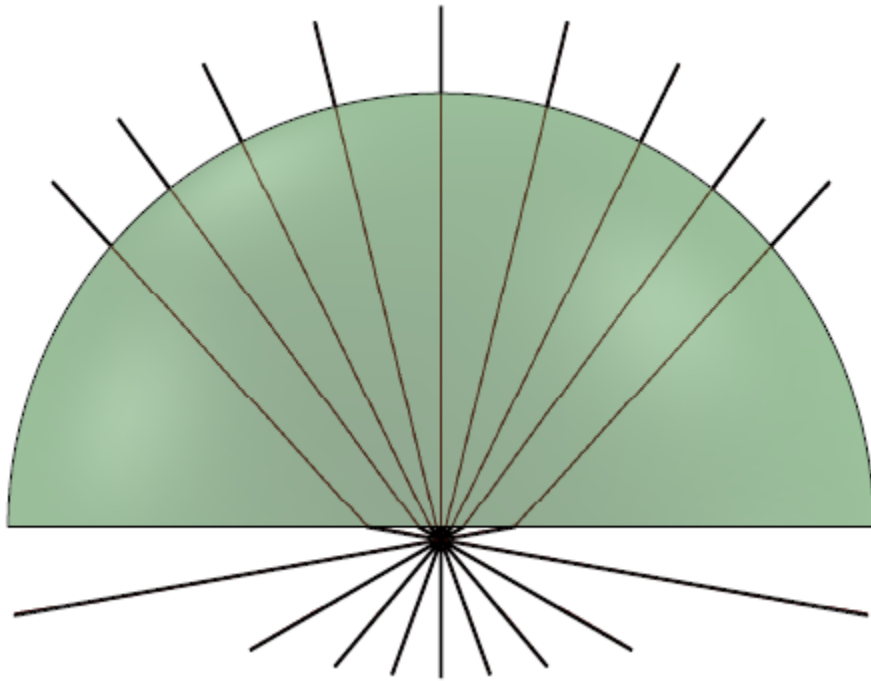
Figure 4.35: Cross section of rays coming from a point source under the hemisphere. Here, we see that all outgoing beams exits almost perpendicular to the surface of the hemisphere. We can use the same objective as if the hemisphere is not there in the first place.

The working principle of a hemispherical lens is shown on the figure above. The hemisphere acts as a first lens, located very close to the atom. The distance from the atom to the center of the hemisphere is greatly exaggerated on the figure above. Here, the atom emits light ray isotropically. Half of the light is captured by the hemisphere while

the other half shines downward and not captured. The flat surface of the hemisphere refracts the incoming beams from the atom into a smaller angle due to the Snell's law:

$$\sin \theta_t = \frac{\sin \theta_i}{n}$$

For the most extreme angle that can be captured by the hemisphere ($\theta_i = 90°$), assuming the index of refraction is $n = 1.46$, the transmitted angle is now $\theta_t = 43.2°$. In the limit that the atoms are very close to the surface to the flat surface of the hemisphere compared to the radius of the hemisphere, the beams always exit normally to the curved surface of the hemisphere. Therefore, the outgoing rays out of the hemisphere is still radial, exactly the same as in the case when there is no hemisphere. This greatly simplifies the objective requirement since we can simply use commercial objective directly without any modification. The only thing the hemisphere lens does is to reduce all of the angles due by using refraction effectively increasing the numerical aperture of the next lens by a factor of n. Note that there is not much point of having the hemisphere opening larger than the critical angle θ_t since there will be no beam coming out at larger angles.

The hemisphere approach in imaging ultracold gases has been recently demonstrated [65, 93, 77].

Pinning Lattice

The first step in the imaging process is to ramp the power of all lattice beams to maximum which increases the confinement, effectively pins each atom to a single lattice site. There are two possible configurations for the vertical pinning. The first method is to turn on the vertical lattice to increase the trapping frequency in the axial direction. The downside is that the auxiliary lattice spacing is short, thus increasing the tunneling rate. The second
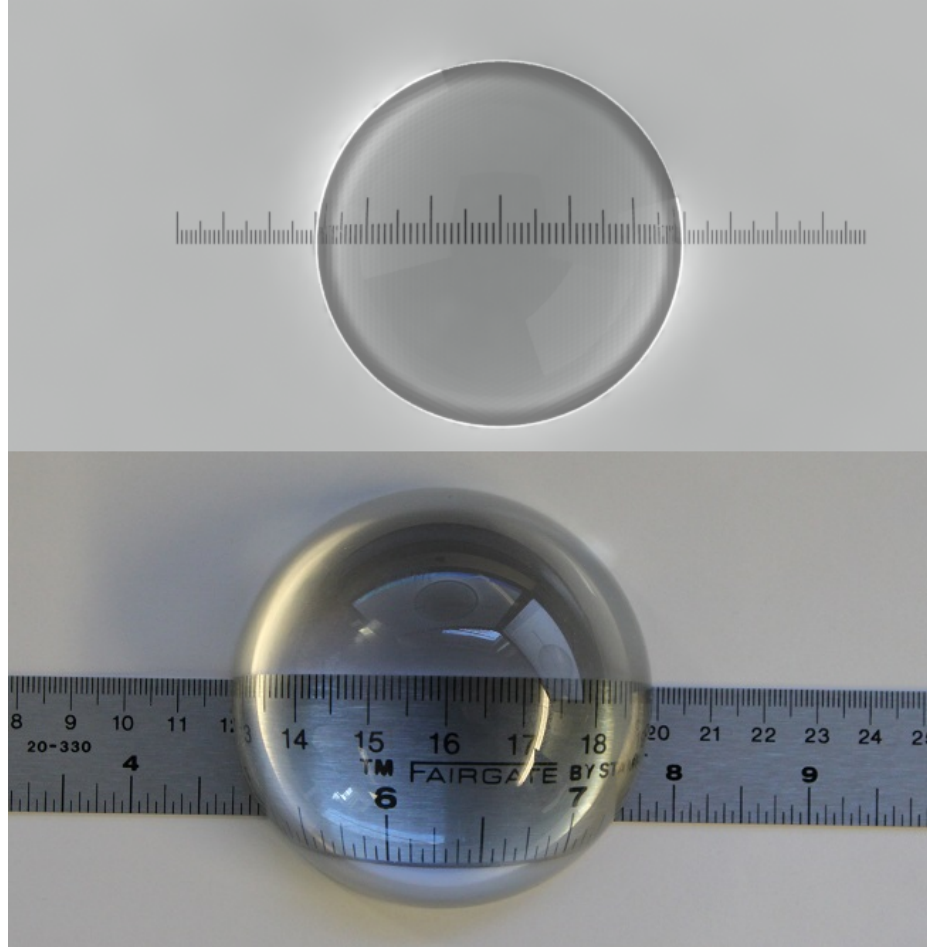
Figure 4.36: The top figure is a simulated hemisphere, while the bottom figure is a real hemisphere. We can clearly see that the hemisphere magnifies the object under it significantly. We use the same effect the enhance our imaging resolution.

method is to pin the atoms without this vertical lattice. Instead, we rely the pinning on the residual vertical lattice generated by the horizontal lattice beams. Although the trapping frequency is weaker, the lattice spacings is far larger, reducing the tunneling rate considerably. At maximum power, the expected lattice parameters is shown on Table 4.5.

Since the goal of the pinning lattice is to prevent the atoms from tunneling to different lattice sites, it might actually be beneficial to switch the vertical lattice beam off. The vertical lattice beam has lattice spacing of around 532 nm, while the vertical lattice

Property	Horizontal	Vertical Residual	Vertical Lattice
Frequencies (MHz)	2.2	1.2	1.7
Lattice spacings (nm)	560	1550	532
Recoil energy / h (kHz)	26	3.4	29
Lattice Depth (E_r)	1840	3680	431
Non-tunneling bands	18	29	6
Max temperature (μK)	1900	1670	490

Table 4.5: Expected pinning lattice parameters

spacing generated by the horizontal lattice has spacing of 1550 nm.

