Towards Dipole Blockade Controlled-NOT Gate Using Ultracold Molecules

by

Megan Byres

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

Waterloo, Ontario, Canada, 2024

© Megan Byres 2024

Author's Declaration

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

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

Abstract

Quantum computing is a promising field that aims to achieve large increases in computational speed by taking advantage of the unique properties of quantum physics. There are many proposals for how it can be implemented in the real world, one of these being the use of Rydberg atoms. Rydberg atoms are limited by the instability of the highly excited Rydberg states, resulting in lifetimes measured in the hundreds of microseconds. Molecules can be used to perform quantum gates with a similar method to Rydberg atoms, and their lifetimes can be several orders of magnitude longer than the lifetimes of Rydberg atoms. This thesis builds on a previous work in which the ideal fidelity of this method was calculated by investigating various real world factors and their implications for the feasibility of molecules as a platform for quantum computing. Additionally, it discusses many changes and improvements to the ovens and larger vacuum system designed to perform these experiments.

Acknowledgements

I would first like to thank Professor Alan Jamison for his guidance, tutelage, and advice, without which this work would not have been possible.

I would also like to thank Andrew Urschel for teaching me all I know about machining and always being available to give guidance and advice for my projects in the shop.

In addition, I would like to thank Paul, Forouzan, Omar, Andrew, Nabeel, Max, Addison, and Marlyn for their help and friendship in the lab.

Finally, I would like to thank my parents, whose support enabled me to get to where I am today. Thank you for always believing in me and pushing me to do my best.

Table of Contents

Author's Declaration ii

Abstract iii

Acknowledgements iv

List of Figures vii

1 Introduction 1

2 Background 3
- 2.1 Laser cooling and trapping . . . 3
  - 2.1.1 Two level atom . . . 3
  - 2.1.2 Our method . . . 4
- 2.2 Ultracold molecules . . . 10
  - 2.2.1 Molecular physics overview . . . 10
  - 2.2.2 Feshbach molecules and STIRAP . . . 12
- 2.3 Quantum Computing . . . 13
- 2.4 Dipole blockade controlled-NOT gate . . . 14
  - 2.4.1 Dipole blockade . . . 14
  - 2.4.2 Three level picture and ideal fidelity . . . 15
  - 2.4.3 CNOT with Rydberg atoms . . . 16
  - 2.4.4 CNOT with Molecules . . . 16

3 Control system 19
- 3.1 Hardware . . . 19
- 3.2 Software . . . 20
- 3.3 Build guide and progress . . . 21
  - 3.3.1 DE10-Nano Setup . . . 21
  - 3.3.2 Digital to Analog Converter Setup . . . 25
  - 3.3.3 Clock Setup with Atticus . . . 26
- 3.4 Cicero and Poseidon quick start guide . . . 27

4 Vacuum system **30**
4.1 Initial design . . . 30
4.2 Vacuum tech rules and guidelines . . . 32
4.2.1 Assembly . . . 33
4.2.2 Baking . . . 36
4.2.3 Operation . . . 38
4.3 Differential pumping calculations . . . 38
4.4 Problems and solutions . . . 40
4.4.1 The ovens . . . 40
4.4.2 The NEG . . . 45
4.5 Summary and future steps . . . 47

5 Dipole blockade CNOT fidelity calculations **48**
5.1 Introduction . . . 48
5.2 Vibrational state energies for $X^1\Sigma$ LiCs . . . 49
5.3 DC electric field method . . . 51
5.3.1 Polarizability and Stark shifts . . . 52
5.3.2 Stray electric fields and noise . . . 56
5.4 Summary and future steps . . . 57

6 Conclusion **58**

Bibliography **60**

List of Figures

2.1 a) Laser configuration for optical molasses. b) Laser frequencies as seen by a stationary atom. c) Laser frequencies as seen by a moving atom. Figures modified from [6] in [9]. . . . 5
2.2 Left: Magnetic field from anti-Helmholtz coils. Right: Coil and laser configuration with left (σ^-) and right (σ^+) polarized light for a MOT. Figures modified from [6] in [9]. . . . 6
2.3 Energy diagram of Zeeman splitting leading to damping force in a MOT. Figures modified from [6] in [9]. . . . 7
2.4 Left: Sisyphus (polarization gradient) cooling [26]. Right: Gray molasses cooling [20]. . . . 8
2.5 Evaporative cooling [31]. . . . 9
2.6 Energy level diagram of a molecule showing electronic, vibrational, and rotational states with respect to internuclear separation [18]. . . . 11
2.7 Energy vs magnetic field diagrams for assembling a molecule with a Feshbach resonance. The right side shows a scenario with hyperfine coupling. Credit: Alan Jamison . . . 12
2.8 Example STIRAP path for NaLi [22]. . . . 13
2.9 Example energy level diagram of a dipole blockade [9]. . . . 15
2.10 Rydberg atom dipole blockade CNOT gate schematic [10]. . . . 17
2.11 Molecule dipole blockade CNOT gate schematic [9]. . . . 18

3.1 DE0-Nano SoC with components labeled [27]. . . . 20
3.2 DAC on evaluation board [1]. . . . 21
3.3 FPGA configuration mode switch set for EPCS configuration mode. . . . 22

4.1 CAD drawing of the full vacuum chamber [9]. . . . 31
4.2 CAD drawing of the oven (top) and its components (bottom) [9]. . . . 32
4.3 The result of accidentally using a copper gasket in a lithium vapor cell. . . . 34
4.4 Trace from a residual gas analyzer (RGA) showing pressure vs mass (in AU) of gases in the vacuum chamber. This trace was taken with the old Cs oven design at roughly 400C. . . . 40
4.5 The new design for the bottom cap of the crucible, showing the hole for the thermocouple and the notches for the macor rods. . . . 41

4.6 The clamp designed to hold the weak joint stable temporarily. Shown in place on the joint in the right side image. 42
4.7 Top: double sided blank flanges with holes drilled to mount the microchannel nozzle. Bottom: The new design for the Cs oven. 44
4.8 Model of the conductance around the hexagonal casing of the NEG pump. 45
4.9 Non-evaporable getter (NEG) pump in a new, larger, nipple. 46

5.1 Potential data (left) and fit (right) for the $X^1\Sigma^+$ state of LiCs. 50
5.2 Energies for the bound vibrational states in the $X^1\Sigma^+$ state of LiCs. 52
5.3 Polarizability data (blue dots) and fit (orange line) for the $X^1\Sigma^+$ state of LiCs. 53
5.4 Electric dipole moment data (blue dots) and fit (orange line) for the $X^1\Sigma^+$ state of LiCs. 54
5.5 Polarizabilities (top) and second order Stark shifts (bottom) for the $X^1\Sigma^+$ state of LiCs. 55

Chapter 1

Introduction

Quantum computing as a field has been drawing a lot of interest in recent years. Quantum computers have the potential to out-perform classical computers by making use of quantum phenomena to perform calculations exponentially faster. However, they are very hard to build in the real world as they are very sensitive to any outside noise or interactions and are prone to errors that are difficult to eliminate. One promising platform for quantum computers is Rydberg atoms, which use a dipole blockade to perform entangling gates. This is explained in chapter 2. However, it has been proposed that ultracold molecules have the potential to out-perform Rydberg atoms as a platform for a dipole blockade gate due to their longer lived and more stable states.

Previous work calculated the fidelity and gate time achievable in an ideal situation [9], but there are many real world factors that need to be considered to evaluate the potential effectiveness of ultracold molecules as a platform for quantum computing. The choice of molecule and states to use, the magnitude of stray fields in the lab, and the stability of the lasers and fields used in the setup can all have a big impact on this. We must also consider the method we will use to polarize the molecules in order to perform the dipole blockade. Two methods have been proposed: using a large DC electric field or using circularly polarized microwaves to polarize the ground state. The impacts of these factors and resulting hardware requirements for desired performance are analyzed in chapter 5.

In order to perform this gate using ultracold molecules in the lab, we need a vacuum chamber that produces a vapor of the desired atoms, then traps and cools the atoms before assembling them into molecules. The initial design of this system was proposed in [9], but the process of building it is not simple. To this end, I wrote a detailed guide on vacuum work, with guidelines and instructions both applicable to general systems and specific to the unique design of our system. This will be a good resource for future students in our lab as it explains the procedures used in depth and outlines known problems and precautions to take to avoid issues. I also analyzed the expected pressure differences in the three chambers of our system due to differential pumping so that we know what we can expect the final pressure in the main chamber to be. Finally, the new design for the

ovens that produce the Lithium and Cesium vapors had not yet been implemented, so there were many small changes and a few large ones I made after encountering issues with them. These changes included modifications to the hardware design for cleanliness, ease of assembly, and long-term stability, as well as a different method to produce Cesium vapor while keeping some of the benefits of the previous design when it proved impractical. I also designed and machined clamps to support several weak joints in the system, which have so far been successful in preventing leaks, and analyzed the conductance of gases to a pump that was under-performing due to being confined in a tight space. Moving this pump to a larger pipe resulted in notable improvement in its performance. The vacuum guide, differential pumping calculations, and detailed descriptions of the problems encountered and the changes made to fix them are all found in chapter 4. In addition to the vacuum chamber, we will need multiple lasers, magnetic fields, and other hardware that can turn on or off or ramp in specific ways at specific times to perform the gates. To control these components in a sequence, we will use a system composed of individual controller boards linked to a central program via Ethernet. While much of this system has already been designed, I added instructions for the inclusion of a specialized clock to ensure all the individual boards act in sync. This setup is explained in detail in chapter 3.

Chapter 2

Background

2.1 Laser cooling and trapping

2.1.1 Two level atom

In the dressed atom approach to the two level atom described in chapter 5 of [3], the atom and the laser photons are considered as a single system. First, we consider an atom with only a ground state $|g\rangle$ and an excited state $|e\rangle$ described by the Hamiltonian

$$H_A|e\rangle = \hbar\omega_A|e\rangle \tag{2.1}$$

$$H_A|g\rangle = 0 \tag{2.2}$$

This atom interacts with a laser with photons described by the Hamiltonian

$$H_L = \hbar\omega_L a^\dagger a \tag{2.3}$$

The total Hamiltonian for the system is then written as

$$H = H_L + H_A + V_{AL} \tag{2.4}$$

where V_{AL} represents the atom-laser interaction. If we set this interaction to be zero, we can find the energy difference between the uncoupled states $|g, N+1\rangle$ and $|e, N\rangle$, where N is the number of photons.

$$E_{g,N+1} - E_{e,N} = \hbar(\omega_L - \omega_A) = \hbar\delta \tag{2.5}$$

Here, δ is the detuning of the laser. Now, we consider the interaction between the atom and the laser, V_{AL}. An atom in the ground state can absorb a photon and transition to

the excited state. In a system with only one photon, we define the vacuum Rabi frequency Ω_0 as follows:

$$\frac{\hbar\Omega_0}{2} = \langle e, 0|V_{AL}|g, 1\rangle \tag{2.6}$$

This can then be generalized to a system with N photons as

$$\langle e, N|V_{AL}|g, N+1\rangle = \frac{\hbar\Omega_{N+1}}{2} \tag{2.7}$$

$$\Omega_{N+1} = \Omega_0\sqrt{N+1} \tag{2.8}$$

If we consider stimulated emission, the state $|g, N+1\rangle$ can couple not only to $|e, N\rangle$ but also $|e, N+2\rangle$, but as this is far off resonance, it can be neglected [3]. The 2x2 matrix for the full Hamiltonian is then

$$\hbar \begin{bmatrix} 0 & \frac{\Omega_{N+1}}{2} \\ \frac{\Omega_{N+1}}{2} & \delta \end{bmatrix} \tag{2.9}$$

This can be used to find the wavefunctions and energies of the two coupled states, called dressed states. The dressed atom approach is useful because it eliminates the time dependence found in the two level atom approach used in chapter 7 of[6].

2.1.2 Our method

Laser cooling and trapping is a fundamental part of any ultracold atoms experiment. The most simple laser cooling method uses the scattering force from a photon interacting with an atom as described in chapter 9 of [6]. Due to conservation of momentum, when an atom absorbs a photon, it gets a momentum kick in the direction the photon was traveling. If the photon and atom were traveling opposite directions, the atom will be slowed. When the atom re-emits the photon, it will again get a momentum kick in the direction opposite from the photon. However, spontaneous emission is equally likely to occur in any direction, so the momentum changes from many spontaneous emission events will average to zero. Therefore, if an atom is traveling toward a laser and absorbs and re-emits multiple photons, the net effect will be to slow the atom down.

This slowing, however, only works when the photon is at (or close enough to) the right frequency to be absorbed, aka "on resonance". Since the atoms are moving quite fast, there will be a Doppler shift in the frequency of the light they see, so the laser must be tuned accordingly. The Doppler shift is dependent on the velocity of the atom, though, so the frequency needed to be on resonance will change as the atom slows down. We can counteract this with a Zeeman slower. The Zeeman effect shifts the energy levels of the atom, and this shift can be used to compensate for the Doppler shift, allowing the atoms

to continue being cooled as they slow down [6]. The Zeeman slower is a solenoid with a varying number of coils, which produce a magnetic field that gets stronger as the atoms move along its length. The slower solenoid is designed so that the magnetic field strength is just right to counteract the Doppler shift at any given point for an atom with a certain initial velocity.

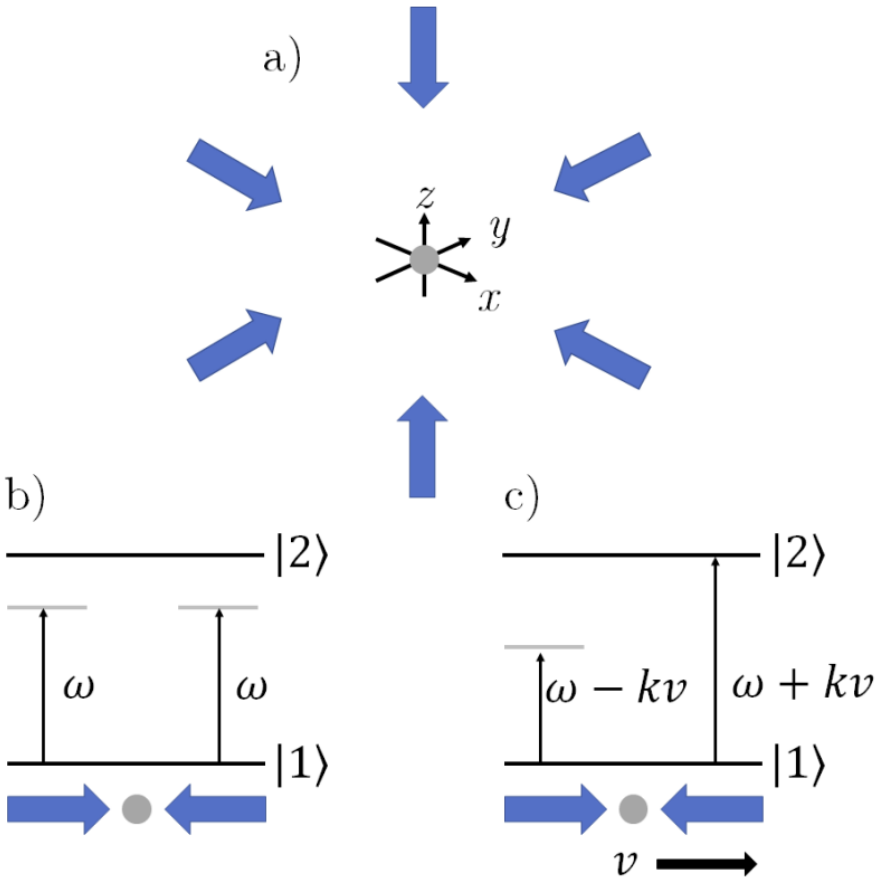


Figure 2.1: a) Laser configuration for optical molasses. b) Laser frequencies as seen by a stationary atom. c) Laser frequencies as seen by a moving atom. Figures modified from [6] in [9].

Once the atoms have passed through the Zeeman slower, we want to trap them, which requires cooling in all three directions. The optical molasses technique, described in chapter 9 of [6] and shown in Figure 2.1, uses three pairs of counter-propagating red-detuned laser beams to accomplish this. For a stationary atom, the scattering forces from these beam pairs will cancel out, but for a moving atom, the Doppler shift will cause the frequency of the laser the atom is moving towards to increase, bringing it closer to resonance, and shift the other laser in the counter-propagating pair further from resonance. Since the atom is more likely to scatter light the closer it is to being on resonance, the atom will be more

likely to scatter light from the laser it is moving towards, which will then push it back towards the center of the trap. Therefore, optical molasses creates a velocity dependent damping force that cools atoms in all three directions. However, this method on its own is not enough to create a complete trap, as there is no force present to move a stationary atom towards the center.