Raman Sideband Cooling

Once we pin the atoms to single lattice sties, we then can proceed with the Raman-sideband cooling process. The Raman sideband cooling functions both for cooling and for generating fluorescent light to be imaged. If everything goes according to plan, we should see images of the atoms similar to Figure 4.37.
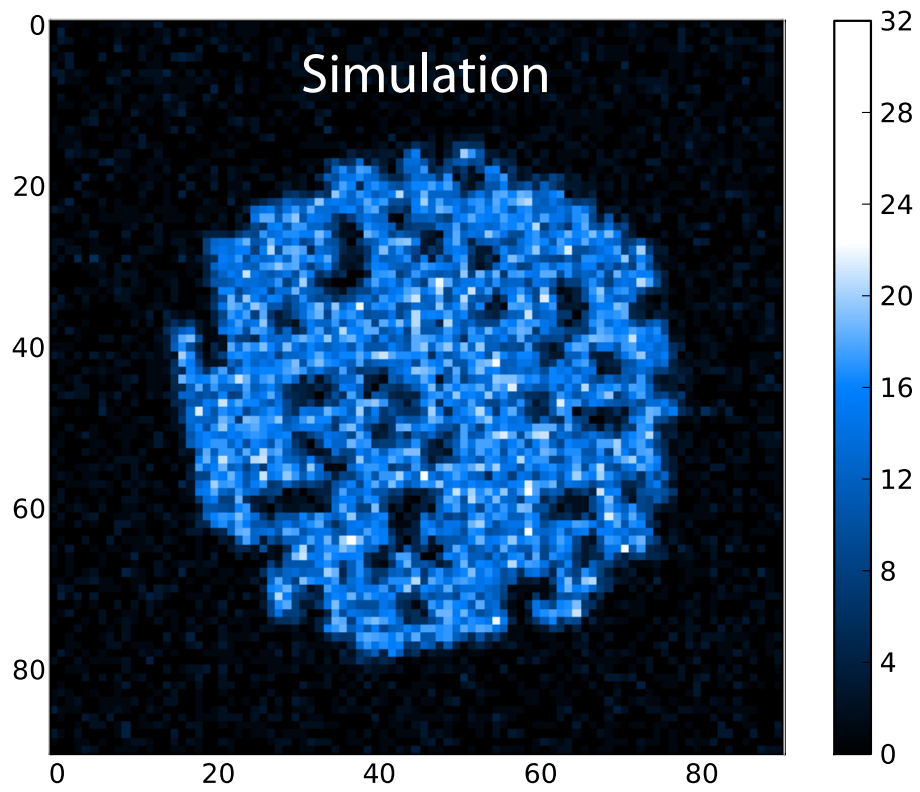
Figure 4.37: Simulation of what the single site resolution images should look like.

Chapter 5

Outlook

5.1 Backup Plans

Although we expect that the experiment sequence proposed on the previous chapter to work as planned, here I will discuss several backup plans if for reasons that we have not considered, the single-site imaging using Raman sideband cooling is not feasible.

5.1.1 No Cooling Imaging

As mentioned in Section 2.4.1, our pinning lattice can pin the atoms enough so that each atom can emit 10 resonant photons before the atom starts to tunnel. However, if a non-saturated light is used, around 500 photons can be scattered before the atoms start to tunnel away. With improved imaging resolution and data analysis, it is plausible that we can get single site imaging resolutions with only 500 photons.

5.1.2 Doppler Cooling Imaging

Instead of using Raman sideband cooling, it might be possible to cool the atom using Doppler cooling. As currently designed, the pinning lattice can hold atoms up to $1900\,\mu\mathrm{K}$ without tunneling as mentioned in Table 4.5. The Doppler cooling limit for lithium is around $140\,\mu\mathrm{K}$, which is well within the pinning lattice limits. The pinning lattice temperature limit can also be increased by using smaller beam waist for the pinning lattice beams.

5.1.3 Hyperfine Polarization Gradient Cooling Imaging

Standard polarization gradient cooling relies on m_F state dependent light shift[94]. In lithium however, this differential light shift is much smaller because the excited state hyperfine state is not resolved. Another recent idea in the atomic physics community is to use hyperfine state dependent light shift instead of m_F state dependent lattice. I have not figured out the detail for this implementation, but it probably involves two lattice beams with well defined relative spatial phase between them generated using EOM.

5.1.4 Increasing Effective Resolution

As mentioned in Section 2.4.4, the number of photons that we need to distinguish atoms in different lattice sites greatly depends on the ratio between the lattice spacing and the optical resolution. One method is to increase the imaging NA. The current design has imaging NA is 0.8. The hemisphere method should allow us to increase the NA to 1 by replacing the objective. Another method is to increase the lattice spacing by reducing the angle between the lattice beams. The downside is that the timescale of the experiment

will go down as we increase the lattice spacings. A third method is by imaging the lithium

using the blue 323 nm transition, doubling the effective resolution.

5.1.5   Photoionization
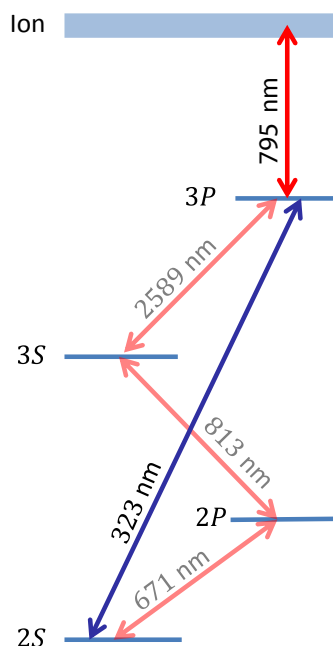
Figure 5.1:   Photoionization level diagram. The photoionization uses 323 nm and 795 laser simultaneously. One concern is the decay from the ^{3}P state to the ^{3}S state where the atom will be ejected from the trap as mentioned in Section 2.2.3.

During the development of this experiment, there was a time when we seriously consider

imaging the atoms using photo-ionization technique[95]. The basic transition in lithium is

shown in the figure above. Rather than trying to cool the atom in the optical lattice, we

can try detecting the atoms with enough signals to noise without any cooling. One

approach is to ionize the atoms with single site resolution and efficiently detect the

ion-electron pair, analogous to electron microscopy[21].

One reason detecting charged particles is a far more efficient process than detecting

photons because we can accelerate the particles to much higher energy using electric field. As a comparison, a typical photon has energy in the order of $2 - 3\,\mathrm{eV}$, while an accelerated charged particle in typical ion detection scheme would have energy in the order of several thousand electron volts. Another argument is the solid angle argument. With ionization method, we can capture most ions and electrons since we accelerate them to the detectors. With photons however, only around 20% of the photons can be captured.

The basic imaging scheme is the following:

- Do experiment in a lattice as usual

- Ramp up the lattice so that each atom is pinned to a single lattice site

- Shine the 795 nm beam on the atoms. This beam does not have to have single site resolution.

- Shine the 323 nm ionization beam to a single lattice site. This is not too hard to do since the diffraction limit of this beam is 162 nm while the lattice spacing is 532 nm.

- Detect the ion-electron pair using channel electron multipliers.

- Repeat the ionization for every lattice site.

Ionization Beam Power   The ^{3}S state shown in Figure 5.1 is actually an anti-trapping state on the 1064 nm lattice light. Once the atom decays to the ^{3}S state, it is very likely for the atom to be ejected from the trap. In order to achieve high imaging fidelity, it is important to minimize the number of time the atom can decay to the ^{3}S state. This can be achieved by saturating the 795 nm transition so that every time the atom reaches the ^{3}P state, that atom will be ionized instead of decaying to the ^{3}S state.

Ionization Rate Estimation   We can roughly estimate the ionization rate by calculating transition matrix element. The energy level of an Akali atom scales as

$$E(n) = -\frac{E_0}{n^2}$$

The density of state is then

$$\rho(n) = \left|\frac{\mathrm{d}n}{\mathrm{d}E}\right| = \frac{n^3}{2E_0},,$$

where E_0 is the binding energy of the valence electron. If we shine a laser with a linewidth ΔE, the number of states addressed by this laser is $N = \rho(n)\Delta E$. For lithium atom, the matrix element of a single line is given by

$$
\begin{aligned}
d &= \langle \psi(n=2)|z|\psi(n)\rangle \\
&\approx 60(n-2)^{n-3}(n+2)^{-n-3}n^{9/2}a_0 \\
&\approx 60n^{-3/2}a_0,
\end{aligned}
$$

where a_0 is the Bohr radius. The total matrix element can be calculated by summing up the individual matrix elements

$$D = \sqrt{d^2 N} = 60a_0\sqrt{\frac{\Delta E}{2E_0}}$$

From here, the transition rate is

$$\Omega = \frac{De}{\hbar}\sqrt{\frac{2I}{c\epsilon_0}},$$

where e is the elementary charge, I is the laser intensity, and c is the speed of light.

Ionization Test Setup   We constructed an ionization test setup to test the feasibility in ionizing and detecting lithium atoms. The test setup is shown in Figure 5.2. The test setup shows that ionizing lithium atoms is possible as shown in the coincidence counts on Figure 5.3.
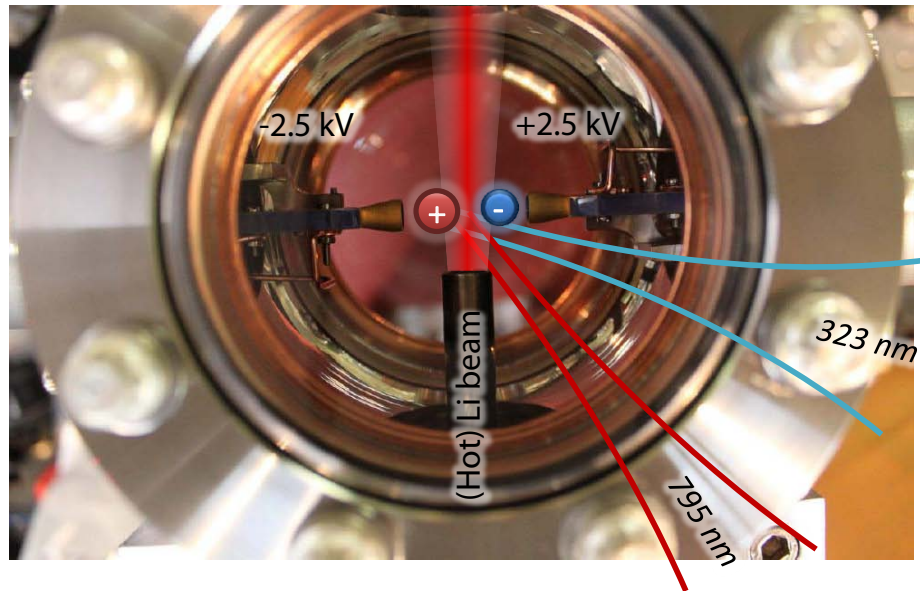
Figure 5.2: The ionization test setup. The atoms are ionized by the two laser beams. The ion and electron pairs are then attracted to the two channel electron multipliers where they will be detected.

5.2 Future Goals

5.2.1 Spin Gradient Demagnetization Cooling

One of the first things that can be done in the experiment is to perform spin gradient demagnetization cooling on a spin mixture in order to get a low enough temperature for anti-ferromagnet ordering[7, 96]. Here, a system with single lattice site resolution greatly helps since the onset of anti-ferromagnet ordering can be observed well before full anti-ferromagnetic ordering appears by possibly using some correlation measurements. This anti-ferromagnetic ordering could be a good starting point to simulate d-wave superfluidity.

Figure 5.3:   Ionization coincidence data between the electron detector and the ion detector. The ions arrive around 160 ns after the electrons since the ions are significantly heavier than the electrons.

5.2.2   Spin Control Using Double Well
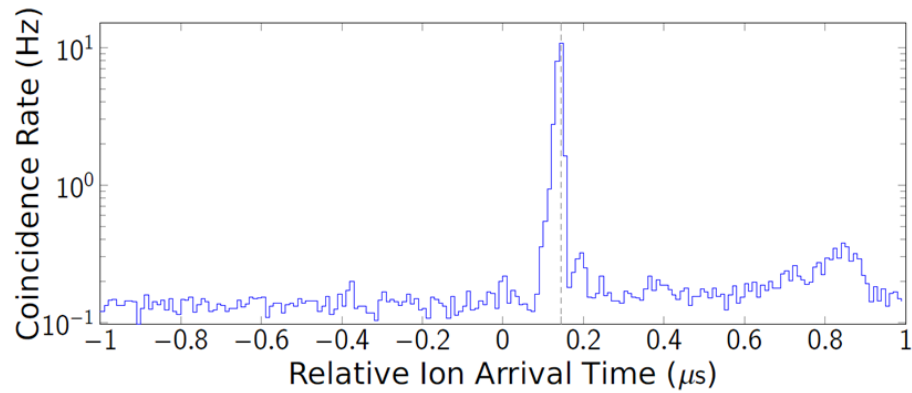
Another method that could be feasible is to control the spin states using double well

method[97, 98, 15].

Bibliography

[1] Wikipedia. Churchturing thesis, May 2012.
`http://en.wikipedia.org/wiki/Church-Turing_thesis`.

[2] S. Barry Cooper. *Computability Theory*. Cambridge University Press, New York, 1999.

[3] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical review*, 47(10):777, 1935.

[4] E. Schrödinger. Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(04):555–563, 1935.

[5] A. Georges, G. Kotliar, W. Krauth, and M.J. Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics*, 68(1):13, 1996.

[6] J. Quintanilla and C. Hooley. The strong-correlations puzzle. *Physics World*, 22(6):32–37, 2009.

[7] Subir Sachdev. *Quantum Phase Transitions*. Chapman and Hall, New York, 2003.

[8] J. Zaanen and G. A. Sawatzky. Band gaps and electronic structure of transition-metal compounds. *Phys. Rev. Lett.*, **55**:4, 1985.

[9] J. G. Bednorz and K. A. Müller. Possible high T_c superconductivity in the Ba-La-Cu-O system. *Zeitschrift fur Physik B Condensed Matter*, 64:189–193, 1986.

[10] H. Maeda, Y. Tanaka, M. Fukutomi, and T. Asano. A new high-T(c) oxide superconductor without a rare earth element. *Japanese Journal of Applied Physics*, 27:L209, 1988.

[11] Z. Z. Sheng and A. M. Hermann. Bulk superconductivity at 120 K in the Tl-Ca/Ba-Cu-O system. *Nature*, 332:138, 1988.

[12] J. Bardeen, L.N. Cooper, and J.R. Schrieffer. Theory of superconductivity. *Physical Review*, 108(5):1175, 1957.

[13] A. J. Leggett. What DO we know about high T_c? *Nature Physics*, 2:134–136, March 2006.

[14] Michael P. Marder. *Condensed Matter Physics*. Wiley, New Jersey, 2010.

[15] M. Lubasch, V. Murg, U. Schneider, J.I. Cirac, and M.C. Bañuls. Adiabatic preparation of a heisenberg antiferromagnet using an optical superlattice. *Physical Review Letters*, 107(16):165301, 2011.

[16] AM Rey, R. Sensarma, S. Fölling, M. Greiner, E. Demler, and MD Lukin. Controlled preparation and detection of d-wave superfluidity in two-dimensional optical superlattices. *EPL (Europhysics Letters)*, 87:60001, 2009.

[17] W. Hofstetter, JI Cirac, P. Zoller, E. Demler, and MD Lukin. High-temperature superfluidity of fermionic atoms in optical lattices. *Physical review letters*, 89(22):220407, 2002.

[18] Richard Feynman. Simulating Physics with Computers. *International Journal of Theoretical Physics*, 21:467–488, June 1982.

[19] Seth Lloyd. Universal Quantum Simulators. *Science*, 273:1073–1078, August 1996.

[20] C. Weitenberg, P. Schauß, T. Fukuhara, M. Cheneau, M. Endres, I. Bloch, and S. Kuhr. Coherent light scattering from a two-dimensional mott insulator. *Physical Review Letters*, 106(21):215301, 2011.