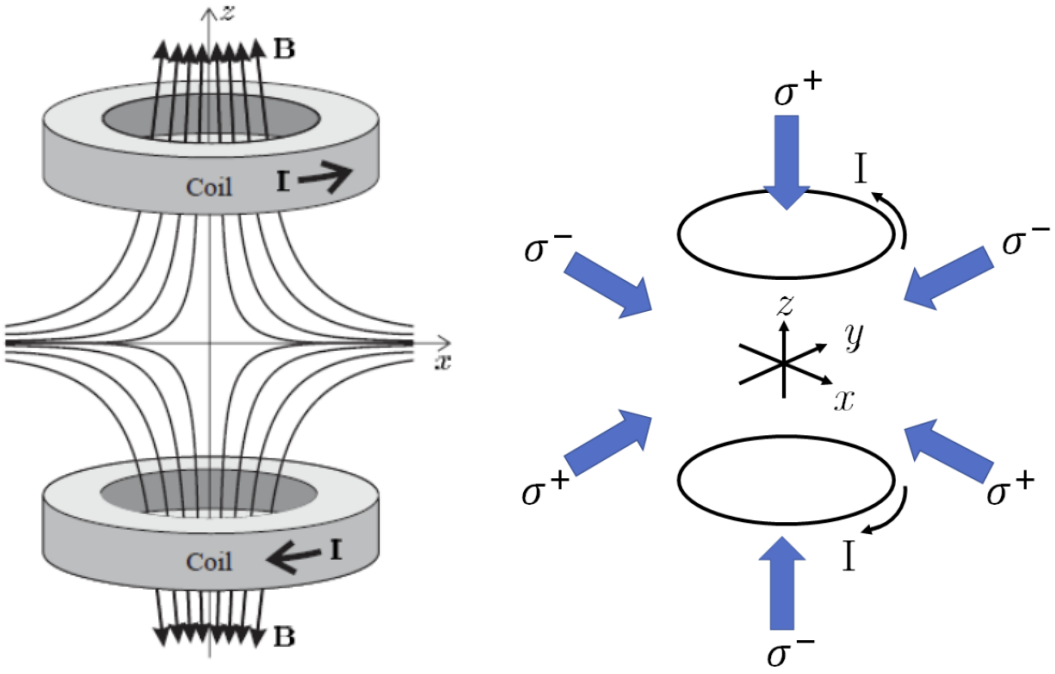


Figure 2.2: Left: Magnetic field from anti-Helmholtz coils. Right: Coil and laser configuration with left (σ^-) and right (σ^+) polarized light for a MOT. Figures modified from [6] in [9].

The trap we use in our system is called a MOT (magneto-optical trap) [6]. It uses optical molasses, with three pairs of counter-propagating left- and right- circularly polarized laser beams, in combination with anti-Helmholtz coils, to create a trap as shown in Figure 2.2. The coils produce a magnetic field that is zero in the center and approximately linear near the center. The further an atom is from the center of the trap, the larger the magnetic field it will experience, which leads to a linear increase in Zeeman splitting as the atom moves further from the center of the trap, shown in Figure 2.3.

The polarization of the counter-propagating beams in Figure 2.2 is chosen such that an atom at a negative x coordinate will likely be in the state M_J = -1 and an atom at a positive x coordinate will likely be in the state M_J = 1. The larger the negative x coordinate, the larger the downwards shift in energy of the M_J = -1 state, and similarly, the larger the positive x coordinate, the larger the downwards shift in energy of the $M_J = 1$ state, shown in Figure 2.3. As the energy of the state shifts downwards, it is brought closer to resonance with the corresponding laser. This can be seen in Figure 2.3: the laser energy

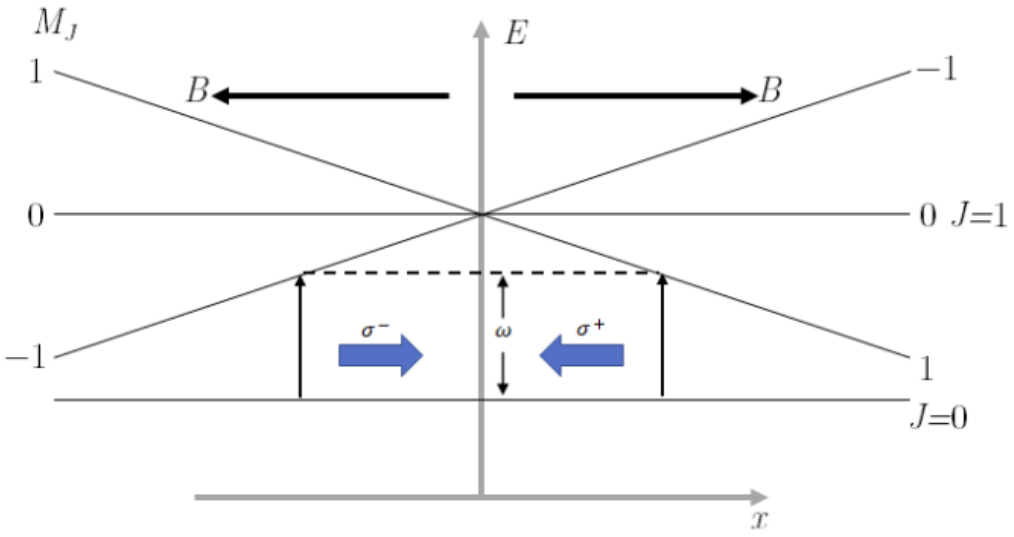


Figure 2.3: Energy diagram of Zeeman splitting leading to damping force in a MOT. Figures modified from [6] in [9].

is represented by a dotted line, and when the Zeeman level energies intersect this line, the level is on resonance with the laser. As the atom gets closer to resonance, it becomes more likely to scatter the light, and this scattering force gives it a kick towards the center of the trap. Since the shift in energy here is only dependent on the atom's position, even stationary atoms will experience a scattering force toward the center of the trap. In this way, there is both a velocity dependent damping force from optical molasses and a position dependent restoring force from the Zeeman shift due to both the beams and the magnetic field. The atoms will be cooled in the same way as optical molasses, and the cooled atoms will be pushed to the center of the trap by the position dependent force of the MOT [6]. This allows a cloud of atoms to be trapped and cooled to the Doppler cooling limit.

In order to increase the density of atoms in the trap, we want to compress the MOT. This is done by turning the current in the coils up to increase the magnetic field gradient and changing the detuning of the laser to be closer to resonance as described in chapter 14 of [3]. The result of this is that the energies of the Zeeman levels shift more quickly with changes in position, i.e. the lines for $M_J = \pm 1$ in Figure 2.3 will have a steeper slope. A smaller laser detuning means a smaller Doppler or Zeeman shift is needed to bring the atom on resonance with the laser. The combination of these two methods results in both the position and velocity dependent forces becoming stronger for a given deviation from zero, which compresses the MOT. However, the smaller detuning of the laser increases the scattering rate of atoms in the trap, which can cause unwanted heating. To compensate for this, as the laser is moved closer to resonance, its power is decreased.

The Doppler cooling limit [6] is the lowest temperature achievable using optical molasses for

most atoms. It arises due to spontaneous emission in the cooled cloud of atoms, resulting in a recoil force that gives atoms a momentum kick in a random direction each time they emit a photon. Since there will always be some absorption followed by spontaneous emission while the cooling beams are on, there will be a balance between the velocity gained from spontaneous emission and the damping from the cooling beams. This results in a nonzero minimum temperature from this cooling method. In order to cool below this limit, different techniques are needed.

For Cesium, we can use Sisyphus cooling, also known as polarization gradient cooling. This cooling method uses two orthogonally polarized counter-propagating beams to form a standing wave with a spatially varying polarization [6]. An atom moving through this standing wave will experience a shift in the energies of its lower level states that is dependent on the polarization. Shown in Figure 2.4, the energies of the lower level states M_J= 1/2, -1/2, vary with the position of the atom in the standing wave. If the atom absorbs the light from the top of a "hill", or the highest energy of the state it is in, and subsequently decays to the bottom of a "valley", or the state with the lowest energy, the atom loses energy and cools down. The red detuning of the light is chosen so that this process happens more often than the opposite, an atom moving from a "valley" to a "hill", and so the atoms overall get cooler.

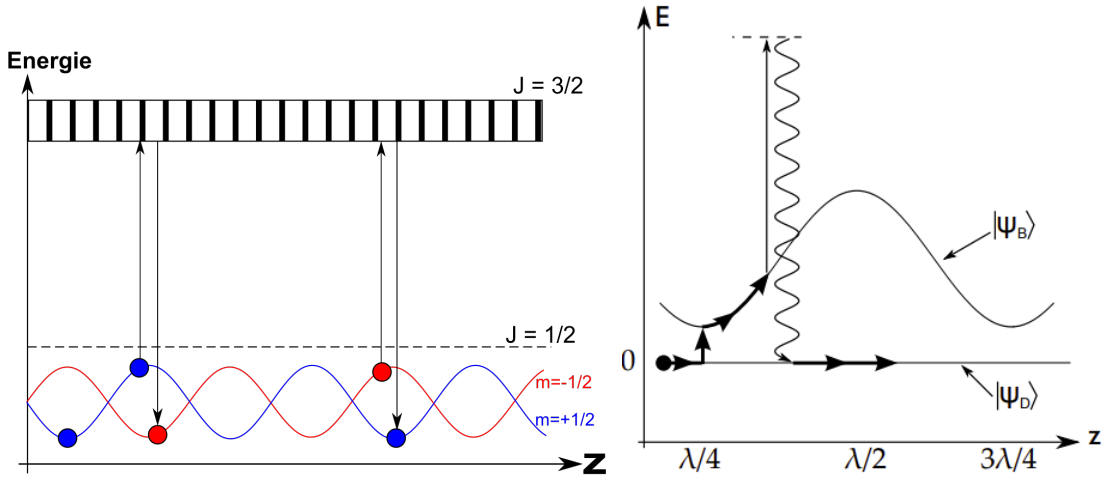


Figure 2.4: Left: Sisyphus (polarization gradient) cooling [26]. Right: Gray molasses cooling [20].

For Lithium, however, the splitting between hyperfine levels is too small for the Sisyphus cooling method to work. Instead, we will use gray molasses cooling [20]. This method is similar to the Sisyphus cooling method in that it also uses a polarization gradient, but the ground state is instead split into a "bright" state and a "dark" state. The bright state energy is spatially varying with the polarization, but the dark state is unaffected. The result of this is that only the bright state couples to the excited state used for cooling. Again, the detuning is selected so that the transition to the excited state is more likely to

happen from the top of a "hill", and if the atom then decays to the dark state, the overall effect is cooling. The specifics of this method for Lithium are explained in [7].

To transfer this cooled cloud from the MOT chamber to the main chamber, we use a technique called moving molasses, described in chapter 14 of [3]. If one laser along the axis of interest is shifted toward resonance while the other is shifted away by the same amount, the atoms will move in the direction of the laser shifted away from resonance at a velocity proportional to the frequency shift. Because of the Doppler shift, this frequency difference essentially cools the atoms in a moving reference frame, hence the name moving molasses.

Once in the main chamber, the atoms will be captured again in an optical dipole trap (ODT) [3], [6]. The presence of an electric field, in this case the field from the laser light, induces a dipole moment in atoms. The interaction between this dipole and the electric field results in a Stark shift of the energy levels of the atoms. Since the electric field of laser light is oscillating, the Stark shift will also oscillate, which is known as an AC Stark shift. This effect gives rise to a force proportional to the derivative of the AC Stark shift. The atom then moves in a potential proportional to this shift. In order to avoid scattering the light, a far off-resonance red detuned laser is used for this trap, and the atoms will be attracted to the region of highest intensity. To trap in three dimensions, we use two beams focused to the same point, which is called a crossed beam dipole trap [8].

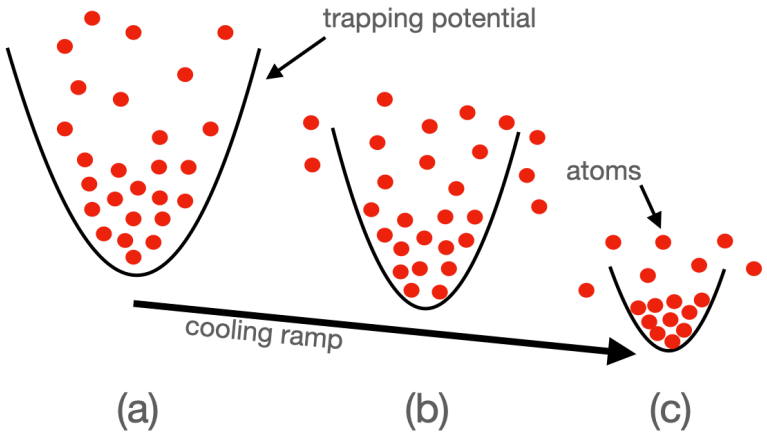


Figure 2.5: Evaporative cooling [31].

Once the cloud of cold atoms has been trapped again in the main chamber, we will cool it even further using a technique called evaporative cooling [6], [13], [25]. The atoms in the trap reach thermal equilibrium through collisions over time, with energies described by a Boltzmann distribution. If the trap depth is then decreased, the atoms with the highest energies escape, as shown in Figure 2.5. This is essentially cutting the high energy tail off the distribution. Allowing the atoms to re-thermalize then results in a new Boltzmann distribution of energies with a lower overall temperature, as the atoms that escaped were those with above average energy. This can be repeated to cool the atoms to the desired temperatures, on the order of 100 nK.

2.2 Ultracold molecules

2.2.1 Molecular physics overview

This overview of the physics of diatomic molecules is largely a summary of chapter 2 of [5]. In the simplest picture, a molecule can be represented as a rigid structure, with the nuclei at fixed positions that minimize the total energy. The electrons bind the nuclei together, and their average distribution determines the minimum energy equilibrium position of the nuclei. This model is called the rigid molecule, and it is a good starting point for understanding molecular states.

The Hamiltonian for a free molecule at rest has three components: the kinetic energy of the electrons, the kinetic energy of the nuclei, and the potential energy. The potential energy is comprised of the repulsion between the electrons, the repulsion between the nuclei, and the attraction between the electrons and nuclei. This Hamiltonian is so complicated that the Schrodinger equation has no exact solution even for the simplest molecule, so some approximations are needed to get useful information from it.

Electrons are much lighter than nuclei and therefore move much faster. In the adiabatic approximation, this is taken to mean that the electron cloud follows the movement of the nuclei adiabatically as long as the kinetic energy of the nuclei is small compared to the energy of the electrons. This means that the electron distribution depends only on the positions of the nuclei, allowing for the separation of the nuclear and electronic wavefunctions. This results in a set of coupled equations for the nuclear and electronic wavefunctions.

In the Born-Oppenheimer approximation, the coupling between the movement of the nuclei and the electron distribution is ignored, allowing the coupled equations for their wavefunctions to be completely separated. These equations can now be solved for the electronic and nuclear wavefunctions, although the only system simple enough to solve exactly is the H_2^+ molecule. This approximation is only strictly valid for nuclear frameworks that are rigid and non-rotating. This allows the electronic states to be represented by potential energy curves found from the wavefunction equations.

The electronic state potentials are quite complicated to solve for all but the most simple molecules. For diatomic molecules, there is only one internuclear distance, R, which simplifies things somewhat. The limit R=0 can be approximated as a united atom with a mass determined by the sum of the two nuclei of the molecule, and this "atom" will have well-defined electronic states. At the other end, R$\rightarrow \infty$, the molecule can be viewed as two separate atoms, and the molecular electronic states are then combinations of the states of the two atoms. While these two limits are comparatively simple, the question of what the potential looks like in between is much more complicated.

To approximate the shape of the potential curve, the total electronic wavefunction can be written as a product of the wavefunctions of the individual electrons if interactions between electrons are neglected. The individual electron wavefunctions, called molecular orbitals, are modeled as linear combinations of atomic orbitals of the two separated atoms, which

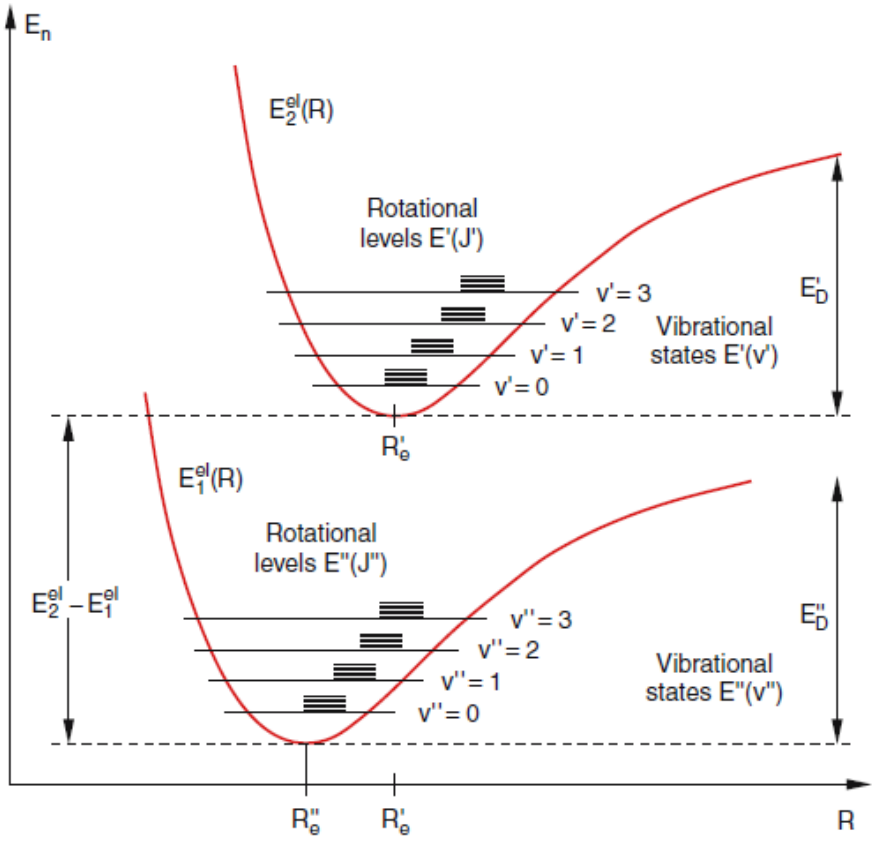


Figure 2.6: Energy level diagram of a molecule showing electronic, vibrational, and rotational states with respect to internuclear separation [18].

is called the LCAO approximation. This molecular orbital depends only on the position of the individual electron for a fixed internuclear distance R. Neglecting the interactions between electrons is inaccurate, but it can be indirectly taken into account in the potential. The electrons move in a potential determined not only by the nuclei, but also the average distribution of all other electrons. This combined potential is known as the Hartree approximation, and while it still neglects the effect of the electron in question on the rest of the distribution, it is an improvement compared to neglecting the other electrons entirely.

In the Born-Oppenheimer picture discussed so far, we have been considering the rigid molecule where the nuclei are fixed at their equilibrium positions, but this is not the case for real molecules. The nuclei can oscillate around their equilibrium separation distance and rotate around multiple axes, resulting in differences in energy. Therefore, within the electronic states of the molecule, there are also vibrational and rotational states with different energies, as shown in Figure 2.6. The further the nuclei move from their equilibrium position in each oscillation, the higher the energy of the vibrational level. Within each vibrational level are many rotational levels, but for our purposes we will be mainly consid-

ering the ground rotational state, i.e. a non-rotating molecule.

2.2.2 Feshbach molecules and STIRAP

Once we have cold cesium and lithium atoms, we want to assemble them into molecules. To do this, we use a Feshbach resonance as explained in chapter 16 of [3], which occurs when a state with two free atoms has the same energy as a bound state at a specific magnetic field, as shown in the left plot of Figure 2.7. The energy of the two free atoms decreases with increasing magnetic field, while the energy of the bound molecule increases with magnetic field. If there is some kind of coupling between the states, there will be an avoided crossing, as shown in the right plot of Figure 2.7. In our case, this avoided crossing arises from coupling in the hyperfine states. Starting with a magnetic field above the resonance and decreasing it adiabatically will cause the atoms to follow the lower curve in the figure, resulting in a molecule.

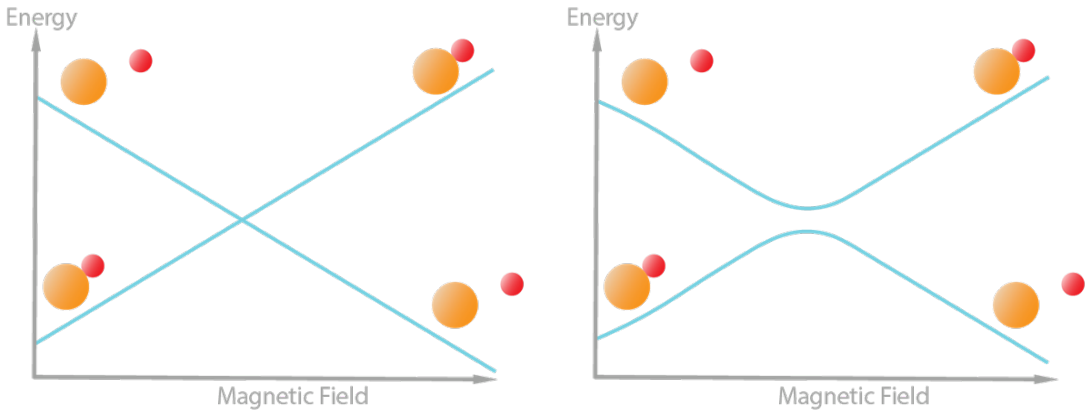


Figure 2.7: Energy vs magnetic field diagrams for assembling a molecule with a Feshbach resonance. The right side shows a scenario with hyperfine coupling. Credit: Alan Jamison

Feshbach molecules are formed in a very loosely bound state, often a highly excited vibrational state. In order to move these molecules to their ground state, we use a technique called STIRAP, which stands for stimulated Raman adiabatic passage [30], [21], [22]. STIRAP uses three levels: the initial high vibrational state (state 1), an intermediate vibrational state in an excited electronic state (state 2), and finally a more deeply bound state (state 3), as shown in Figure 2.8. The goal is to transfer the molecules from state 1 to state 3 through the intermediate state 2 without significantly populating it. To do this, we use two lasers to couple states 1 and 3 to state 2. Initially, only the laser coupling state 3 to 2 is turned on, making state 1 an effective dark state as discussed in gray molasses cooling. As the 3-2 laser is ramped down and the 1-2 laser ramped up, the dark state becomes first a mix of states 1 and 3 before eventually becoming state 3 as the 3-2 laser is turned off. Because the dark state does not interact with the lasers, if this state changes adiabatically, the states of the atoms will follow. Therefore they can be transferred from

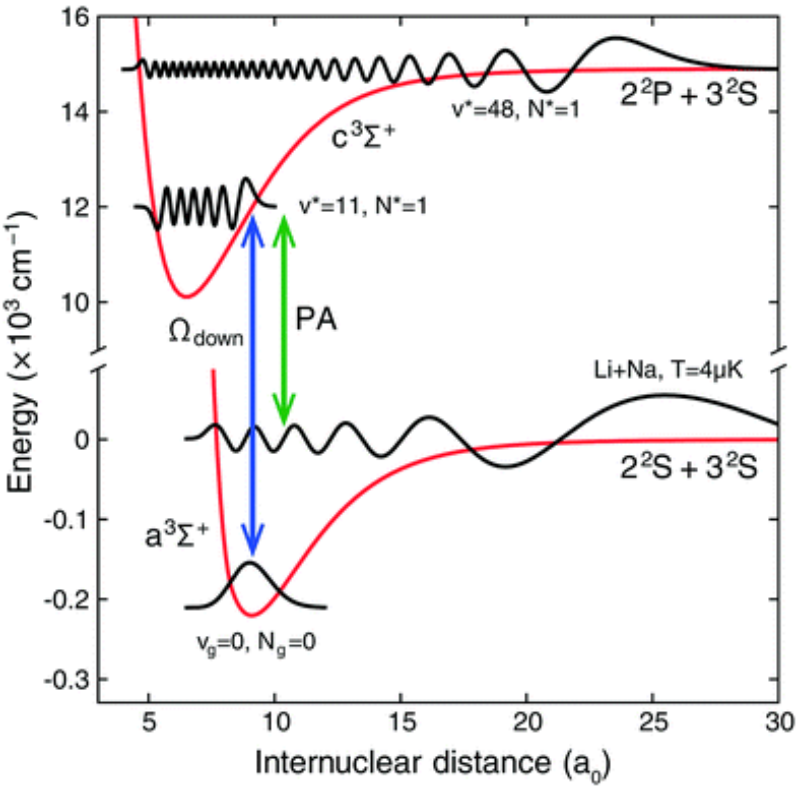


Figure 2.8: Example STIRAP path for NaLi [22].

state 1 to state 3 without significantly populating state 2. The exact levels that will be used for this process in our experiment have yet to be determined.

2.3 Quantum Computing

Quantum computers have a theoretical advantage over classical computers that arises from the way information is stored and handled. Classical bits store one value at a time, either 0 or 1. Quantum bits, or qubits, can store two values using a superposition of the two basis states. Qubits can also be entangled with one another, so that an operation on one qubit can also affect the other. The disadvantage to this system is that any interference or measurement will cause the superposition to collapse to a single state, so to gain advantage over classical computers, an ideal quantum computer would have to be completely noiseless and error-free. This is incredibly hard to achieve in the real world, and we are currently nowhere near a quantum computer advanced enough to execute something as complex as Shor's algorithm. Many current quantum computers fall into the NISQ (noisy intermediate-scale quantum) category [2]. NISQ systems contain on the order of 100 imperfect qubits, but even this can have an advantage over classical computers in some applications. For example, lattice QFT simulations in nuclear physics have a high computational cost, but a NISQ system in conjunction with a classical computer could reduce this. One proposal is

to use a variational quantum simulator algorithm to reduce the number of parameters in the classical algorithm, improving the computational cost [14]. Ultracold molecules could be a good platform for a NISQ system thanks to strong interactions and comparatively long lifetimes.

Quantum algorithms are composed of a series of quantum gates, and a complete set of gates can be used to construct any quantum algorithm. A complete set of gates must include single qubit gates, which are generally easy to perform, and at least one two qubit gate. The two qubit gate we will be using is the controlled NOT gate, written in matrix form as

$$CNOT = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{2.10}$$

In this gate, the target qubit has a bit flip dependent on the state of the control qubit. If the control qubit is in state 0, the target is unchanged, but if the control is in state 1, the target is flipped. This gate is chosen as it can be accomplished using a dipole blockade, discussed in the next section.

2.4 Dipole blockade controlled-NOT gate

2.4.1 Dipole blockade

Two dipoles will interact with each other according to the dipole-dipole interaction potential

$$V_{dd}(\mathbf{r}) = \frac{1}{4\pi\epsilon_0}\frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|^3}(\mathbf{d}_1 \cdot \mathbf{d}_2 - 3(\mathbf{d}_1 \cdot \hat{\mathbf{r}})(\mathbf{d}_2 \cdot \hat{\mathbf{r}})) \tag{2.11}$$

If one atom or molecule is in a state with a large dipole moment, called the dipole state, it will cause the dipole state energy level of a nearby atom or molecule to shift proportional to this interaction energy. This can be used to form a dipole blockade, in which one atom being in the dipole state will prevent the other from transitioning to this state. A laser tuned to the ground to dipole state transition will no longer be on resonance if the dipole state energy is shifted in this way, and therefore the transition to the dipole state will be suppressed [28]. This can be used to perform a CNOT gate, as illustrated in Figure 2.11 and described in [11]. If the control qubit begins in state 1, it will not transition to the dipole state, allowing the target qubit to transition through the dipole state and flip. However, if the control qubit begins in state 0, it will transition to the dipole state, shifting the dipole state energy off resonance for the target. The target then does not transition to the dipole state, instead remaining in its initial state. Therefore we have a controlled

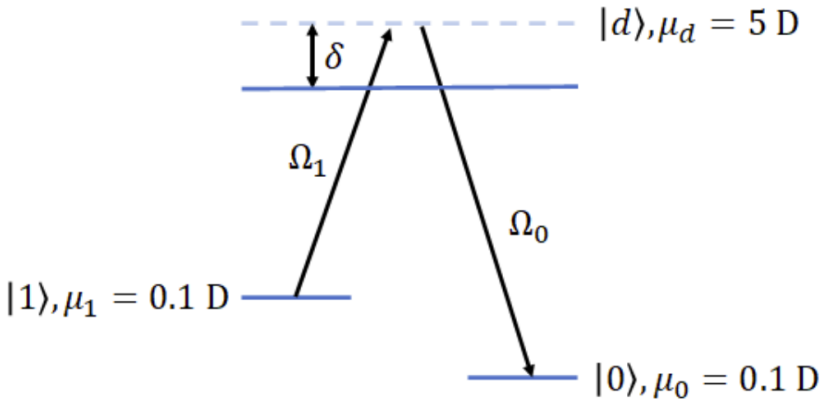


Figure 2.9: Example energy level diagram of a dipole blockade [9].

NOT gate: a NOT gate for the target qubit that is dependent on the state of the control qubit.

2.4.2 Three level picture and ideal fidelity

The three level picture is similar to the two level picture, except instead of only a ground and excited state, we have two storage states, $|0\rangle$ and $|1\rangle$, and a third state, in this case the dipole state $|d\rangle$. We also have two lasers, tuned to the $|0\rangle$-$|d\rangle$ and $|1\rangle$-$|d\rangle$ transitions, with respective Rabi frequencies Ω_0 and Ω_1 defined as follows:

$$\frac{\hbar\Omega_0}{2} = \langle 0|V_{AL}|d\rangle \tag{2.12}$$

$$\frac{\hbar\Omega_1}{2} = \langle 1|V_{AL}|d\rangle \tag{2.13}$$

Several approximations are made for simplicity. First, we assume no coupling between the states $|0\rangle$ and $|1\rangle$. We also assume the two lasers have the same detuning, δ. Using the same method discussed in section 2.1.1, we get the three level atom Hamiltonian matrix

$$\hbar \begin{bmatrix} 0 & \frac{\Omega_0}{2} & 0 \\ \frac{\Omega_0}{2} & \delta & \frac{\Omega_1}{2} \\ 0 & \frac{\Omega_1}{2} & 0 \end{bmatrix} \tag{2.14}$$

However, the assumption that there will be no transition between $|0\rangle$ and $|1\rangle$ is not the case in a real scenario, as the dipole blockade energy shift is finite. In [9], the fidelity due to this finite detuning is calculated as:

$$F = |\langle 0|\hat{U}(t_G)|0\rangle|^2 \tag{2.15}$$

This is simply the probability that an atom that starts in the state $|0\rangle$ is still in that state after the gate time t_G. The calculations shown in [9] find the gate time and fidelity:

$$t_G = \frac{\sqrt{2}\pi}{\Omega_0} \tag{2.16}$$

$$F = (1 - \frac{\Omega_0^2}{2\delta^2}) \cos^2 (\frac{\Omega_0^2}{4\delta} t_G) \tag{2.17}$$

An additional assumption made in these calculations is that, since the transition is limited by the lower of the two Rabi frequencies, we can set $\Omega_1 = \Omega_0$. This result shows that, for CsLi with a detuning of 200 kHz, we can achieve 99.9 percent fidelity with a gate time of 2 ms, which is very good compared to the coherence time of 1 to 10 s. While this fidelity takes into account the finite dipole blockade energy shift, it does not consider other factors such as which states are used in a real molecule and noise in fields and lasers. The effects of these and other factors on fidelity are discussed in chapter 5.

2.4.3 CNOT with Rydberg atoms

Rydberg atoms are atoms, usually alkalis, that are excited to a very high electronic state so that the valence electron is on average quite far from the nucleus. This results in a dipolar interaction between two neighboring Rydberg atoms that can be used to perform a CNOT gate [11], [24], [28], [10] as shown in Figure 2.10.

If the control atom starts in state $|0\rangle$, pulse 2 will not excite it to the Rydberg state $|r\rangle$, leaving the target able to be excited. When the target is excited to $|r\rangle$ and back by pulse 3, it gains a phase π. However, if the control atom starts in state $|1\rangle$, pulse 2 will excite it to the Rydberg state $|r\rangle$. This will shift the energy levels of the target atom such that it cannot be excited to $|r\rangle$ by pulse 3 and therefore its phase is unchanged. Pulse 4 then returns the control to $|1\rangle$. This is called a controlled phase (Cz) gate, as the phase shift gained by the target is dependent on the state of the control [10]. The Cz gate can easily be converted to a so-called H-Cz CNOT gate by applying a Hadamard rotation to the target atom before and after the Cz gate (pulses 1 and 5 in Figure 2.10).

2.4.4 CNOT with Molecules

A similar scheme has been proposed for ultracold molecules in [4]. Although even the most polar molecules have smaller dipole moments than Rydberg atoms, they have the advantage of longer coherence times and a more stable dipole state, so they could potentially perform better overall as a platform for a dipole blockade CNOT gate. The dipole blockade energy shift for Rydberg atoms is commonly on the order of 1 GHz while molecules can only

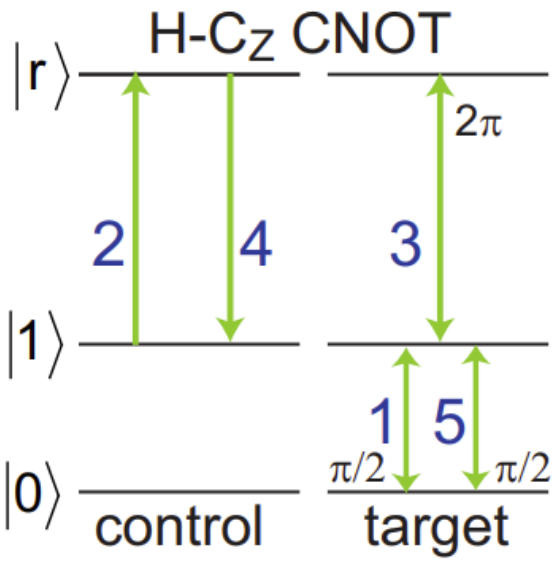


Figure 2.10: Rydberg atom dipole blockade CNOT gate schematic [10].

achieve around 100 kHz, but the Rydberg state's instability results in coherence times of only around 100 μs while molecules can have coherence times up to around 10 s. If we look at the ratio of coherence time to gate time as a measure of how good a qubit is, molecules have the potential to be 10 or even 100 times better than Rydberg atoms, making them a promising platform to study. With molecules, the dipole state is usually the rovibrational ground state, and the stability of this state allows for the CNOT gate to be performed directly rather than through the H-Cz method as we are no longer trying to minimize time spent in this state. Although the energies of the levels shown are reversed (the dipole state will have a lower energy than the storage states), Figure 2.11 provides a good visual for how the CNOT gate is executed. A π pulse tuned to the $|0\rangle$ to $|d\rangle$ transition is applied to the control qubit, followed by a π pulse tuned to the $|1\rangle$ to $|d\rangle$ transition applied to the target qubit. If the control qubit was initially in $|0\rangle$, it will transition to the dipole state, shifting the energy of the target's dipole state so that the transition will be blocked and the target will remain in state $|1\rangle$. However, if the control qubit was initially in state $|1\rangle$, the first pulse will not push the transition to the dipole state, allowing the target to be excited to the dipole state, performing a bit flip. In this way, the state of the control qubit determines whether the target will flip, resulting in a CNOT gate.