[21] T. Gericke, P. Würtz, D. Reitz, T. Langen, and H. Ott. High-resolution scanning electron microscopy of an ultracold quantum gas. *Nature Physics*, 4(12):949–953, 2008.

[22] Alain Aspect and Massimo Inguscio. Anderson localization of ultracold atoms. *Physics today*, 62:30, 2009.

[23] B. Deissler, E. Lucioni, M. Modugno, G. Roati, L. Tanzi, M. Zaccanti, M. Inguscio, and G. Modugno. Correlation function of weakly interacting bosons in a disordered lattice. *New Journal of Physics*, 13:023020, 2011.

[24] W.S. Bakr, P.M. Preiss, M.E. Tai, R. Ma, J. Simon, and M. Greiner. Orbital excitation blockade and algorithmic cooling in quantum gases. *Nature*, 480(7378):500–503, 2011.

[25] I. Bloch, J. Dalibard, and S. Nascimbène. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8(4):267–276, 2012.

[26] Markus Greiner. *Ultracold quantum gases in three-dimensional optical lattice potentials*. PhD thesis, Ludwig-Maximilians-Universität München, 2003.

[27] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. *Rev. Mod. Phys.*, 80:885–964, Jul 2008.

[28] Jonathan Simon et al. Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472:307–312, April 2011.

[29] Gyu-Boong Jo et al. Itinerant ferromagnetism in a fermi gas of ultracold atoms. *Science*, 325(5947):1521–1524, 2009.

[30] K. et al Kim. Quantum simulation of frustrated ising spins with trapped ions. *Nature*, 465(7298):590–593, 2010.

[31] DC McKay and B. DeMarco. Cooling in strongly correlated optical lattices: prospects and challenges. *Reports on Progress in Physics*, 74:054401, 2011.

[32] F. Werner, O. Parcollet, A. Georges, and SR Hassan. Interaction-induced adiabatic cooling and antiferromagnetism of cold fermions in optical lattices. *Physical review letters*, 95(5):56401, 2005.

[33] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. *Reviews of Modern Physics*, 80(3):885, 2008.

[34] U. Schneider, L. Hackermüller, S. Will, T. Best, I. Bloch, TA Costi, RW Helmes, D. Rasch, and A. Rosch. Metallic and insulating phases of repulsively interacting fermions in a 3d optical lattice. *Science*, 322(5907):1520–1525, 2008.

[35] Rudolf Grimm, Matthias Weidemller, and Yurii B. Ovchinnikov. Optical dipole traps for neutral atoms. *Advances in Atomic, Molecular and Optical Physics*, **42**:95–170, 2000.

[36] Ad Lagendijk, Bart van Tiggelen, and Diederik S. Wiersma. Fifty years of anderson localization. *Phys. Today*, 62(8):24–29, 2009.

[37] M. Bajcsy, S. Hofferberth, V. Balic, T. Peyronel, M. Hafezi, A.S. Zibrov, V. Vuletic, and M.D. Lukin. Efficient all-optical switching using slow light within a hollow fiber. *Physical review letters*, 102(20):203902, 2009.

[38] J. Simon, H. Tanji, J.K. Thompson, and V. Vuletić. Interfacing collective atomic excitations and single photons. *Physical review letters*, 98(18):183601, 2007.

[39] Immanuel Bloch. Ultracold quantum gases in optical lattices. *Nature Physics*, 1(1):23–30, 2005.

[40] Cheng Chin, Rudolf Grimm, Paul Julienne, and Eite Tiesinga. Feshbach resonances in ultracold gases. *Rev. Mod. Phys.*, 82:1225–1286, Apr 2010.

[41] N. Gemelke, X. Zhang, C.L. Hung, and C. Chin. In situ observation of incompressible mott-insulating domains in ultracold atomic gases. *Nature*, 460(7258):995–998, 2009.

[42] W.S. Bakr, J.I. Gillen, A. Peng, S. Fölling, and M. Greiner. A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice. *Nature*, 462(7269):74–77, 2009.

[43] J.F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr. Single-atom-resolved fluorescence imaging of an atomic mott insulator. *Nature*, 467(7311):68–72, 2010.

[44] Charles Kittel. *Introduction to Solid State Physics*. John Wiley & Sons, Inc, New Jersey, 2005.

[45] J.H. Boer and E.J.W. Verwey. Semi-conductors with partially and with completely filled 3d-lattice bands. *Proceedings of the Physical Society*, 49:59, 1937.

[46] A.K. Geim and K.S. Novoselov. The rise of graphene. *Nature materials*, 6(3):183–191, 2007.

[47] Tilman Esslinger. Fermi-hubbard physics with atoms in an optical lattice. *Annual Review of Condensed Matter Physics*, 1:129–152, 2010.

[48] Niels Strohmaier. *Exploring the Hubbard model with ultracold fermionic atoms in an optical lattice*. PhD thesis, ETH Zurich, 2010.

[49] J.K. Chin, DE Miller, Y. Liu, C. Stan, W. Setiawan, C. Sanner, K. Xu, and W. Ketterle. Evidence for superfluidity of ultracold fermions in an optical lattice. *Nature*, 443(7114):961–964, 2006.

[50] Robert Jordens. *Metallic and Mott-insulating phases in fermionic quantum gases*. PhD thesis, ETH Zurich, 2010.

[51] K. et al Dieckmann. Decay of an ultracold fermionic lithium gas near a feshbach resonance. *Phys. Rev. Lett.*, 89:203201, Oct 2002.

[52] W. Ketterle and M.W. Zwierlein. Making, probing and understanding ultracold fermi gases. *Arxiv preprint arXiv:0801.2500*, 2008.

[53] CH Schunck, MW Zwierlein, CA Stan, SMF Raupach, W. Ketterle, A. Simoni, E. Tiesinga, CJ Williams, and PS Julienne. Feshbach resonances in fermionic^{6} li. *Physical Review A*, 71(4):045601, 2005.

[54] Michael E. Gehm. *Properties of Li*. PhD thesis, University of Arizona, 2003.

[55] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, New York, 2002.

[56] David J. Griffiths. *Introduction to Quantum Mechanics*. Pearson Prentice Hall, New Jersey, 2005.

[57] Daniel Adam Steck. *Rubidium 87 D Line Data*. PhD thesis, University of Oregon, 2008.

[58] Daniel Adam Steck. *Rubidium 87 D Line Data*. PhD thesis, University of Oregon, 2008.

[59] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. Wiley, New York, 1991.

[60] Modern Atomic and Optical Physics. *Rubidium 87 D Line Data*. PhD thesis, Harvard University, 2005.

[61] A. Lindgard and S.E. Nielsen. Transition probabilities for the alkali isoelectronic sequences. *Atomic Data and Nuclear Data Tables*, 19(6):533–633, 1977.

[62] Rodney Loudon. *The Quantum Theory of Light*. Oxford University Press, Oxford, 2010.

[63] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. W. B. Saunders, Philadelphia, 1976.

[64] Joseph W. Goodman. *Introduction to Fourier Optics*. Roberts and Company, Greenwood Village, 2005.

[65] Jonathon Isaiah Gillen. *The Quantum Gas Microscope*. PhD thesis, Harvard University, 2009.

[66] Vladan Vuletić, Cheng Chin, Andrew J. Kerman, and Steven Chu. Degenerate raman sideband cooling of trapped cesium atoms at very high atomic densities. *Phys. Rev. Lett.*, 81:5768–5771, Dec 1998.

[67] Andrew James Kerman. *Raman Sideband Cooling and Cold Atomic Collisions in Optical Lattices*. PhD thesis, Stanford University, 2002.

[68] X. Zhuang. Nano-imaging with storm. *Nature photonics*, 3(7):365, 2009.

[69] Wikipedia. Internet protocol suite, May 2012.
http://en.wikipedia.org/wiki/Airy_disk.

[70] Wikipedia. Transmission control protocol, May 2012.
http://en.wikipedia.org/wiki/Transmission_Control_Protocol.

[71] Wikipedia. User datagram protocol, May 2012.
http://en.wikipedia.org/wiki/User_Datagram_Protocol.

[72] Wikipedia. Internet protocol suite, May 2012.
http://en.wikipedia.org/wiki/Internet_protocol_suite.

[73] Analog Devices. *AD9910: 1 GSPS, 14-bit, 3.3V CMOS Direct Digital Synthesizer*, 2010. http://www.analog.com/.

[74] A.A. Oswald. Early history of single-sideband transmission. *Proceedings of the IRE*, 44(12):1676–1679, 1956.

[75] Tobias Schuster. *A Novel Apparatus for Experiments with Ultracold Sodium and Lithium*. PhD thesis, Harvard University, 2008.

[76] Florian G. Huber. *A Novel Apparatus for Experiments with Ultracold Fermions.* PhD thesis, Harvard University, 2009.

[77] Waseem S. Bakr. *Microscopic Studies of Quantum Phase Transitions in Optical Lattices.* PhD thesis, Harvard University, 2011.

[78] Selim Jochim. *Bose-Einstein Condensation of Molecules.* PhD thesis, University of Innsbruck, 2004.

[79] Timo Bastian Ottenstein. *Few-body physics in ultracold Fermi gases.* PhD thesis, University of Heidelberg, 2010.

[80] Claudiu Andrei Stan. *Experiments with Interacting Bose and Fermi Gases.* PhD thesis, MIT, 2005.

[81] CA Stan and W. Ketterle. Multiple species atom source for laser-cooling experiments. *Review of scientific instruments*, 76:063113, 2005.

[82] Dallin S. Durfee. *Dynamic Properties of Dilute Bose-Einstein Condenstates.* PhD thesis, MIT, 1999.

[83] M. Greiner, I. Bloch, T.W. Hänsch, and T. Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Physical Review A*, 63(3):031401, 2001.

[84] KB Davis, M.O. Mewes, M.R. Andrews, NJ Van Druten, DS Durfee, DM Kurn, and W. Ketterle. Bose-einstein condensation in a gas of sodium atoms. *Physical Review Letters*, 75(22):3969–3973, 1995.

[85] TL Gustavson, AP Chikkatur, AE Leanhardt, A. Gorlitz, S. Gupta, DE Pritchard, and W. Ketterle. Transport of bose-einstein condensates with optical tweezers. *Arxiv preprint cond-mat/0108496*, 2001.

[86] Tobias Tiecke. *Feshbach resonances in ultracold mixtures of the fermionic quantum gases.* PhD thesis, Universiteit van Amsterdam, 2009.

[87] John Howard Huckans. *Optical Lattices and Quantum Degenerate $^{8}7Rb$ in Reduced Dimensions.* PhD thesis, University of Maryland, 2006.

[88] M.W. Zwierlein, C.A. Stan, C.H. Schunck, S.M.F. Raupach, S. Gupta, Z. Hadzibabic, and W. Ketterle. Observation of bose-einstein condensation of molecules. *Physical review letters*, 91(25):250401, 2003.

[89] TL Gustavson, AP Chikkatur, AE Leanhardt, A. Görlitz, S. Gupta, DE Pritchard, and W. Ketterle. Transport of bose-einstein condensates with optical tweezers. *Physical Review Letters*, 88(2):20401, 2002.

[90] Ananth P. Chikkatur. *Colliding and Moving Bose-Einstein Condensates: Studies of superfluidity and optical tweezers for condenstate transport.* PhD thesis, MIT, 2002.

[91] Aaron E. Leanhardt. *Microtraps and Waveguides for Bose-Einstein Condensates.* PhD thesis, MIT, 2003.

[92] A. Einstein et al. On the electrodynamics of moving bodies. *Annalen der Physik*, 17(891):50, 1905.

[93] Amy Wan-Chih Peng. *Quantum Gas Microscope With Optical Lattice.* PhD thesis, Harvard University, 2010.

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

[95] F. Henkel, M. Krug, J. Hofmann, W. Rosenfeld, M. Weber, and H. Weinfurter. Highly-efficient state-selective sub-microsecond photoionization detection of single atoms. *Arxiv preprint arXiv:1008.1910*, 2010.

[96] P. Medley, D.M. Weld, H. Miyake, D.E. Pritchard, and W. Ketterle. Spin gradient demagnetization cooling of ultracold atoms. *Physical Review Letters*, 106(19):195301, 2011.

[97] S. Fölling, S. Trotzky, P. Cheinet, M. Feld, R. Saers, A. Widera, T. Müller, and I. Bloch. Direct observation of second-order atom tunnelling. *Nature*, 448(7157):1029–1032, 2007.

[98] A.M. Rey, V. Gritsev, I. Bloch, E. Demler, and MD Lukin. Preparation and detection of magnetic quantum phases in optical superlattices. *Physical review letters*, 99(14):140601, 2007.

[99] John David Jackson. *Classical Electrodynamics.* Wiley, New Jersey, 1999.

[100] Wikipedia. Ground loop, May 2012. http://en.wikipedia.org/wiki/Ground_loop_(electricity).

[101] Wikipedia. Transmission line, May 2012. http://en.wikipedia.org/wiki/Transmission_line.

[102] Unknown Editor. Microwaves 101 encyclopedia, May 2012. http://www.microwaves101.com.

[103] John Rice, Dirk Gehrke, and Mike Segal. Understanding noise-spreading techniques and their effects in switch-mode power applications. *Texas Instruments*, 2008.

Appendix A

General Electronics

A.1 Introduction

In basic electronic design, we usually assume that the potential across a wire is the same
everywhere since, to a very good approximation, the wire has zero resistance.
Unfortunately, this simplification quickly breaks down once we are dealing with a
changing signal. Maxwell equation says[99]:

$$\vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}$$

The changing current in the system create changing magnetic field, which implies that
$\vec{\nabla} \times \vec{E} \neq 0$. Unfortunately, the mathematical requirement for the existence of a potential
field such that $\vec{E} = \vec{\nabla}V$ is that $\vec{\nabla} \times \vec{E} = 0$. This implies that the electrical potential is not
a well-defined quantity. The AC ground loop and radio frequency problem can be traced
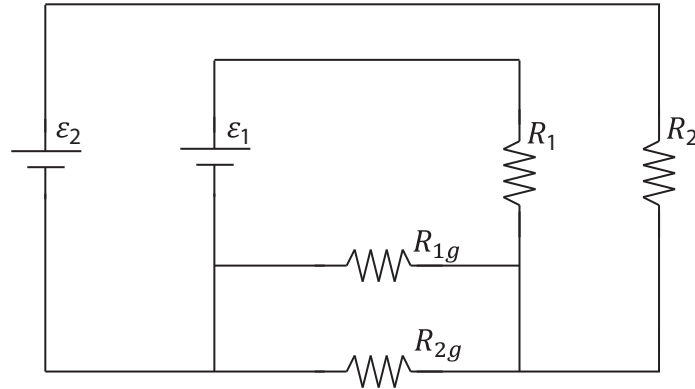back to this particular effect.

A.2   DC Ground Loop

Figure A.1:   A simplified schematic of the source of DC ground loop. The current on resistor R_1 changes based on the current on resistor R_2 due to the shared ground connections between the two circuits.

The DC ground loop effect can be easily understood by observing the circuit shown in Figure A.1. Imagine that we have two circuits: circuit 1 and circuit 2. R_1 and R_2 represent sensors, for example photo diodes. The two circuits are independent except on the connected return line: R_{1g} and R_{2g}. In this configuration, the potential drop across R_1 changes depending on the current on circuit 2[100].