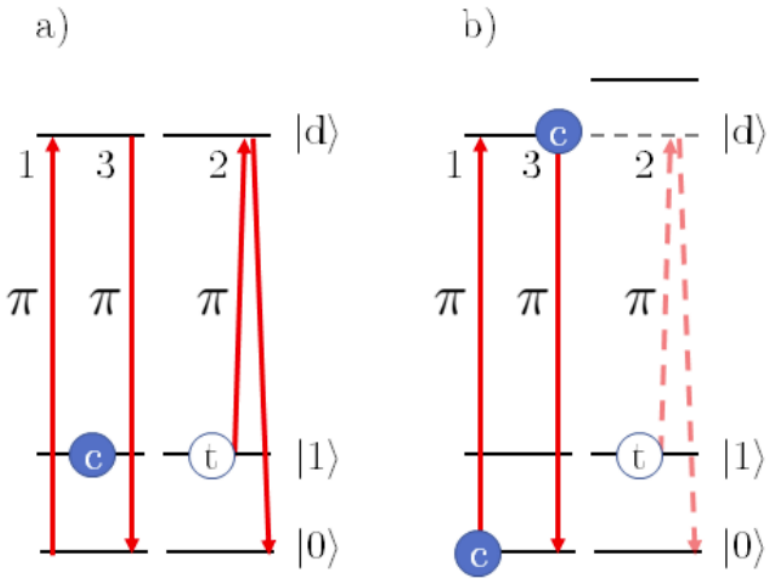


Figure 2.11: Molecule dipole blockade CNOT gate schematic [9].

Chapter 3

Control system

In our experiment, we will need precise timing for several things, such as ramping electric and magnetic fields and turning lasers on and off. Our control system will consist of a series of controllers run by a central program that will do all of the necessary changes with the correct timing.

3.1 Hardware

Each individual module of the control system has two main components- a SoC (system on a chip) and a DAC (digital to analog converter). We use a DE0-Nano SoC, as shown in Figure 3.1. It includes a FPGA (field-programmable gate array), which enables us to configure it to do exactly what we need for our system. We configure it to be able to communicate with Cicero (discussed in the software chapter) so that we can control many electronic components in the lab, such as magnetic field coils and laser shutters, remotely from the same computer. Control from the same place allows the components to be controlled in sync with one another, for example closing two shutters at the exact same time. the FPGA has an Ethernet port which is used to communicate with the computer, and the GPIO output pins are used to communicate with the DAC. More details can be found in its manual [27].

The DAC used is part of the EVAL-AD5791SDZ board, shown in Figure 3.2. Its purpose, as the name suggests, is to convert the digital signals from the SoC to analog voltages needed to control things in the lab. More details can be found in its manual [1].

In order to ensure all SoCs output to their DACs at the correct times, they need a clock signal. A simple function generator won't work as the clock itself needs to start and stop at the right times in the program sequence, so we use the NI PXIe-6535 digital IO card to create the clock signal as it has the ability to communicate with the computer. This will be explained more in depth later in the chapter. Since all SoCs are connected to the same clock signal, it is important to prevent crosstalk between them through the clock wires.

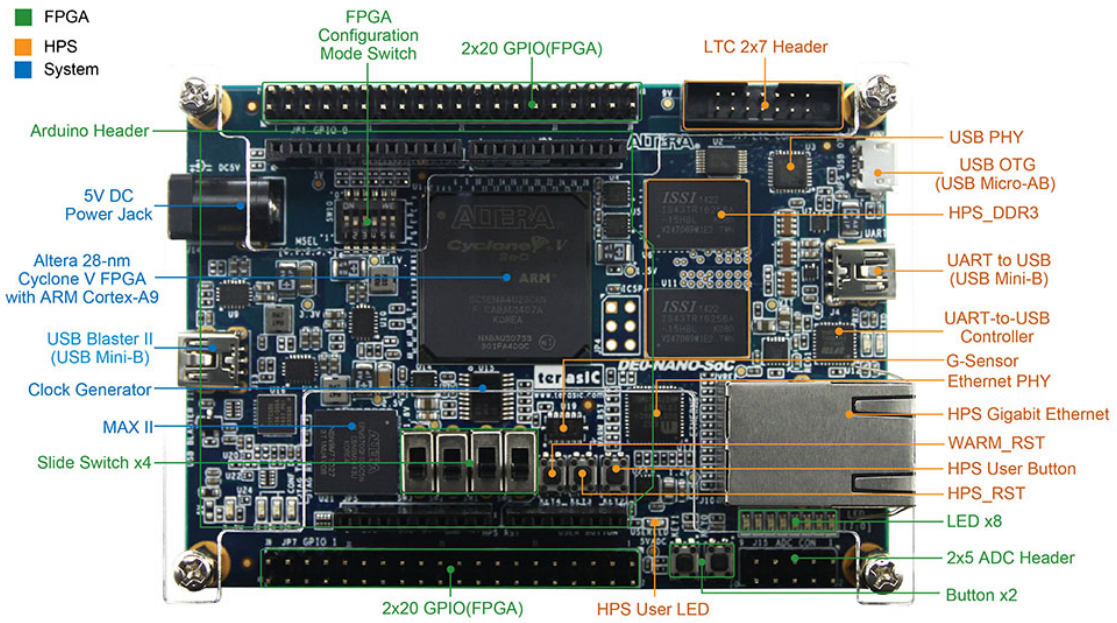


Figure 3.1: DE0-Nano SoC with components labeled [27].

To do this, we use clock isolators for each SoC. The clock isolators are opto-isolators that transfer the signal using light, preventing voltage noise from being transmitted.

3.2 Software

There are two main kinds of software needed for the control system: the software needed to set up and program the boards, and the software needed to coordinate and run the system.

Quartus Prime is a software that can be used to develop and install hardware designs for FPGAs. We use it to compile and load the design we want onto our SoCs. Aside from programming the FPGA, we also want to set up and program a Linux system on the SoCs. For this, we use balenaEtcher to flash a disk image onto the DE0-Nano's microSD card. This establishes the Linux system. Next, we use a software called PuTTY, which enables connection to the Linux terminal of the DE0-Nano using either USB or Ethernet. Here, we finish the Linux system setup and add the correct passwords and permissions. Finally, WinSCP enables the computer to transfer the necessary files to this Linux system. This process will be explained in depth in the next section.

The main software used to run the control system is called Cicero Atomic Suite. Cicero connects to other servers on the computer and can run custom voltage sequences through this. The software Poseidon connects to the DE0-Nano SoCs over Ethernet and establishes a server that Cicero can communicate with. Similarly, the software Atticus establishes a server that Cicero uses to communicate with the NI clock. Cicero acts as the central control



Figure 3.2: DAC on evaluation board [1].

system, and communicates with the individual boards through the servers established by Poseidon and Atticus. The details of establishing this communication can be found in the next section.

3.3 Build guide and progress

This guide contains the detailed step by step instructions to set up and use the control system discussed in chapter 3, written by undergraduate student Aryan Prasad for the system design outlined in [16]. I contributed the entire section on the NI clock setup along with some small updates to the rest, but the guide is here in its entirety as having the NI clock section alone would not make sense without the context.

3.3.1 DE10-Nano Setup

- Using a pen or other pointed tip, set the board's dip switches to EPCS configuration mode MSEL $[4\ldots0] = 10010$. The state of switch 6 does not matter. The MSEL configuration informs the device of how a hardware design will be loaded onto the FPGA.
- Download the folder FIFO_Version from the group's Dropbox. This folder contains the hardware design that is programmed onto the FPGA.
- Open Quartus. Under File/Open Project, find soc_system.qpf and open it. Note that the linked project was built in Quartus Standard version 21 for the FPGA on

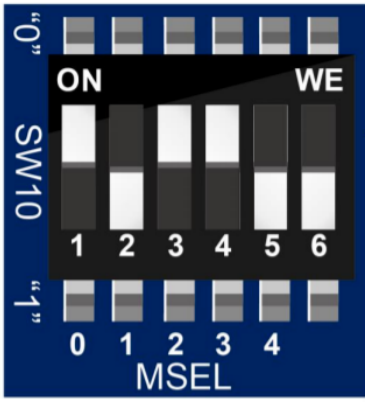


Figure 3.3: FPGA configuration mode switch set for EPCS configuration mode.

the DE10-Nano or in Quartus II 14.1 for the DE0-Nano.

- Compile the project. You will see many warnings throughout the process: these can often be ignored and the design will still work, but it should be a long-term goal to address them.
- Once compilation is finished, Quartus will generate a .sof file. This can be used to directly program the FPGA, but we don't want that. Instead we convert it to a .jic file, which programs the EPCS device. Go to File/Convert Programming Files.
- Select JTAG Indirect Configuration File (.jic) as the file type.
- Click the ellipses (...) by Mode. Our device family is Cyclone V. Select that and find EPCS64 in the list. Click that and click OK. This lets the converter know which device we are targeting.
- Click the ellipses on the same line as File name. This allows you to name and select where the .jic file will be outputted. Historically, we've called the file output_file_EPCS64.
- Towards the bottom, click Flash Loader and then Add Device... . Find the Cyclone V menu and select 5CSEBA6. That is the model of FPGA we are working with.
- Click SOF Data and then Add File... . Find soc_system.sof and select it. This provides the converter with a programming file to convert. The other settings don't matter for our purposes, so just leave them alone.
- Click Generate and then close the menu.
- Power up the DE10-Nano with a 5V jack and connect it to the computer via a mini USB cable (ports J14 and J13 on the board, respectively). The connections are the same for the DE0-Nano).
- In the Quartus navigation bar, select Tools/Programmer. This is where we will actually load the hardware design onto the DE10-Nano.

- Click the button labeled Hardware Setup... and a new window will open. In the drop down select DE-SoC.
- Click Auto Detect. If a popup prompts you to select a device, pick 5CSEBA6 (or 5CSEMA4 for the DE0-Nano). This process finds the devices on the DE10-Nano board that you can program. If this fails, you can manually add them through Add Device.... The FPGA will be found under the Cyclone V family as 5CSEBA6U23 (5CSEMA4U23 for the DE0-Nano) and the hard processor under Soc Series V as SOCVHPS.
- Double click the "none" in the File section next to the FPGA device (5CSEBA6U23) and select the .jic output file generated from the previous steps.
- Check the box Program/Configure and click Start to flash the design.
- Once this has finished, power off the DE10-Nano and unplug the mini USB. Remove the microSD card from the underside of the device and insert it into the computer via a microSD card reader.
- Download the Linux Console disk image from Terasic, the manufacturer of the board. Make sure you have selected the .iso file and not the .img file with the same name. Loading an operating system onto the board will allow us to interact with the hard processor directly.
- Use balenaEtcher to flash the disk image to the SD card. Skip the validation or it will fail.
- Once finished, put the SD card back into the DE10-Nano and power up the board. At this point, the EPCS device will program the FPGA.
- Connect the device to the computer using mini USB port J4.
- Open Device Manager on the PC and look for the serial port ID for the board under Ports. It should be of the form COM followed by a number.
- Open PuTTY and select serial mode. Type the serial port found in the previous step and set the Speed (rate of communication) to 115200.
- Click Open. A blank window will appear. Power cycle the DE10-Nano (turn if off and on again). Following some bootup text, a Linux terminal will appear in the window. Log in as root.
- Move to the directory /etc/init.d.
- Type vi runcode.sh. This creates a file named runcode.sh and opens it in the text editor Vi.
- Copy the following code. At a high level, this script sets the IP address of the device, allowing us to communicate with it through the Cicero Suite.

  #!/bin/bash

  udhcpc

ifconfig eth0 192.168.1.2XX netmask 255.255.255.0 up

route add default gw 192.168.1.1

echo “Hi”

- With the code copied to the clipboard, type i in the Linux terminal and then right click anywhere in the terminal window to paste the code.
- Each device needs to have a unique IP address. As such, change XX in the above code to a different number for each board. If starting from scratch, begin with 00, then 01, 02, and so on (skip 22).
- Press escape, type :wq, and press enter. This saves the file and exits the text editor.
- Be sure to make a list somewhere of which device has which IP address. Each board comes with a sticker ID that starts with “SoC”. You can use them to distinguish boards.
- Back at the command line, type chmod 777 runcode.sh. This command gives all users the permission to read from, write to, and execute the file.
- Change the password to w0lfg4ng by typing passwd root. Wolfgang Ketterle is the Director of the MIT-Harvard Center for Ultracold Atoms, where the Cicero Suite and FPGA control system were developed. This is the password that the Cicero Suite is programmed to use when attempting to connect to the boards.
- For whatever reason, the DE10-Nano is very slow to connect through SSH over Ethernet compared to the DE0-Nano. We can mitigate this by changing a setting in the file /etc/ssh/sshd_config. Return to the home directory, then type cd /etc/ssh, and finally vi sshd_config to open the file.
- Find the line that says UseDNS and change it to UseDNS no. Ensure that there is not a # at the start of the line, as that indicates that the line is commented out and thus ignored. This should make the connection sufficiently fast to work with the Cicero Suite, but if it doesn’t, you can remove the password from the system by typing passwd -d root.
- Return to the home directory, then type cd /etc/ssh.
- Now type update-rc.d -s runcode.sh defaults. This makes the script from earlier run on system startup, which establishes a server on the DE10-Nano that can be connected to. The Sending discover... messages show that the server is active and seeking a connection.
- Unplug the DE10-Nano from the computer, and then connect it and the computer to the lab’s internal router via Ethernet cables.
- Open WinSCP and connect to the DE10-Nano using the IP address specified for the board in runcode.sh.
- Navigate to the /home/root directory and upload fifo_test. The Cicero Suite checks this file whenever it attempts to connect to a board.

- Right click fifo_test, select Properties, and change rights to 7777.

3.3.2 Digital to Analog Converter Setup

On the EVAL-AD5791SDZ board, there are pins positioned throughout the board that can be connected in different ways through caps that connect two pins when inserted on top of them. These pins are labelled as LK1 through LK11 (link number). Some pins have multiple "on" configurations (labelled A, B, C, etc.). Configure the pins as follows:

- LK1 Inserted on A
- LK2 Removed
- LK3 Removed to output from V0, Inserted to output from V0 BUF. The default is to output from V0, but in the case of a faulty or broken V0 connector, V0 BUF can be used instead with no other changes to the setup. However, if one is switched to V0 BUF, all of them must be switched to V0 BUF, otherwise there will be weird delays in the signal. Also of note is that the signal from V0 BUF is very noisy compared to that from V0, so using V0 is always the first choice.
- LK4 Removed
- LK5 Inserted
- LK6 Removed
- LK7 Removed
- LK8 Inserted on A. The datasheet says connecting the pin on A results in power being provided from the IOVCC pin. However, we are instead providing power from port J11. Why not connect on C, then? It turns out that connecting on A and C both allow the system to get power from J11. Connecting on A has the added bonus that VCC is seen on the IOVCC pin, which provides power for the isolator circuits on the bridge.
- LK9 Inserted on A
- LK10 Inserted on A
- LK11 Removed

To power the DAC, make the following connections: +3.3 V and its ground to port J11 and +15 V, -15 V, and their shared ground to port J13. In electronic circuits, VCC and VDD denote positive power supplies, while VEE and VSS denote negative ones. Connect accordingly.

The DE10-Nano needs to tell the DAC what kind of output to generate. It sends digital signals from its GPIO pins to the J12 port of the DAC. The communication is based off of SPI protocol. The inputs are described in detail in the DAC's datasheet. Use jumper wires to make these connections:

- GPIO 0 Pin 1 to SYNC
- GPIO 0 Pin 2 to SCLK
- Pin 3 to DGND
- Pin 4 to SDIN
- Pin 5 to LDAC
- Pin 7 to CLR
- Pin 8 to DGND
- Pin 9 to RESET

All of these pin outputs are programmed in the FPGA design.

3.3.3 Clock Setup with Atticus

The DE10-Nano requires a clock signal to know when to output to the DAC. Without it, it won't do anything. Note that a previous version of this guide said that if NI hardware was not available, a function generator could be used instead. This does NOT work for multiple boards as the function generator clock cannot stop and start at the right points in the sequence, which results in funky delays for multiple signals (each will start at a different time when they are meant to start in sync). As such, we require more sophisticated hardware from National Instruments (NI) to properly synchronize all the sequences. Here we describe that. For the full guide (not specific to our setup), see [12].

We are using a NI PXIe-1073 chassis and a NI PXIe-6535 10 MHz digital IO card to create a clock signal. We also use a NI CB/SCB-2162 single ended digital IO accessory to interface with the card. Below are the steps to configure Atticus.

- Install NI MAX (and MAX DAQ drivers if needed) and open to detect the hardware.
- Rename the card to Dev1 (this is the name Atticus is expecting).
- Click the "configure" tab and set all to line state. Note that this card can only go up to 3.3V (this is fine because the clock isolators will step it up to 5V).
- Install and open Atticus (comes with Cicero).
- Refresh hardware and choose Dev1.
- In the "Hardware Settings" section:
  - Under "Global": set DeviceEnabled to True
  - Under "Digital": set DigitalChannelsEnabled to True.
  - Under "Timing": set ClockEdge (when the FPGAs make changes) to Rising, set MasterTimebaseSource to OnBoardClock, set MySampleClockSource to DerivedFromMaster, set SampleClockRate to 500000 (500 kHz, or whatever frequency corresponds to the time resolution you are using in Cicero), and set UsingVariabletimebase to False.

  - Under "Triggering": set SoftTriggerLast to True and set StartTriggerType to SoftwatreTrigger.

- In the "Server Settings" section:
  - Set the server name to Jlab.
  - Under "Hardward Timing - Old Style": set VariableTimebaseMasterFrequency to 500000 (the same as the sample clock rate) and VariableTimebaseOutputChannel to /Dev1/port0/line0 (or whichever line and port you are taking the output from, any can work).
  - Under "Hardware", click the ellipses next to ExcludedChannels and add the channel you specified for VariableTimebaseOutputChannel in the previous step. This prevents it from being used for anything other than the clock signal.
  - Under "Hardware", click the ellipses next to Connections and set SourceTerminal to /Dev1/DO/SampleClock and DestinationTerminal to /Dev1/PXI_Trig7. Note that you can see the other available device routes for this in NI MAX.
  - All true/false settings under Network Clock and Opal Kelly FPGA should be set to False.
- Click the Connect button to start the server. Leave this window open and connected while using it with Cicero.
- Use jumper wires and BNC cables to take the clock signal from the output channel you specified and run it through separately powered (5V) clock isolators to the DE-0 NANOs. The clock signal goes to GPIO 0 Pin 13 and the ground to Pin 12.
- Note: When you are done, be sure to shut down the computer before powering off the chassis. The chassis won't power off while the computer is on, and if you force it to power off anyway, the computer will crash.