One solution is to reduce the ground resistance of one or both circuits. When the R_{1g} and R_{2g} resistance go to zero, the ground loop problem virtually disappear. Another solution is to separate the return connection of both circuits so that the two circuits are completely decoupled.
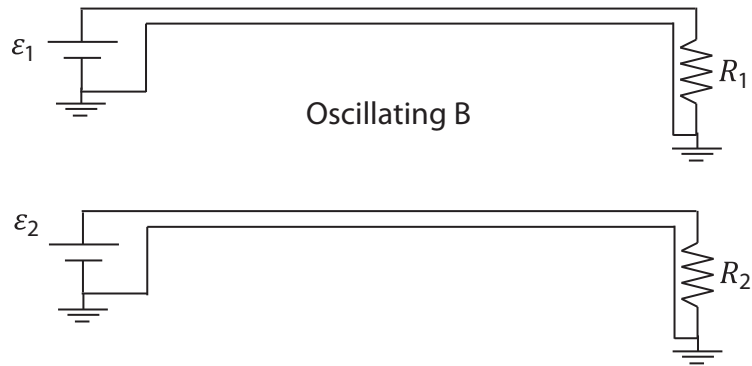
Figure A.2: A simplified schematic of AC ground loop. Consider the case when there is oscillating current in the top circuit. Since the ground connection on both circuits are shared, some of the current from circuit 1 flows through the ground of circuit 2, creating a big loop. The current in the big loop then create an oscillating magnetic field in the loop, effectively coupling the two circuits together through mutual inductance.

A.3 AC Ground Loop

The AC ground loop problem is more severe. Even if we take the ground resistance to be zero and we separate the two circuits completely, there is still mutual inductance between the two circuits which effectively coupled the two circuits together.

Part of the general solution in solving AC ground loop problem is by using a coaxial cable where the signal is located inside the return line. This minimizes the amount of mutual inductance between coaxial lines. However, there is still a potential problem even with coaxial cable as shown in Figure A.2.

As before, R_1 and R_2 represent photo-diodes. The two photo diodes are mounted on an optical table so they share a common ground. The two power supplies represent different channels of the oscilloscope, so they also both share a common ground. Furthermore, the ground of the optical table is the same as the ground of the oscilloscope. In this problem, we assume that the coaxial cables are perfect conductors since the current flowing through

them is relatively small.

However, there is a residual 60 Hz background magnetic field in the entire room. On the circuit above, the magnetic field would induce current across the two ground connections. Since the current on the photodiode is small, this induced current can be a significant fraction of the signal themselves.

Although in principle we can avoid this problem altogether by not grounding anything, in practice, this is very difficult. Fortunately, we do not have to disconnect all of the grounds. As long as there is no loop in the system (a maximum of one ground connection), the ground loop problem would disappear. This can be done by using a dedicated floating power supply for each circuit and by not connecting the ground of the power supply to the ground of the room.

A.4   Impedance

We all know that a BNC cable has $50\,\Omega$ impedance. But if you get a piece of BNC cable, and measure the resistance between the inside and outside layer, you will find that the resistance is infinite. What does this mean? The answer is that if the BNC cable were infinitely long, and you measured the resistance between the inner and outer layer, you would get $50\,\Omega$. This makes sense since if the cable were infinitely long, there would be no way for the ohm meter to find out whether then other end of the cable is either shorted or open-ended, or anything in between. The best the ohm meter can do is to measure the effective impedance of the cable locally, which is $50\,\Omega$. Note that although there is no resistive element in the BNC cable, the impedance is still real.

Consider a voltage wave travelling on a transmission line from left to right. The wave can

be written as:

$$y_1(x, t) = A_1 \sin(\omega t - kx)$$

Whenever a signal in the transmission line hits a boundary with another transmission line, say at $x = 0$, with different impedance, a reflection occurs. The voltage amplitude of the reflection is given by:

$$\Gamma = \frac{Z_2 - Z_1}{Z_1 + Z_2}$$

The reflected wave can be written as

$$y_2(x, t) = A_2 \sin(\omega t + kx) = \Gamma A_1 \sin(\omega t + kx)$$

The total voltage is just the sum of the incoming wave and the reflected wave:

$$y(x, t) = A_1 \sin(\omega t - kx) + \Gamma A_1 \sin(\omega t + kx)$$

When $Z_1 = Z_2$, the reflection coefficient is zero. This is called the impedance matching condition, where there is no reflection at all.

In real life, the problem is not as simple as it seems. Consider a case where we measure the voltage exactly at the boundary, $x = 0$. Furthermore, let us keep the boundary to be open-ended: do not connect anything after the first wire. This corresponds to $Z_2 = \infty$. The voltage at this location is given by:

$$y(x = 0, t) = A_1 \sin(\omega t) + \Gamma A_1 \sin(\omega t) = 2A_1 \sin(\omega t)$$

Note that the voltage at the end of the wire is twice the driving voltage! This is really weird. If you actually try to do this, put a 5 V supply on the left end of a wire, and measure the voltage at the right end, you will not measure 10 V. You will still measure 5

V. Where is the problem? It turns out that the impedance calculation above assumes that the wire is infinitely long. In practice, the above equation works when the wavelength of the signal is far shorter than the length of the cable. Instead of sending a 5V signal on the left end, you could send a 1 ns, 5V pulse on the left end. Sure enough, at the right end you will see a 10V signal. A simple rule of thumb is that if the transmission line is less than $\lambda/10$, we can assume that we do not have to consider transmission line effect.

A.5   RF Electronics

One of the most important things to realize when dealing with RF electronics is that voltage is not a mathematically defined entity anymore. The whole concept of voltage relies on the fact that $\nabla \times E = 0$. However, in the RF case this is not satisfied because according to the Maxwell Equations,

$$\nabla \times E = -\frac{\partial B}{\partial t}$$

Which is clearly non-zero for an oscillating electromagnetic wave.

The major implication is that we cannot treat all ground as equal. The effective potential at any ground level is only marginally related to the ground potential in another place. Instead, we should use transmission lines to guide our intuition in RF electronics.

A.6   Transmission Lines

Here I will discuss several types of transmission lines that are usually encountered in atomic physics laboratories[101].

A.6.1 BNC and SMA Coaxial Cables

The most common transmission lines are the BNC and SMA coaxial cables. The impedance is either $50\,\Omega$ or $75\,\Omega$ depending on the wire. The general formula to calculate the impedance of a coaxial cable is:

$$Z_0 = \frac{1}{2}\pi\sqrt{\frac{\mu}{\epsilon}}\ln\left(\frac{D}{d}\right),$$

where:

- $\mu=$ Magnetic permeability of the insulator. Usually, $\mu \approx \mu_0$

- $\epsilon = \epsilon_r\epsilon_0 =$ Dielectric constant of the insulator

- $D =$ outside diameter of the inner conductor

- $d =$ inside diameter of the outside conductor
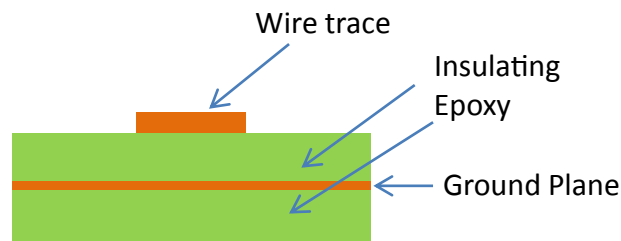
A.6.2 Microstrip

Figure A.3: Microstrip waveguide.

The most common type of transmission line on a PCB board is the microstrip[102]. In this setup, a wire trace is placed on top of an insulating layer. Under the insulating layer is a conductor plane, usually the ground plane. The wire itself is surrounded by air on the

top side. The formula for calculating microstrip impedance is complicated. However, it is usually good to know the typical value of the impedance in standard pcb board:

- Insulation layer thickness for typical 6 layer board: $0.2\,\mathrm{mm}$

- Insulator dielectric constant for FR4: $\epsilon_r = 4.7$

- Wire width: $0.15\,\mathrm{mm}$

- Impendance: $71\,\Omega$

A.6.3 Twisted-Pair Cable

The most common uses of twisted-pair cable is Ethernet cables and Firewire cables. The characteristic impedance is around $100\,\Omega$ to $150\,\Omega$.

A.7 Power Supplies and Regulators

Power supplies and voltage regulators are a necessity in virtually all experiments. Although technically incorrect, here I assume that a power supply is a device that converts wall outlet AC voltage to desired DC voltages, while a voltage regulator is a device that converts a small DC voltage to one or more small DC voltages. More specifically, typical ultracold experiments use power supplies for two distinct categories:

- To power electronics, equipment, lasers, and computers.

- To power magnetic coils for controlling the Zeeman shift on the atoms.

This chapter mainly describes power supplies in the first category.