3.4 Cicero and Poseidon quick start guide

- Connect the DE10-Nano(s) and the computer to the same router via Ethernet.
- Open Poseidon (Poseidon.exe). Under Output Servers, you will probably see a list of devices and their IP addresses. We need to update this to match our setup. Click "SOCList" and then the ellipses by (Collection).
- Remove all of the existing members.
- Click "Add". Change the IP Address to the one you set for the device(s) you are trying to connect to. Poseidon likes to follow the naming convention of SOCx, where x is an integer. Leave the rest alone.
- Click "OK" once all the boards have been added to the list.

- You can name the server if you wish. Historically, we have called it Percy, out of affection for Percy Jackson.
- You can save these settings by going to File/Save Settings. If you want them to be automatically loaded when you start Poseidon, save them as poseidonSettings.psf. Otherwise, you can just load them through File/Load Settings after opening Poseidon.
- Click "Establish Communication with SOCs". Poseidon goes through the SOC list, trying to establish contact and set up a server with each of them. If successful, the statuses should read "Connected". If it doesn't work the first time or if only one SOC fails, give it a few moments and try again. Power cycling the SOCs can sometimes help as well (just make sure to wait a bit after powering them back on before attempting to connect).
- When a connection is established, the DAC will output -10V on its VO SMA connector. You can check if this is the case using an oscilloscope.
- Click "Connect" on Poseidon and ensure that the log indicates that it was successful. This step establishes Poseidon's server and allows it to be used in other programs, notably Cicero.
- Leave Poseidon running. Closing Poseidon closes the server and terminates the connections to the DE10s, so only close it once you're done with everything.
- Open Cicero (Cicero Atomic Suite Source/Cicero/Cicero/WordGenerator.exe).
- Click "Server Manager". We need to tell Cicero where to find the Poseidon server we just established and the Atticus server from earlier. Click "Servers" and then click the ellipses on the same line. Remove any existing servers.
- Click "Add" to add each server. Set the server address to 127.0.0.1 for both. Poseidon and Atticus are running on your PC and that IP address (known as localhost) signifies that. Set ServerEnabled to True for both, and set ServerPort to 5681 for Poseidon 5678 for Atticus.
- Click "OK" and then click "Connect to enabled servers". Ensure that Cicero has successfully connected to both Poseidon and Atticus and then close the window.
- Now that Cicero has found the Poseidon server, it needs to know what devices are connected to it. Open Channel Manager and click "Add Logical Device".
- Under "Type", select analog (since we are outputting analog voltages). "Hardware Channel" lists the devices that were connected to in Poseidon. Pick one. You can also give the device a name to help distinguish it. Click "OK".
- Add all the boards and then exit the window.
- Right-click in the open space below RS232 (F5) and select "Add New Timestep". A timestep is a specific voltage pattern that we want to output for a specific duration. Click "Disabled" to toggle it to "Enabled".

- Go to the Analog (F3) tab. Select "Create New Group", enter a descriptive name for it, and hit "Rename". This tab is where we can specify the setpoint patterns we want to output. A group allows us to generate different setpoints for each SOC at the same time.
- The connected devices are listed in the bottom left (often called "channels"). Click "Continue" next to a device to specify what output it should generate.
- Every enabled device will have a corresponding graph showing its output. Clicking a graph allows you to set the output via the dropdown menu "Type". The "Equation" type allows you to input any arbitrary equation, including constants. The values listed on the graphs are the actual values that will be outputted over time. You can also set the duration of the sequence you want for each device.
- Note the "Time Resolution" setting on the left side. This specifies the precision of the outputted waveform (i.e. a 1 ms resolution means that a new voltage value is outputted every millisecond). This resolution MUST match the sample clock rate you set earlier in Atticus. The time resolution is calculated as 2/f, where f is the sample clock rate. Why it is 2/f and not 1/f? Unlike in a function generator where one cycle has both a high and low voltage, the NI clock regards the high and low voltages of the square wave as two separate cycles, so a full period is actually two cycles.
- Once you've set patterns for all of the devices, return to the Sequence (F1) tab. Under the "Analog Group" dropdown, select the group you just made. Click "Update" and ensure that the graphs match what you inputted. You can click on the graphs to see the specific values. You can also set the duration of the timestep.
- We're now ready to output to the DAC(s). In the top left, click "Run Iteration 0 (F9)". Watch the readings on the oscilloscope and see if they match what you expect. If they don't, try running through the Cicero and Poseidon setup again, and check that your clock settings are all correct. There are a lot of small details that can be missed in the setup.

Chapter 4

Vacuum system

4.1 Initial design

The initial design for our vacuum chamber, shown in Figure 4.1, is comprised of four main parts: two ovens, a MOT chamber, and a main chamber. The ovens each have a 10 L/s TiTan ion pump, and are expected to reach a pressure on the order of 10^{-9} Torr. In summary from [29], these ion pumps work by using strong magnets and high voltage to generate a cloud of electrons, which then collide with and ionize gas particles. Due to the high voltage, these ions then fly off and collide with the cathode plate at high enough speeds that they become embedded. The 2D MOT chamber has a 150 L/s NEXTorr z100 non-evaporable getter (NEG) pump, described in [23], as well as an ion gauge to read the pressure. The NEG uses a stack of discs made of a TiZrVAl alloy, which will trap any gas particles that collide with it. The attached ion pump functions in the same way as the ones on the ovens. The pressure in this chamber is expected to reach on the order of 10^{-10} Torr. The main chamber also has an ion gauge, along with a 75 L/s TiTan ion pump and a titanium sublimation pump. The titanium sublimation pump has three filaments of thick titanium wire, and when a large current is run through them, it causes them to heat up and give off titanium vapor. The titanium vapor will then settle on the walls of the chamber, and gas particles that collide with this titanium coating will stick to it. The main chamber is expected to reach a pressure on the order of 10^{-11} Torr. The order of magnitude difference between successive chambers is possible due to differential pumping, which will be explained in depth later in this chapter.

The ovens, one for Lithium and one for Cesium, heat their respective metals to produce a vapor. In summary from chapter 2, this vapor is slowed with Zeeman slowers before it reaches the MOT chamber, where as the name suggests, it is trapped in a MOT. This is then compressed and sub-Doppler cooled before being moved to the main chamber via moving molasses, There it is trapped in an optical dipole trap and evaporatively cooled before it is ready for use. While parts of the design have already been done in other labs, our ovens are a new design, done by a previous masters student in the group [9].

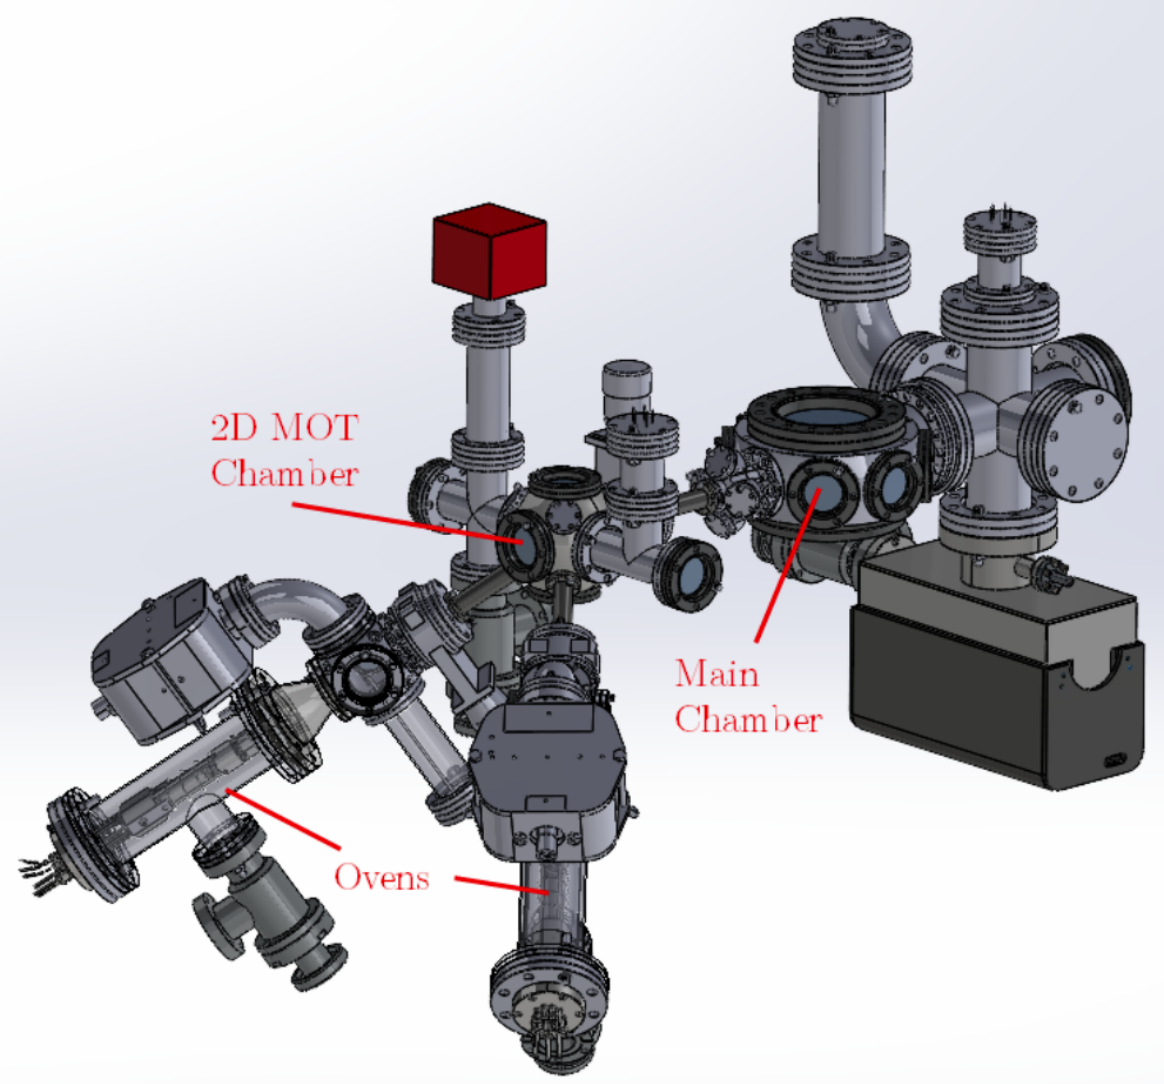


Figure 4.1: CAD drawing of the full vacuum chamber [9].

The initial oven design, shown in Figure 4.2, was a part of previous student Zach Hinkle's masters thesis [9]. The metal sits in a stainless steel pipe crucible (4) with two end caps (1 and 5). The bottom cap (1) can be bolted to an I beam which in turn is bolted to the flange with the thermocouple and power feedthroughs. The top cap (5) has a 0.2" hole in the center to allow the vapor to escape and three smaller holes to bolt the microchannel nozzle (6) into place. The microchannel nozzle is a triangular array of 0.25" long hypodermic needles, and it is designed to allow only the atoms moving in approximately the right direction to pass. This results in a beam that is much better collimated than other alkali oven designs, which normally don't try to control the direction the atoms exiting the crucible are traveling. The metal in the pipe crucible is heated by passing current of 2-3 A through a coil of nichrome wire (2) that is insulated from the crucible itself by four macor

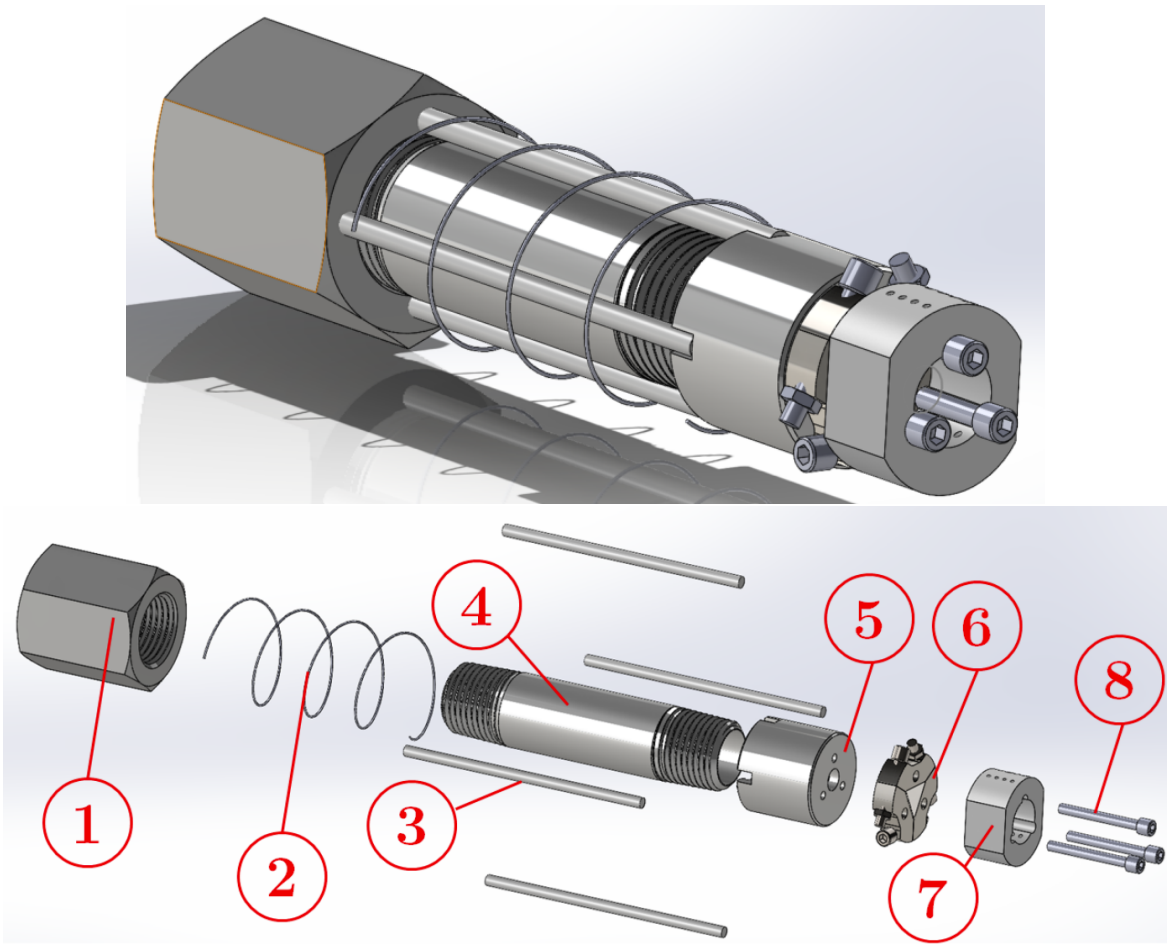


Figure 4.2: CAD drawing of the oven (top) and its components (bottom) [9].

ceramic rods (3). The rods fit into notches in the top cap (5). Similarly, the microchannel nozzle (6) is heated by a coil of nichrome wire looped through a hollow macor ceramic frame (7). Heating the two parts separately allows the nozzle to be kept hotter than the crucible, which is essential to avoid buildup of metal clogging the needles in the nozzle. Two thermocouples (not shown in the diagram) are also used to monitor the temperature of the crucible and nozzle. The crucible temperature needed for a good vapor pressure is approximately 350 C for Lithium and 50 C for Cesium.

4.2 Vacuum tech rules and guidelines

Throughout my time building this system, I have been working to document the general rules of vacuum tech, as well as some specific things about our system we learned through trial and error that are important to know.

4.2.1 Assembly

For any vacuum assembly, there are some basic things that are important to know. The first and most important thing is that we want our vacuum chamber to be as clean as possible. To this end, always wear nitrile gloves when handling any vacuum components, to avoid getting the oils from your skin on them. Oils of any kind are particularly bad to get in a vacuum chamber, as they will outgas slowly over months or years and poison your vacuum the whole time. To get rid of any dirt that may be on the components from manufacturing or previous use, we clean them using a kim wipe and methanol. This is also what you should do if you accidentally touch a component without gloves. Methanol evaporates quickly and is effective at cleaning oils, making it a good choice for cleaning vacuum components. If there is visible dirt on the kim wipe after cleaning, get a new one and repeat until there is no more dirt. For particularly stubborn dirt, isopropanol or acetone can also be used.

There are, however, two special cases to keep in mind for cleaning. The first is components that are extremely dirty, have holes or other places too small to clean by hand, or both. For these, we place them in a beaker of methanol or acetone and sonicate them for 20-30 minutes. Sonication involves placing the beaker in a water bath, then transmitting sound waves through the water to essentially vibrate the dirt off. Many parts of the crucible are cleaned this way, as well as anything I have made or modified in the machine shop as coolant grease would be incredibly bad for our vacuum. The second special case is viewports. In order to avoid damaging the surface of the viewport, they should be cleaned like any other optical component, using the special optics wipes and a hemostat rather than a kim wipe by hand, and the methanol from the glass bottle rather than the plastic bottle. Kim wipes are not designed for optics and can scratch the surface of the viewport, and the methanol in the plastic bottle can leave a residue on the surface that will impede transmission through the viewport.