A.7.1 Linear Regulator

Often, different parts of a single electronic circuit need to be powered by different DC voltages. A common way to solve this problem is by powering the board using the highest DC voltage required by any of the components, and by generating the remaining lower DC voltages using linear regulators. This method reduces the number of DC voltages that have to be generated by the AC/DC power supply, which in turns reduces the number of wires that have to be connected from the AC/DC power supply to the electronic circuit.

A linear regulator works like a tunable voltage divider. The primary component of a linear regulator is a power transistor that can adjust the effective resistance such that the desired output voltage is achieved. A major advantage of linear regulator is its simplicity. A typical linear regulator chip only has 3 pins: an input pin, an output pin, and a ground (reference) pin. Another advantage is that due to its analog nature, linear regulator has none of the switching behavior seen in digital circuits. This allows linear regulator to operate with very little additive noise.

The simplicity of the linear regulator comes at a price: efficiency. Due to its resistive nature, linear regulator dissipates a lot of heat. For example, a linear regulator with an input voltage of 15 V and an output voltage of 5V will dissipate 66% of the electrical energy on the regulator itself, and only use the remaining 33% on the actual components. This is the primary reason why linear regulator can usually only be found in devices that consume very little of power.
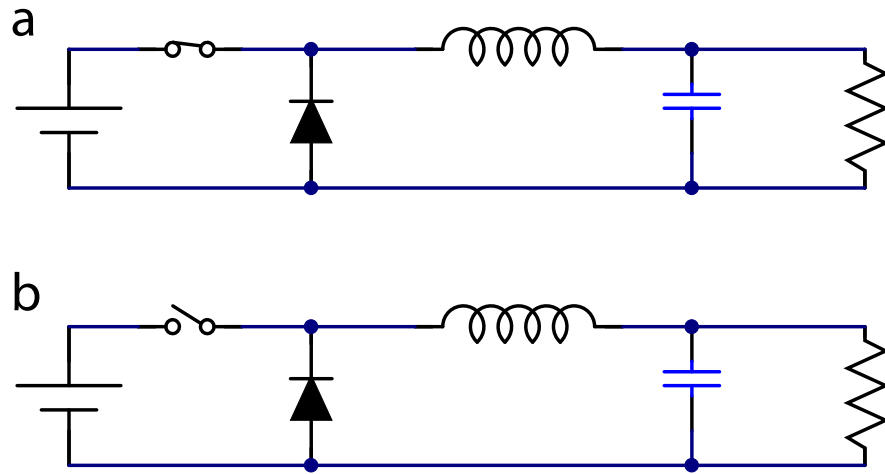
Figure A.4: Switch-mode regulator. In state a, the current flow from the power supply, to the inductor, and to the load resistor. In state b, the current flows from the diode, to the inductor, and to the load resistor. The switching duty cycle is controlled by a feedback system to get a constant output voltage at the load resistor. The capacitor acts as a filter to reduce the switching noise.

A.7.2 Switch-Mode Regulator

Switch-Mode regulator was created to solve the efficiency problem. The idea is that instead of operating the transistor in analog mode where there is significant power dissipation, we operate the switch in digital mode: that is either fully on or fully off. Since the power dissipated in the transistor is $P = IV$, when the transistor is fully on, the ΔV on the switch is zero, which means that there is no power dissipated. Similarly, when the switch is fully off, the I is zero, which again means that there is no power dissipated.

One of the most common types of switch-mode regulator is what is called the buck converter, as shown in Figure A.4. When the switch is turned on, the current starts to flow and some energy is stored in the inductor while some charges the capacitor. Then, when the switch is turned off, the energy stored in the inductor is released to charge the capacitor. The timing of the switching controls the final voltage of the capacitor. Note

that theoretically, 100% of the energy is transformed to the desired voltage since the there are no dissipative elements except for the load resistor.

In order to minimize the ripple voltage at the output, we want the switching frequency of the regulator to be as high as possible. Good switching regulators in the 20 Watt range can operate at around 1 MHz. More specialized regulator for low power application in 0.5 Watt range can operate at much higher frequency of 180 MHz as mentioned in Section 3.4. The higher frequency can be both a blessing and a curse. We want high frequency to minimize ripple and to move the switching noise to a freuency higher than any relevant components in the circuit. However, without proper shielding, higher frequency also means that it is easier to radiate the noise throughout the lab.

Different from typical linear regulators, the output voltage of a switch-mode regulator can be higher than the input voltage, or it can even be negative. The opens up a lot of flexibility in electronic circuit design. With switch-mode regulator, and electronic board only needs to be powered by a single DC voltage, while the board itself will generate the remaining voltages, with high efficiency.
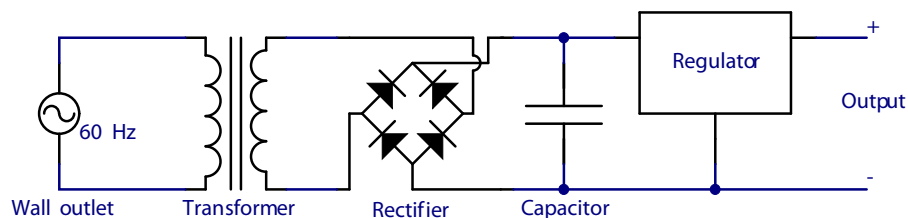
A.7.3 Linear Power Supply

Figure A.5: Linear power supply. A typical AC/DC linear power supply consists of a transformer, a full-wave rectifier (also called the bridge rectifier), one or more filtering capacitors, and a linear regulator.

The power from the wall outlet goes to the transformer. The transformer is there for two purposes. First, the wall outlet is usually at 110 Volt, while the final DC voltage that we want is around 3 Volt to 15 Volt depending on the application. The transformer acts as the first, in principle lossless, voltage reducing stage for this power supply. The second purpose is to electrically isolate the relatively high 110 Volt outlet voltage from the low voltage electronics. The rectifiers then transform the AC voltage into a DC voltage. The capacitor and the voltage regulator then filters this DC voltage to a very stable value.
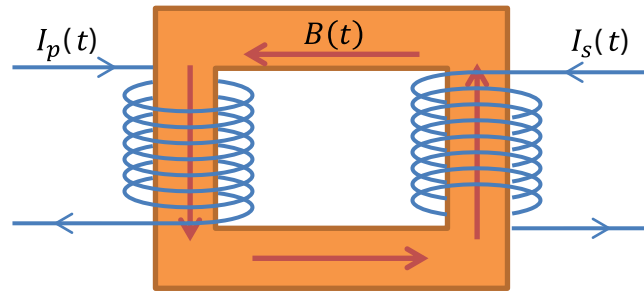

Figure A.6: Transformer.

Transformer One limiting factor in linear power supply is the transformer. Consider an ideal transformer as shown in Figure A.6, powered with a sinusoidal current with angular frequency ω. The magnetic field inside the solenoid of the output stage can be written as

$$
\begin{aligned}
B(t) &= B_0 \sin(\omega t) \\
B_0 &= \mu \frac{N I_0}{L},
\end{aligned}
$$

where L is the length of the transformer, N is the number of primary winding which is proportional to the voltage at the transformer V_0. We can deduce that

$$
B_0 \propto P,
$$

where P is the power at the input stage. For a fixed size transformer, the power of the transformer is directly proportional to the magnetic field of the transformer.

The core of the transformer acts to increase the magnetic field in the transformer by around $50,000$, reducing the required number of windings significantly. The downside is that these transformer cores have a saturation magnetic field an the order of $2\,\mathrm{Tesla}$, which is the maximum magnetic field the core can produce. This in turns limit the amount of power that can be transmitted through the transformer.

The flux of the solenoid with N windings is

$$\Phi(t) = NAB_0 \sin(\omega t)$$

Faraday's law says that the induced voltage is

$$\varepsilon(t) = -\frac{d\Phi}{dt} = -NAB_0\omega \cos(\omega t)$$

One important aspect is that the amplitude of the output voltage depends linearly with the frequency ω, so the higher the frequency, the higher the output voltage is, and the higher the power that can be converted.
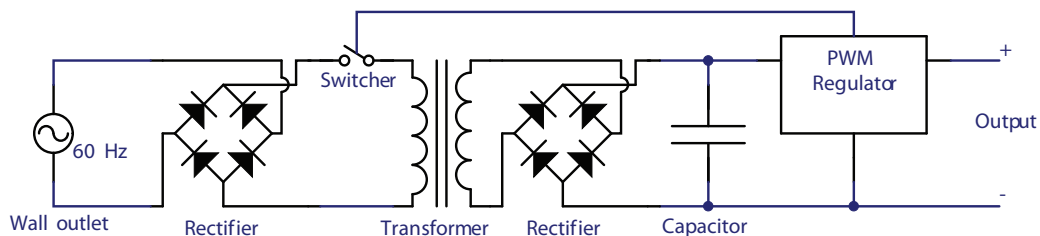
A.7.4 Switch-Mode Power Supply

Figure A.7: Switch-mode power supply. Switch-mode power supply overcomes the power limitation of the transformer by increasing the frequency at the transformer substantially.

A switch-mode power supply is used whenever a large amount of energy has to be converted. The main limitation of linear power supply, is the size of the transformers. A switch-mode power supply overcomes this limitation by switching the input voltage at higher frequency than the standard 60 Hz wall outlet frequency. This in turns reduces the linear dimension of the transformer substantially.

A.7.5   Can I Use Switching Power Supplies in a Lab?

Some atomic physics labs try to forbid switching power supplies in the lab due to the possibility that the switching noise will interfere with the experiment. Here, I want to argue that with proper design, switching power supplies can live peacefully with high precision physics experiments.

Virtually all experiments use switch-mode power supplies or regulators in one form or another. This first example is the high current power supplies used to power the magnetic coils. This power supplies generally output in the multiple-kilowatt range. A linear power supply with this much power will invariably have high cost, take a large amount of space while generating excessive amount of heat at the same time. Therefore, most people use switch-mode power supplies to power their magnetic coil, probably out of necessity. Fortunately, switch-mode noise used to power coils rarely interferes with atomic physics experiments.

The second example is computer power supplies. Virtually all computers use both switch-mode power supplies and switch-regulator. The switch-mode power supplies generate the 5 Volt and 12 Volt signal to power the motherboard and other components. However, most of the chips on the motherboard require a much lower voltage, probably

around 1 Volt or even less. Therefore, several switch-mode regulators, located on the motherboard itself, convert the 5 Volt and 12 Volt voltages to 1 Volt required by the processor. One might wonder why the primary AC/DC power supply does not generate the 1 Volt voltage directly. The answer is because the motherboard requires more than 100 Amperes at 1 Volt at peak power. Delivering this much current from the power supply to the motherboard will require a substantial size wire which is not practical.

In many atomic physics experiments, the computer also contains one or more high-precision digital to analog converters that are used to control the experiment. Again, due to proper design, switching noise is rarely an issue even with the high-precision analog components in the computer.

A third and final example that I wanted to bring up is the oscilloscope and spectrum analyzer. These remarkable devices can easily measure electrical signal with more than 150 db of dynamic range with frequency span spanning multiple of Giga Hertz. At the same time, these devices are powered by several switch-mode power supplies and regulators.

The take home message is that with proper design, switch-mode power supplies can easily operate with acceptable amounts of noise.

A.7.6 Reducing Noise in Switch-Mode Regulator

Switch mode regulators naturally produce noise at the switching frequency. This noise has to be filtered to a level that is acceptably for each application. There are 4 ways for the switching noise to propagate:

- Conductive coupling. The noise propagates directly through the wire. This is both

the easiest noise to understand and the easiest noise to filter. A simple LC filter
virtually eliminates the conductive coupling noise. The higher the frequency, the
easier it is to filter the noise from conductive coupling. It is important to design the
LC filter with enough dissipative elements so that the filter does not amplify noise at
any frequency, especially at frequencies close to the switching frequency.

- Capacitive coupling. Parasitic capacitance between conductors in the circuit acts as
  a low impedance connection for the high frequency noise. Here, it is actually better
  to reduce the frequency in order to increase the impedance from the parasitic
  capacitors.

- Inductive coupling. Mutual inductance between different parts of the circuit acts as
  a low impedance path for noise to propagate. This is probably dominated by the
  inductors which are required in all switch-mode power supplies. In this case, it is
  beneficial to increase the frequency since the impedance of the inductor is increased
  at higher frequency.

- Radiative coupling. The circuit acts like an antenna and generates electromagnetic
  waves that can interfere with other equipment.

A.7.7 Spread Spectrum Technology

A basic switcher uses a constant fixed frequency to drive the transistor switch. This poses
a fundamental problem. With enough averaging time, in principle we can always see this
delta function switching frequency noise above the background noise level. A neat trick to
alleviate this problem is to smooth out the switching frequency into a broad range of
frequencies. Even though the total noise power remains more or less constant, the noise at
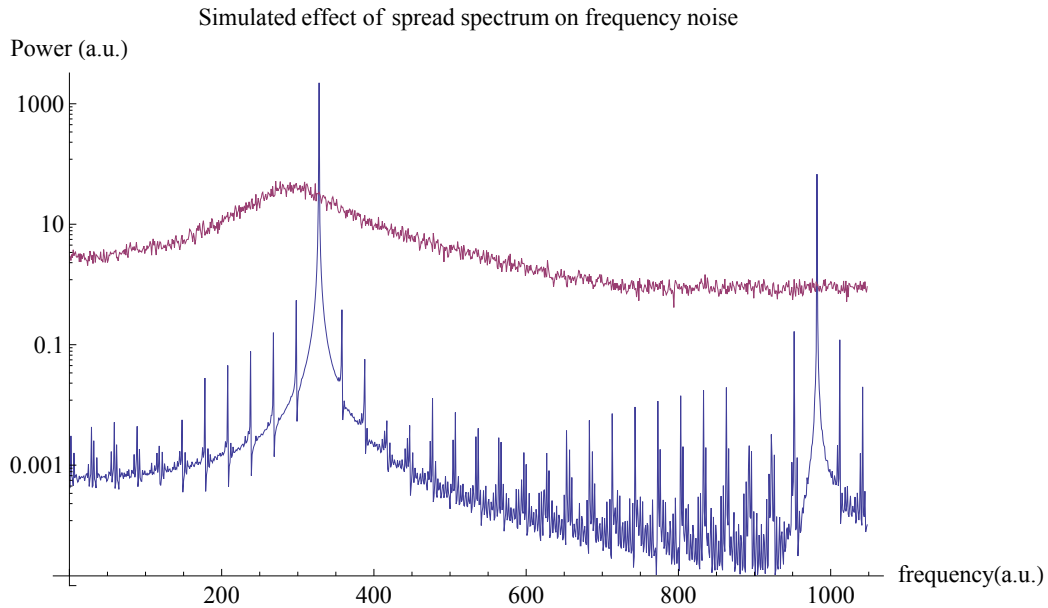
Figure A.8: Spread spectrum simulation. The blue curve shows how the noise of single frequency switch-mode regulator would look like. There are discrete peaks at the switching frequency and its harmonics. The red curve shows how the noise would look like if we randomize the switching frequency also called the spread spectrum technique. Even though the total noise power is still the same, the peak noise at any frequency is significantly lower than the single frequency switch-mode regulator.

any given frequency can be lowered significantly, probably to a level lower than the background frequency noise. This method is called the spread spectrum method.

The spread spectrum technique probably started in 1903 in several patents by Nikola Tesla. He proposed to use a frequency hopping method as a form of secure radio transmission. More relevantly however, spread spectrum technique was recently developed to reduce noise generated by clocked circuits (such as switching power supplies and digital circuits) in order to pass the Electromagnetic Compatibility (EMC) tests.

One way to widen the frequency band width is to use random or pseudo-random clock generation circuitry. This can be easily achieved using linear-feedback shift register[103]. Note that it is not enough to modulate the carrier frequency with another modulation

frequency as this would only create sidebands at the modulation frequency. A simulation

on the noise spectrum of spread spectrum switching power supply is shown in Figure A.8.