Almost all parts of our vacuum chamber attach using metal gaskets and CF (ConFlat) flanges. For our system, we use copper gaskets for the Cs oven, MOT, and main chamber, but silver-plated copper for the Li oven. This is because lithium vapor reacts with copper, which will absolutely destroy the gasket as shown in Figure 4.3. The CF flanges have a "knife edge", a thin ridge that cuts into the softer metal of the gasket, providing a very good seal. However, if the knife edge is damaged, the seal will be much weaker, if it can be sealed at all. Damage to a knife edge often results in the component being unusable, and as such it should be avoided at all costs. When not being used, flanges should be covered with a plastic cap or an old gasket and tinfoil to protect the knife edge from being bumped into other components. Parts should never be left resting on an unprotected knife edge or put together without a gasket in between the two. For parts with rotatable flanges, the bolts should be left loose until the flange is in the final position, as trying to turn it with brute force when the bolts are too tight will severely damage the knife edge.

Additionally, if the gasket has been cut into before, cutting into it again will produce a faulty seal. If pressure is removed from a gasket for any reason, it must be replaced with a

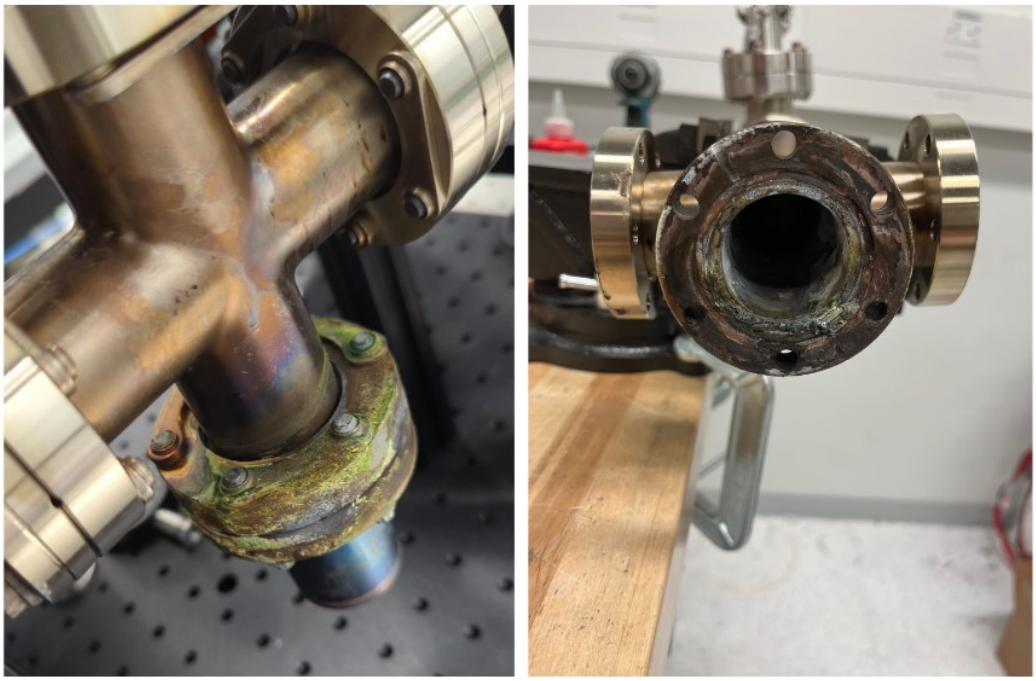


Figure 4.3: The result of accidentally using a copper gasket in a lithium vapor cell.

new one. Even if only one bolt out of six is loosened, the gasket must be replaced to ensure a reliable seal. When assembling, the bolts should all be tightened the same amount, using a torque wrench wherever it will fit as this is more reliable than doing it by hand. The bolts should be tightened incrementally in a star pattern to avoid uneven forces across the flange. For a viewport, the increments of tightening should be much smaller, no more than an eighth of a turn on each bolt, as the glass is much more sensitive to differences in forces across its surface than a regular metal flange.

A common procedure that is used to replace small parts in a vacuum chamber that has already been baked without getting water in it, and in our case, to attach the crucible without the lithium inside reacting with the water in the air, is back-filling the chamber with argon. Argon, as a noble gas, is nonreactive, and by filling the chamber with argon at a pressure slightly above that of the atmosphere, we can largely avoid any air getting into the chamber when it is opened. The procedure for this is to pump the chamber down, fill it with argon, pump it again, and fill it with argon again. This is done to be sure any residual gases are gone. When the chamber is opened, the argon flow should be increased so that the pressure difference between the chamber and the room is maintained, otherwise air will be able to get in.

The crucible uses a lot of non-standard parts, and as such we have developed guidelines for assembling it, sometimes through trial and error.

- All parts that will fit in a beaker should be sonicated, as many have small holes that are impossible to clean manually.

- When putting the crucible together, if the thermocouple leads must be bent for any reason, it is important to use two pairs of pliers. One pair should be used to bend the lead while the other should be used to hold the feedthrough near the ceramic. This is so that the ceramic does not crack from the stresses bending would otherwise put on it.
- Cracked ceramic feedthroughs will leak, but they can be repaired using Torr Seal, a low pressure epoxy. This is not perfect as it can leak again later, especially if the Torr seal is baked above its maximum temperature of 120 C, so it's much better to avoid cracking the ceramic to begin with.
- When attaching the thermocouples and heating wires, make a note of which lead goes to which part of the crucible, as you can't tell from the outside once it is in place.
- As much of the assembly as possible should be done before putting the lithium inside, and the lithium should be added in a glove bag under argon to avoid it oxidizing in the presence of air.
- To minimize oxidization when transferring from the glove bag to the oven, a small amount of hexane can be added to the crucible, as this will create a hexane vapor that keeps the air out. There should only be a few drops of hexane added though, as large amounts of hexane vapor will damage the pump when it is turned on.
- In the end, we decided that this method was impractical for cesium as it reacts much more violently than lithium if exposed to air. The new cesium oven design will be discussed later in this chapter.
- It is also important to check using a multimeter that the heating wires are not shorting inside the machine, since they can sometimes shift during assembly. This can easily be done by measuring the resistance from the feedthrough to the outside of the chamber. If there is any finite resistance measured, there is a short in the system that needs to be fixed. Checking this immediately after installation gives an opportunity to fix it before baking if there is an issue.

After assembly is complete, an important last step is checking for leaks. Depending on the size of the leak, there are a few different methods that can be used to find it. To begin pumping the chamber down, the roughing pump is turned on. A pressure gauge attached between this pump and the turbo pump is an important diagnostic tool for this early stage of pumping. A very large leak can often be heard as a whistling or hissing noise from the leaking flange. These leaks are often found right away, as the pressure gauge mentioned above will often still be reading pressures on the order of 10 Torr in this case. If very large leaks like these are present, they will usually be audible almost as soon as the roughing pump is turned on, and can be easily fixed by tightening bolts on the leaking flange.

If there are no large leaks, the roughing pump should easily be able to get below 30 mTorr as measured by the pressure gauge, at which point the turbo pump can safely be turned on. Since the pressure gauge is on the backing side of the turbo pump, its pressure reading

will be limited by the backing pressure of the turbo pump. This pressure, normally around 5-10 mTorr, is the pressure generated by the turbo pump on the opposite side from the vacuum chamber, and will be there as long as the turbo pump is turned on. Therefore, the value shown by the gauge is limited to this backing pressure, and below that pressure it no longer reflects the actual pressure in the chamber. Once this limit is reached, a pressure measurement from an ion pump or ion gauge is needed in order to be accurate.

If the pressure gets below about 10 Torr but doesn't reach 30 mTorr, there is likely a smaller, silent leak, and these can be found with methanol. Spraying methanol into and around the sniffer ports of the flanges will result in an almost immediate and very noticeable drop and/or spike in pressure if the flange in question is leaking. This method can also be used for even smaller leaks, down to pressures around 10^{-8} Torr. In these cases, the pressure reading from the ion pump or ion gauge will similarly spike or drop suddenly when the leaking flange is found, and can usually be fixed by tightening bolts.

If there is a leak that is too small to find with methanol, the most sensitive test we can do is a helium leak test. The RGA (residual gas analyzer) can detect helium pressures as small as 10^{-11} Torr. Spraying helium around each of the flanges and components will result in a spike of this helium pressure when the hose is near the leaking part. This leak test method can also be used to detect leaks in the glass-metal seal of viewports, or leaks from cracked ceramic feedthroughs. These leaks are less simple to fix, often requiring use of a low pressure epoxy like Torr Seal, or in some cases, necessitating replacing the leaking part completely. If the pressure in the chamber is higher than expected and even the helium leak test shows nothing, there may be a "virtual leak". This type of "leak" is commonly caused by gas trapped in the threads of a valve that works its way out over time. To prevent or fix these, we normally open and close each valve a few times to release any trapped gas.

4.2.2 Baking

In order to fully clean the vacuum chamber after it is assembled, it must be baked. A bake involves heating the chamber and leaving it for a certain amount of time to remove unwanted gases. There are two types of bakes: high and low temperature. The first bake is a high temperature bake, and this is done to remove hydrogen from the stainless steel. Some hydrogen remains in the steel components from the forging process, and if not baked out, it will outgas slowly, poisoning the vacuum. This bake is done at temperatures of 350-400 C, which is much too hot for many of the more sensitive components, such as viewports, gate valves, and ion pump magnets. Therefore, the viewports and any other sensitive components are replaced with blanks and the magnets removed for this bake. Since its purpose is to remove hydrogen from the large steel pipes and chambers, the amount of hydrogen coming from these few components that are removed for the bake is small. The low temperature bake is done once the chamber is fully assembled, and its main purpose is to bake off the water that has settled inside the chamber from exposure to atmosphere. The normal temperature for this bake is 180 C, as the maximum temperature

of the viewports is 200 C and we want to leave a cushion to allow for temperature variations overnight.

For both bakes, there are a few things we found that are important to consider.

- If the heating tapes touch or overlap, there will be localized overheating, so it is important to avoid this.
- The thermocouples should not come in contact with the heating tapes, since when the current is increased, the temperature of the heating tape will increase much faster than the temperature of the steel. If the thermocouple is touching the tape, it will measure that temperature rather than the temperature of the chamber. The characteristic sign of this on the temperature plot is a temperature that increases quickly for a short period of time after the variac is turned up before levelling off.
- If the line on the temperature vs time plot looks "jumpy" rather than smooth, or if the temperature changes abruptly when the outside of the foil is tapped, the thermocouple in question has most likely detached from the metal. There is no way to fix this while the bake is ongoing, but if the thermocouple is detached, the actual temperature of the metal is likely notably higher than the value measured and this should be taken into account when adjusting the temperature.
- Gloves should be worn when putting the heating tapes and thermocouples on to minimize fiberglass splinters from their coatings.
- It is always important to check the maximum heating rate of all components in the bake, and to not exceed the maximum allowed rate of the most sensitive one. For example, the maximum heating rate of the viewports is 3 C/min, so we try to heat and cool no faster than 2 C/min in a low temperature bake to be safe.

There are a few things to know about the high temperature bake specifically.

- Any thermocouples used in high temperature bakes will have damaged coatings at the end, making them prone to letting out a puff of fiberglass dust when moved. Wearing a mask is recommended when handling these to avoid breathing in this dust.
- The temperatures needed for this bake are high enough that the kapton tape used to hold heating tapes and thermocouples in place will break down. Make sure the thermocouples and heating tapes will stay in place without this tape, as after the bake it will have crumbled into a brittle crust.
- After the high temperature bake, the bolts should be replaced, especially any that have turned black. The blackened bolts are prone to coming loose later and as such are unreliable in the long term. Bolts that are in close proximity to a heating tape are more likely to turn black during a bake.

It is also often the case that gaskets will become oxidized and get stuck to the flanges after a high temperature bake, and there are a few methods that can be used to remove them. If two flanges are stuck together, a small pry bar can be pushed between them and used to

pull them apart. If that doesn't work, as a last resort, you can hit one side with a mallet to try to break them apart. Even after getting them apart, the gasket will often be stuck to one side. If it can be lifted off with a pair of pliers, that is the best solution. It is also possible to get them off with the small pry bar, but it is extremely easy to slip and dent the knife edge with the pry bar, so it should be used with great care and only if pliers can't get the gasket off. Whenever a gasket must be forcefully removed, it is important to check the knife edge for damage afterward. This can be done using a flashlight to look for any changes in the way the light reflects off the knife edge, or by running your finger along the knife edge to feel for any defects. If you check with a finger, just be sure to clean the component again afterward.

For the low temperature bake, it is important to consider the more sensitive components that are now in place.

- The viewports should first be covered with a "hat" made of the thicker UHV aluminum foil to avoid scratching them with the heating tapes, thermocouples, or regular foil.
- The heating tapes should not pass over the face of a viewport in order to avoid heating the glass unevenly. The heating tapes should instead pass around the edges of the viewports.
- The magnets for the ion pumps have a self-baking feature, and will reach their intended bake temperature if their heater is plugged in. The ion pumps must be outside the foil for this bake becuase if their heaters are turned on while covered in foil, they will overheat and this could badly damage the magnets.
- The low temperature bake should continue until the RGA (residual gas analyzer) detects no more water in the chamber, which usually takes roughly a week.

4.2.3 Operation

When operating the ovens, it is important that the nozzle is kept at a temperature at least 50 C higher than the crucible. If this difference is not maintained, the lithium especially is able to build up in the needles of the nozzle, eventually clogging them and reducing the flux. The lithium is also prone to building up on the small viewport in the MOT chamber that is used for its Zeeman slower beam. To prevent this, we shine UV light on that viewport constantly whenever the lithium oven is hot. This one viewport, unlike the rest of the viewports in the MOT chamber, is attached with a silver plated gasket so that lithium buildup around the edges of the viewport doesn't react with the copper of the gasket.

4.3 Differential pumping calculations

Since we are relying on differential pumping to maintain a difference in pressure between subsequent chambers, it is useful to calculate the expected difference in pressure for the

length and size of tubes used. The combined length of the Zeeman slower and the small support nipple is 6.5", and they both have an inner diameter of 0.75". Pressures in all chambers are low enough that we are in the molecular flow regime. The conductance through a long round tube is therefore given in chapter 3 of [19] as

$$C_{tube} = \frac{\pi}{12} v \frac{d^3}{l} \tag{4.1}$$

Where d and l are the diameter and length of the tube and v is the average velocity of the gas,

$$v = \sqrt{\frac{8kT}{\pi m}} \tag{4.2}$$

For molecular hydrogen at room temperature, this velocity is 1754 m/s, which gives a conductance of approximately 19.1 L/s. For Argon at room temperature, the velocity is 394 m/s, which gives a conductance of 0.43 L/s. Now, we can use the throughput Q to relate this to the pumping speed S. Throughput is defined in chapter 3 of [19] as

$$Q_{tube} = C(P_1 - P_2) \tag{4.3}$$

$$Q_{pump} = SP \tag{4.4}$$

Where P_1 is the higher pressure (the oven in our case) and P_2 the lower pressure (the MOT chamber). These pressures are constant values that we can measure, so the question of interest is, given a measured value of P_1, what do we expect P_2 to be? Since we expect P_2 to be constant as well, we can say that the net throughput of the MOT chamber is zero:

$$Q_{tube} - Q_{pump} = 0 \tag{4.5}$$

$$C(P_1 - P_2) - SP_2 = 0 \tag{4.6}$$

Solving for P_2, we get:

$$P_2 = \frac{C}{C+S} P_1 \tag{4.7}$$

In the MOT chamber, we have the NEG pump, which has a pumping speed of approximately 150 L/s for hydrogen. The NEG is much less effective at pumping argon, with a speed of roughly 6 L/s. Plugging these values in, along with the conductances calculated earlier, gives $P_2 = 0.113P_1$ for hydrogen and $P_2 = 0.067P_1$ for argon. Although the NEG

is much less effective at pumping argon than hydrogen, the very low conductance of argon through the tube means that we still expect hydrogen to have a higher pressure in the MOT chamber. Of course, since these calculations do not take into account any outgassing in the MOT chamber and are only for two of the possible gases we would have, these values are not going to perfectly match what we see in the chamber. However, we expect the residual pressure to be predominantly due to hydrogen, so they should be a fairly good estimate of the relative pressure. Therefore, the result that $P_2 = 0.113P_1$ for hydrogen confirms that we should be able to sustain an order of magnitude pressure difference between successive chambers as planned.

4.4 Problems and solutions

4.4.1 The ovens

Because the student who designed our ovens graduated before we had a chance to fully implement them, there were several issues we came across in the process of building the ovens that necessitated changes to the original design.

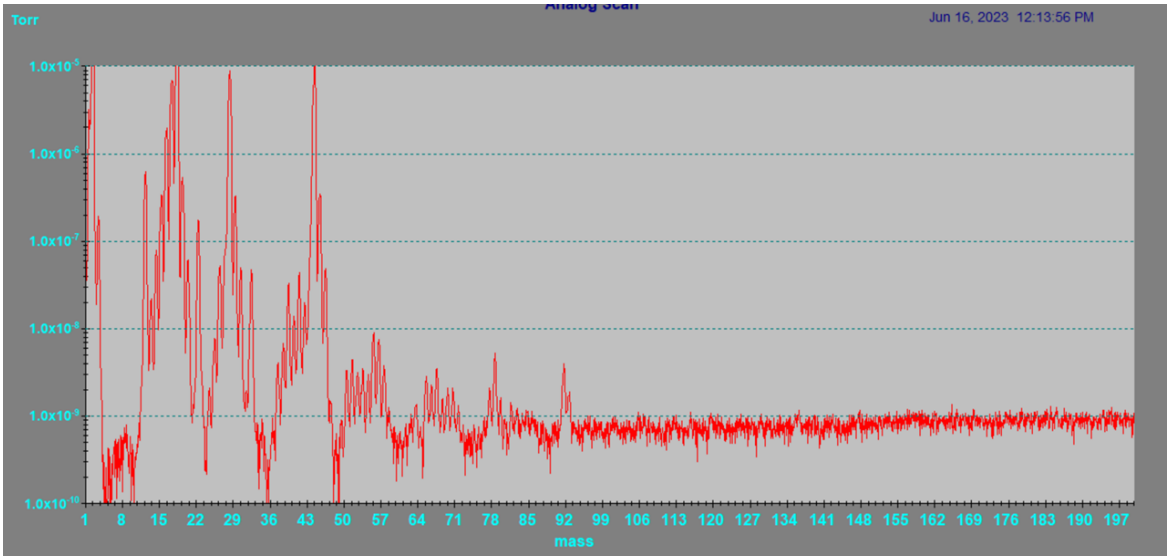


Figure 4.4: Trace from a residual gas analyzer (RGA) showing pressure vs mass (in AU) of gases in the vacuum chamber. This trace was taken with the old Cs oven design at roughly 400C.

Initially, we used nichrome wire for the heating coils, but we were seeing strange high mass peaks on the RGA (residual gas analyzer) traces. The unexpected peaks can be seen in Figure 4.4 between masses of 50 and 100 AU. We found out that, unless the wires are first oxidized in air, they will emit nickel and chromium oxides when heated. Oxidizing them by heating them in air is impractical, so we switched to titanium wires for the heating coils as titanium has a similar resistance but does not emit unwanted junk when heated.

Additionally, we found that the I beam attaching the crucible to the flange was slightly too big, and its corners were shorting to the heating wire feedthroughs. This was easily solved by grinding down the corners with a dremel.

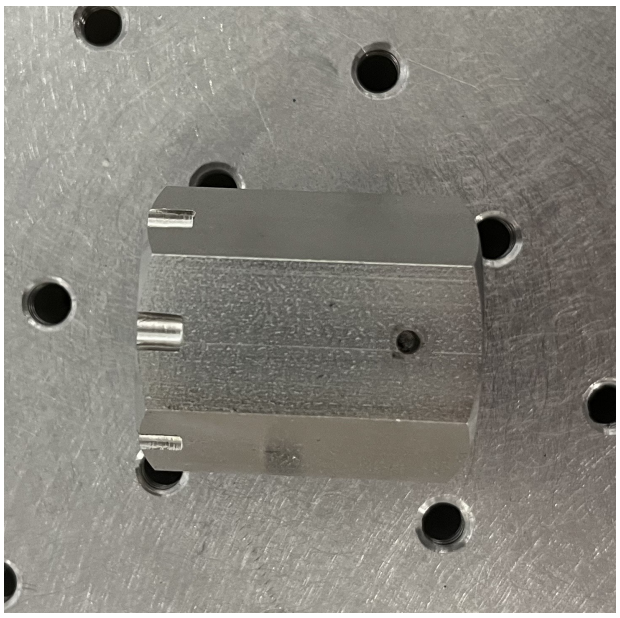


Figure 4.5: The new design for the bottom cap of the crucible, showing the hole for the thermocouple and the notches for the macor rods.

In the initial design, there was a notch in the ceramic nozzle heater for its thermocouple, but the thermocouple for the crucible simply rested on top of the bottom cap. We found that it was very easy for this thermocouple to slip out of place, so I drilled a small hole into the solid base of the bottom cap for the thermocouple to be placed, as shown in Figure 4.5.

The crucible was also designed with notches in the top cap for the insulating macor rods, but this made the top cap difficult to put on without the rods falling out. As a result, we put small pieces of lithium through the hole in the cap rather than adding large chunks and putting the cap on afterwards. Even though we did this in a glove bag under argon, the process took long enough that a thick oxide layer formed on the lithium pieces, giving us much lower vapor pressure than we wanted. To fix this, we redesigned the crucible to have six notches in the bottom cap as shown in Figure 4.5 rather than four in the top cap, so that the top cap can be removed and replaced easily. The reason for six rods rather than four is simple: due to the hexagonal symmetry of the bottom cap, the notches were much easier to drill if we did one in the center of each side.

The original design also had the ovens attached to the 2D MOT chamber directly via the Zeeman slower, which is a regular steel nipple. The issue we ran into here is that every time the ovens were jostled, such as when replacing the crucible at the end, the connection to the 2D MOT would start leaking. The joint was not strong enough to handle the stress

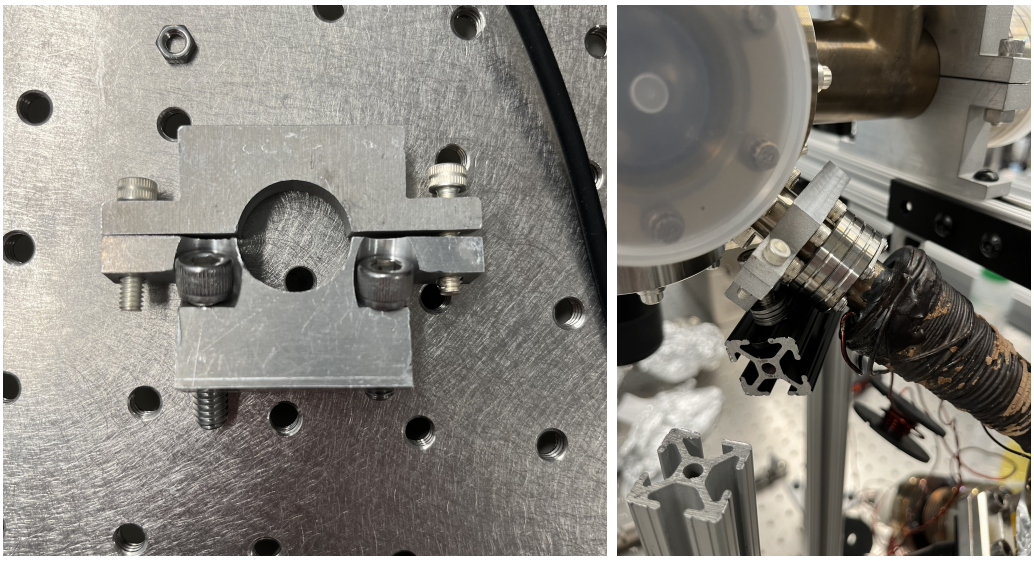


Figure 4.6: The clamp designed to hold the weak joint stable temporarily. Shown in place on the joint in the right side image.

of even small movements. Tightening the bolts would stop the leak, but it would just leak again the next time anything was bumped. Eventually, the bolts could not be tightened any more without risking stripping them, so another solution was needed. In order to take the strain off of that joint, we decided to add small bellows between the slowers and the 2D MOT chamber. The added flexibility will enable small movements of the oven without putting all the stress on the one joint, which will stop the repeated leaks. However, since these bellows need to be custom ordered and will take months to arrive, I designed a clamp to be used to prevent movement of that joint as an interim measure. While not perfect, it should prevent most small movements in the oven from putting stress on the joint between the slower and the 2D MOT.

Initially, we planned to use a reaction to produce cesium rather than using pure cesium because it would be much easier and safer to assemble that way. The reaction used was:

$$2CsCl + Ca \rightarrow CaCl_2 + 2Cs \tag{4.8}$$

However, this reaction proved to be very unreliable, starting at wildly varying temperatures or sometimes just not reacting at all. This was determined using spectroscopy: when the reaction started, cesium was being produced, and we could see the vapor this way. In the final attempt, we were using the RGA as well, and if Cs was being produced we would expect to see a peak in Figure 4.4 at a mass of 133 AU. As seen in the figure, there is nothing above the noise at 133. This trace was taken with the crucible at a temperature upwards of 400 C, which is far above the temperature we would expect the reaction to start. After considering all other explanations for the lack of cesium production, we decided that this reaction was too unreliable to be used to produce our cesium, and we decided to

switch to pure cesium. However, the problem of how to safely and effectively load the cesium into the crucible had no good solution, so we ended up redesigning the cesium oven entirely. Instead of the complicated crucible setup previously discussed, we would simply put a glass ampule of cesium in a bellows. Once the system was under vacuum, the ampule could be broken, so there is no risk of the cesium being exposed to atmosphere. In order to maintain some of the benefits of the old design, we kept the microchannel nozzle, mounting it onto a double sided flange at the end of the bellows, with holes drilled out in the same configuration as the old top cap, shown in Figure 4.7. Since cesium only needs to be heated to about 50 C for a good vapor pressure, a less complex heating system consisting of a simple external heater on the bellows will suffice. Spectroscopy after the final bake of the system measured a more than adequate amount of cesium vapor being produced.

We also found that the stainless steel crucible design, still in use in the lithium oven, produced more hydrogen than expected when heated, and the ion pump is not very effective at pumping hydrogen. In order to combat this, we decided to add a second titanium sublimation pump to the lithium oven.

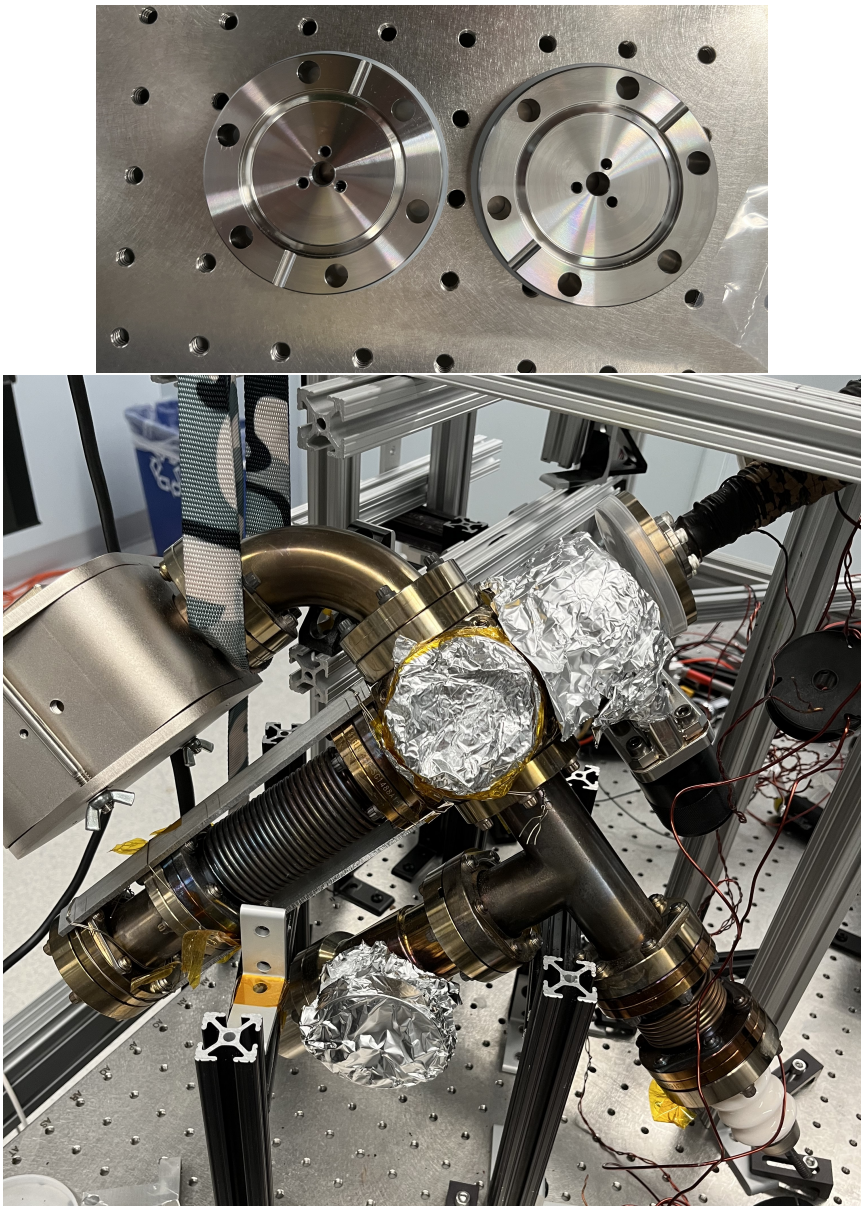


Figure 4.7: Top: double sided blank flanges with holes drilled to mount the microchannel nozzle. Bottom: The new design for the Cs oven.

4.4.2 The NEG

There was one final issue we ran into, but it was in the 2D MOT design rather than the ovens. The NEG (non evaporable getter) pump was performing at less than a tenth of its advertised pumping speed despite multiple attempts to activate it. In addition, the ion pump attached to the NEG was reading pressure an order of magnitude different from that shown by the ion gauge in the same chamber. The issue turned out to be the size of the nipple the NEG was in. The design of the NEG includes a stack of discs of the getter material inside a hexagonal casing, and the diameter of the casing was only slightly smaller than the inner diameter of the nipple. This caused the conductance around the NEG to be very small, and therefore it couldn't pump effectively. The inner diameter of the pipe is 36.1 mm and the hexagon's corner-to-corner diameter, specified in [23], is 35 mm.

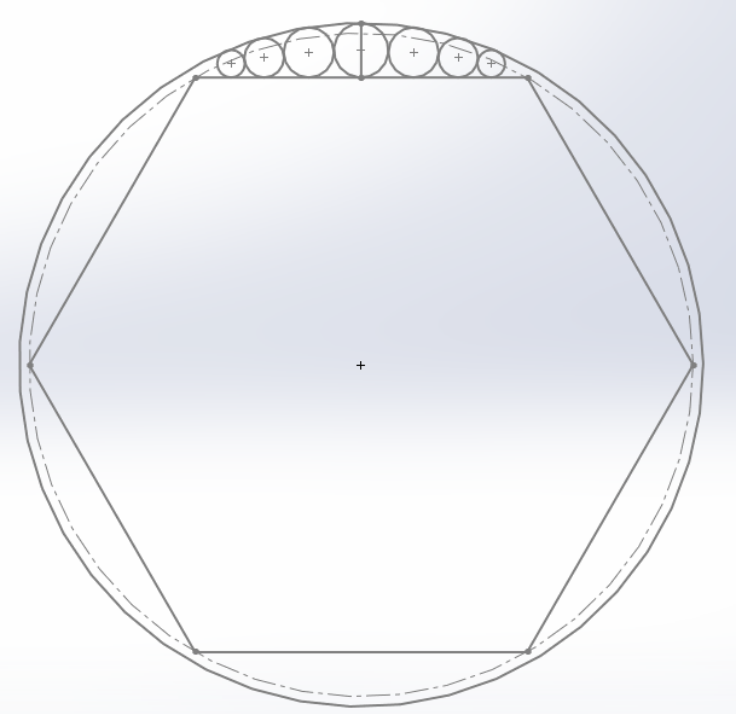


Figure 4.8: Model of the conductance around the hexagonal casing of the NEG pump.

To calculate the conductance past this hexagon, I modelled the gap between the hexagon and the pipe as a set of circular tubes, shown for one side in Figure 4.8. The largest tube in the center has a diameter of 1.45 mm, and the pairs of tubes on either side get successively smaller, with diameters of 1.32 mm, 1.04 mm, and 0.74 mm respectively. From [19], the conductance of a long tube is

$$C_{tube} = \frac{\pi}{12} v \frac{d^3}{l} \tag{4.9}$$

For tubes in parallel, their individual conductances can just be added together for a total

conductance. Therefore, the total conductance past the hexagonal casing can be expressed as

$$C_{total} = \frac{\pi}{12}\frac{v}{l}(6d_1^3 + 12d_2^3 + 12d_3^3 + 12d_4^3) \tag{4.10}$$

where d_1 is the diameter of the largest central tube, and $d_2 - d_4$ are the diameters of the smaller tubes on either side. To get an estimate of the conductance to most of the getter disks, we use a tube length of 1". For hydrogen, with an average velocity of 1754 m/s, we get a conductance of approximately 1.16 L/s, which explains the very low pumping speed. While this conductance is an underestimate as the approximation does not take into account the entire area between the pipe and the hexagon, it does show that the true conductance is nowhere near the 150 L/s the pump would otherwise be capable of. We fixed this simply by putting it into a larger nipple to enable better conductance, as shown in Figure 4.9.

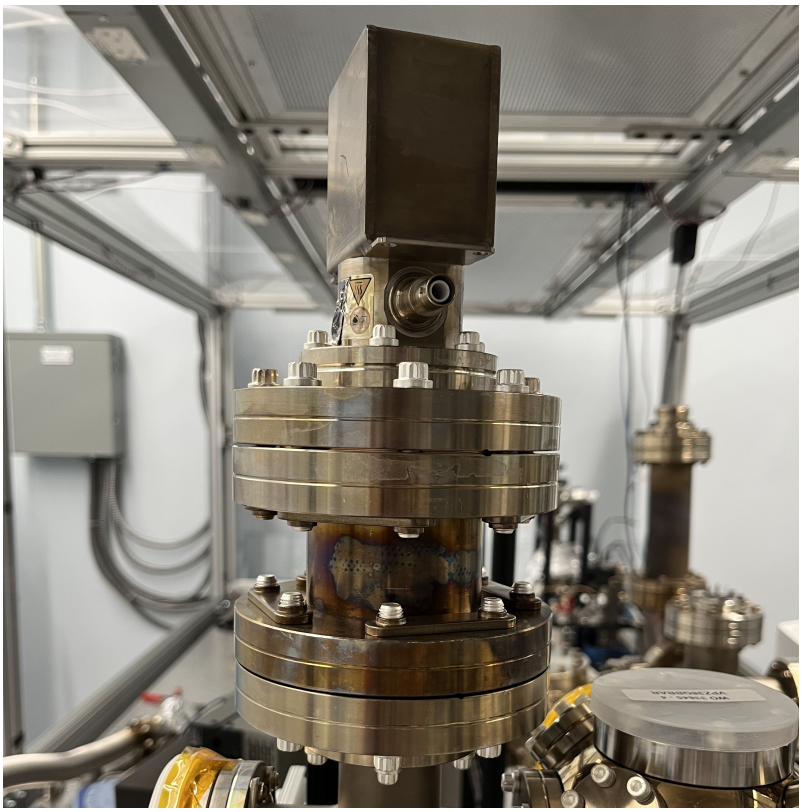


Figure 4.9: Non-evaporable getter (NEG) pump in a new, larger, nipple.

This larger nipple did improve the pressure in the MOT chamber, resulting in a decrease of roughly an order of magnitude compared to the previous pressures measured with similar oven pressures. However, due to an unresolved leak in the MOT chamber, it is not possible as of the writing of this thesis to accurately quantify the improvement.

4.5 Summary and future steps

In conclusion, a lot of progress has been made in the documentation and design of our vacuum system. With the experience of building and troubleshooting our vacuum system, I have written a thorough guide on vacuum technology that will be useful for future students. It covers vacuum guidelines, rules, and procedures, both for general vacuum systems and for specific things unique to our system. I have also made a number of important improvements to the initial design of the system. The changes in the ovens improved ease of assembly, cleanliness of the system, and long term stability. The clamps on the 2D MOT chamber succeeded as an interim leak prevention measure for the Zeeman slower joints, and the NEG is performing much better now that it is in a larger nipple to improve conductance.

The next step for this system will be to capture a cloud of atoms in the MOT to quantify the performance of the ovens. After that, the main chamber will be attached and baked and the surrounding optics put into place, after which the chamber will be ready to run experiments. This chamber is intended to run experiments on ultracold mixtures rather than molecules, so I will use everything I have learned from this assembly to build a second very similar system for the molecules experiments. Since the prototyping and trial and error have already been done for the first system, the second should be assembled much more quickly and smoothly.

Chapter 5

Dipole blockade CNOT fidelity calculations

5.1 Introduction

The ideal fidelity found in [9] and discussed in chapter 2 does not take into account many factors that affect real molecules. It assumes, among other things, no interaction between atoms in storage states and no interaction of the storage states with external electric fields, which is not the case for real molecules.

Applying a large electric field to polarize the ground rovibrational state is one method considered to perform the dipole blockade. Since it has a large polarizability, this ground state acts as the dipole state $|d\rangle$ in the dipole blockade. However, the other vibrational states also have some polarizability, and although it is not as strong as the ground state, they will always be polarized to an extent as well. The effect of this can be minimized by choosing the states with the smallest polarizabilities, which will be the two highest energy vibrational states due to their nuclei being on average the furthest apart. Different states will also have different electric dipole moments, which leads to different Stark shifts in the presence of an electric field. Similarly to the polarizabilities, this effect is expected to be smaller for higher vibrational states.

In addition, these differing energy shifts will cause the lasers driving the $|0\rangle - |d\rangle$ and $|1\rangle - |d\rangle$ transitions to be detuned from their respective transitions by different amounts. One of the key assumptions made in the ideal fidelity calculations in [9] is that these detunings would be the same. To quantify the effect these differences will have on the coherence and gate times and therefore on the fidelity, we need to identify the storage states to be used and calculate their respective polarizabilities and Stark shifts. Using these results, we can determine how the states will respond to stray electric fields and noise in the polarizing electric field and the effect this will have on coherence time.

In order to do these calculations, we begin by calculating the wavefunctions and energies for each vibrational state within the ground electronic state. The ground electronic state

is chosen because the ground rovibrational state within this electronic state has the largest dipole moment of any state of the molecule, and it will therefore be the most effective choice for a dipole blockade using this method. Given the potential curve for this state, the wavefunctions are calculated using the Numerov method to solve the Schrodinger equation for each state, and the energies of the states are calculated from these wavefunctions. Given data for the polarizability and electric dipole moment of the molecule with respect to internuclear distance, we can use the wavefunction to get an overall value for the polarizability and dipole moment of each vibrational state averaged over internuclear distance. The details of these calculations are explained in depth in the next section.

These calculations will tell us how many bound vibrational states exist within the ground electronic state, what their energies and wavefunctions are, and give distance-averaged values for their polarizabilities and Stark shifts. Using this information, we can choose which two states we will use as storage states and calculate the difference in overall energy shift due to external fields between the two. We can also quantify the effect stray electric fields will have on these states and determine how stable the main polarizing electric field needs to be to achieve the desired coherence times. The results of these calculations are discussed later in this chapter.

5.2 Vibrational state energies for $X^1\Sigma$ LiCs

For this calculation, I inherited some code from Alan Jamison that would take a set of potential datapoints, create a potential curve, and find the wavefunctions and vibrational levels for NaLi. I updated it to run on Python 3 since it was written for Python 2, modified it to work for LiCs and calculate some hardcoded parameters dynamically, and added functions to calculate polarizabilities and energy shifts. This will be explained in more detail in the next section.

The first step in this process is to find the wavefunctions and corresponding energies of each vibrational level of the $X^1\Sigma^+$ state. Using datapoints from [15], the potential is splined and extended far beyond where we expect any bound states, shown in the bottom panel of Figure 5.1, in order to ensure they are all found. The potential curve for this state was calculated in [15] as a sum of several contributions. The coupled cluster method with corrections for electron correlations is used at small and intermediate internuclear distances, but this method becomes inaccurate at large distances. For large distances, the result of this method is merged with an expansion of the potential approximated by

$$V_{int}(R) \approx -\frac{C_6}{R^6} - \frac{C_8}{R^8} - \frac{C_{10}}{R^{10}} \tag{5.1}$$

where the C coefficients are long range dispersion coefficients found in [17]. This same long range approximation is used in our calculations to expand the potential given in [15] even further, in order to ensure all bound states are found. This part of the code required no modification aside from changing the potential data and coefficients to be those of LiCs.

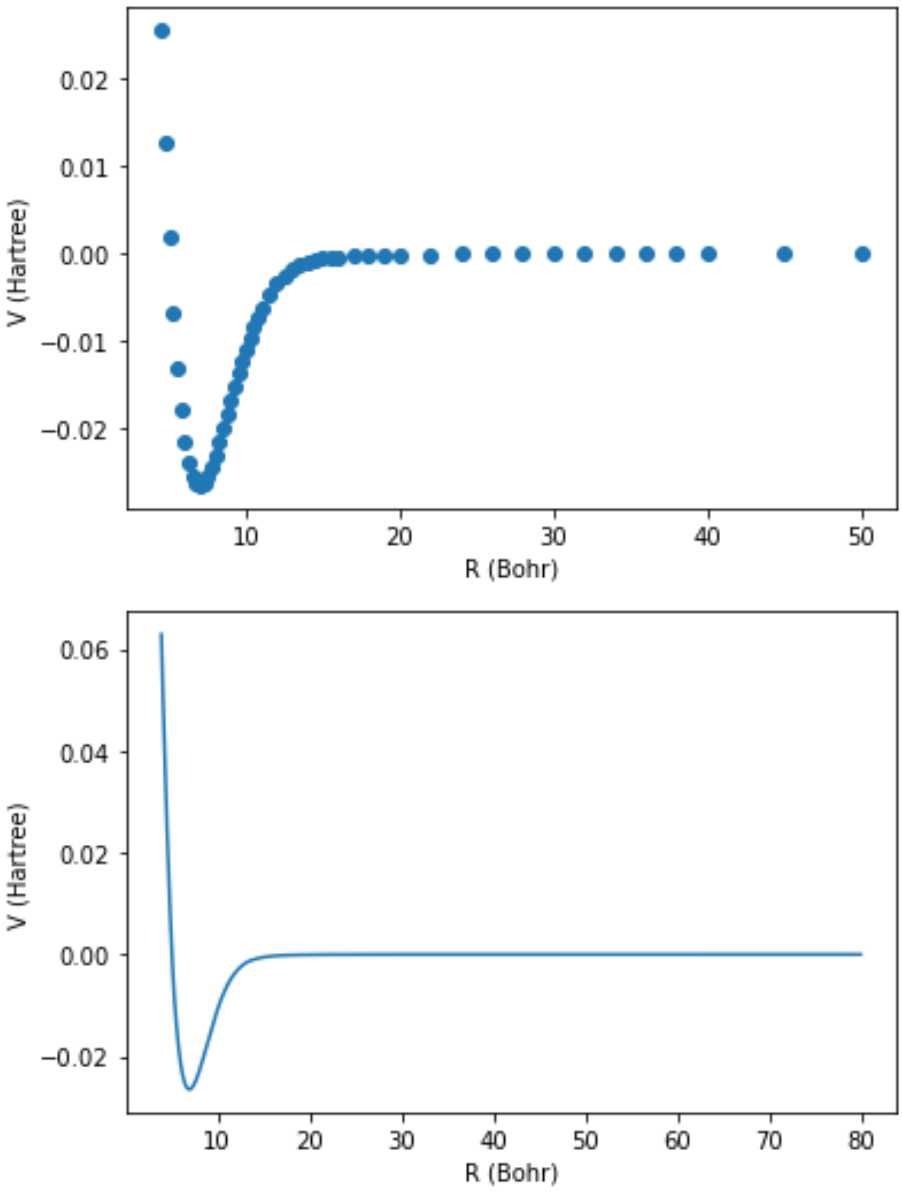


Figure 5.1: Potential data (left) and fit (right) for the $X^1\Sigma^+$ state of LiCs.

The first major change came in the calculation of the wavefunctions and energies. As a part of the calculation, the code finds the outer classical turning point and integrates inward from the end of the potential to match the wavefunction value at the classical turning point. The problem is that at this stage in the calculation, the wavefunction has not yet been normalized and is very large. Integrating inwards from as far as 80 Bohr resulted in values so large that they overflowed Python's float maximum, resulting in infinities. There is no need for the integration to be done from that far out, it was simply hardcoded that way for convenience in the original version, but to fix this problem, I wrote a function to calculate

how far beyond the classical turning point the integration should start dynamically for each level based on the energy of the previous level. This distance was defined to be $x_n = 10/\kappa_{n-1}$, where κ is the decay coefficient of the wavefunction beyond the classical turning point, found for a molecule with reduced mass μ as follows:

$$\kappa_n = \sqrt{-2\mu E_n} \tag{5.2}$$

For $n > 1$, this is easily found using the energy of the previous state, but for the first step, the ground state must be found manually. This is done by fitting the bottom of the potential curve to a fourth order polynomial, starting with 100 points either side of the minimum and decreasing this range until a good fit is found. Since we expect the bottom of the potential to be roughly parabolic, a fit is considered good when the second order term is at least 100 times larger than the fourth order term at both endpoints of the fit range. The energy of the ground state is then calculated as

$$E_1 = -D + \frac{1}{2}\omega \tag{5.3}$$

where D is the potential depth and, for a fit polynomial $ax^2 + bx + c$,

$$\omega = \sqrt{2a\mu} \tag{5.4}$$

This energy is then used to calculate the distance to integrate beyond the classical turning point for the lowest energy level, with the distance for subsequent levels determined by the energy of the previous level. Now, the wavefunctions and energies (shown in Figure 5.2) can be calculated. This method is not specific to the LiCs $X^1\Sigma^+$ state, it can be used to determine dynamic integration limits for any potential with a roughly parabolic shape around the minimum.

5.3 DC electric field method

In this method, a large DC electric field is applied to polarize the absolute ground state, which is then used as the dipole state in the dipole blockade. As calculated in the previous section, this state has binding energy -1.714x10^5 GHz. For the best fidelity, storage states with weak interactions with neighboring molecules are needed, so we will likely use the two highest lying vibrational levels, 48 and 49, for storage states. As found in the previous section, their binding energies are -31.47 GHz and -2.96 GHz respectively. Another aspect of the storage states to consider for coherence time is the shift in their energies resulting from electric fields. In an ideal scenario, they would shift by the same amount, but that is not the case. In order to determine the difference in energy shift, we need to find the polarizabilities and Stark shifts for the vibrational states.

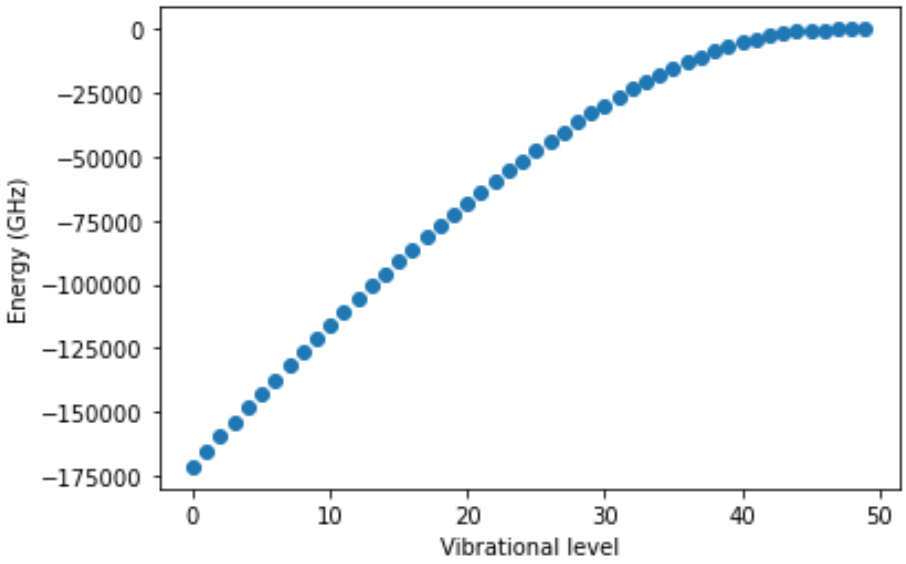


Figure 5.2: Energies for the bound vibrational states in the $X^1\Sigma^+$ state of LiCs.

5.3.1 Polarizability and Stark shifts

An important aspect to consider when selecting storage states is the polarizabilities of these states. In the fidelity calculations summarized in chapter 2 [9], one of the key assumptions is that the detuning of the laser is the same for both storage states. However, this will not be the case for real states, as they will have different energy shifts for a given electric field, and some level of stray fields or field fluctuations are unavoidable. The goal is then to select two storage states that have the smallest difference in energy shift.

With the wavefunctions and energies determined, we now need to find the polarizabilities. The functions to find these were added to the code from the previous section, as it already calculates the energies and wavefunctions of each state, which are needed for this calculation. Molecules in an external electric field will have an induced dipole moment

$$\mu_{ind}^{el} = -\alpha E \tag{5.5}$$

determined by the polarizability, α [5]. The resulting energy shift will be

$$\Delta U = \alpha |E|^2 \tag{5.6}$$

U is used for energy here to avoid confusion with the electric field E. The polarizability α is a function of internuclear distance, and as such needs to be spatially averaged for each vibrational state. This is calculated as

$$\alpha_{tot} = \int \Psi^* \alpha(R) \Psi dR \tag{5.7}$$

where Ψ is the wavefunction of the state.

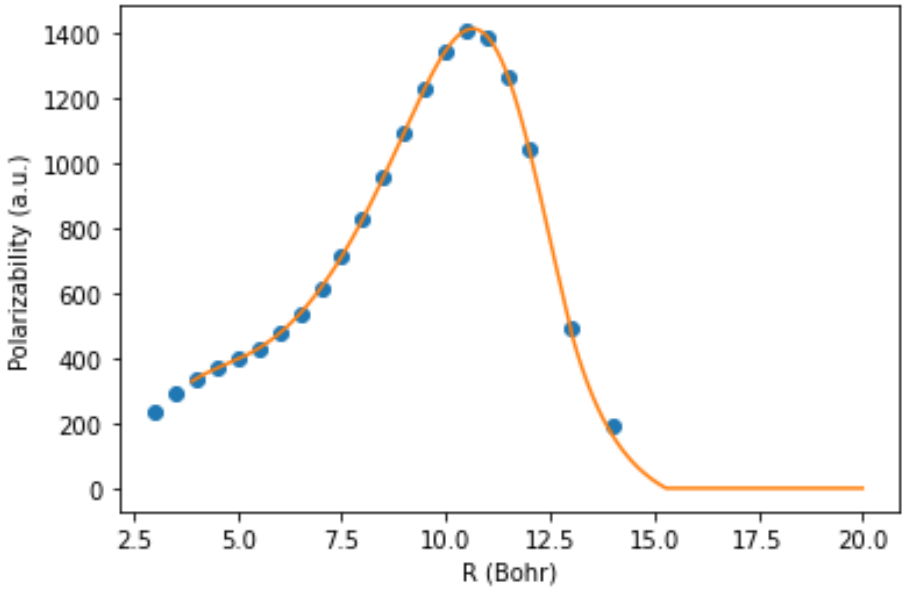


Figure 5.3: Polarizability data (blue dots) and fit (orange line) for the $X^1\Sigma^+$ state of LiCs.

The polarizability data from [15], shown as blue dots in Figure 5.3, is fitted to a curve using a simple spline in the short range, but in the long range ($R > 12$ Bohr), the data is instead fitted to an exponential decay for greater accuracy. Since negative polarizability values are not physically accurate, any negative values from the fit are set to zero. Using this curve and the wavefunction previously calculated, the polarizabilities for each vibrational level are calculated and shown in the upper panel of Figure 5.5.

Molecules in electric fields will also experience a Stark shift to their energies, explained in chapter 10 of [5]. LiCs here, line all linear molecules in a Σ state, has no first order Stark shift due to its symmetry. However, it does have a second order Stark shift due to its orientation that must be taken into account. This second order shift is calculated in [5] as

$$\Delta U_{JM}^{(2)} = \frac{(\mu^{el})^2 E^2}{2hcB_v} \frac{J(J+1) - 3M^2}{J(J+1)(2J-1)(2J+3)} \tag{5.8}$$

Since we are using the ground rotational state, the rotational quantum number J and its projection M are both equal to zero, so this simplifies to

$$\Delta U_{JM}^{(2)} = \frac{(\mu^{el})^2 E^2}{2hcB_v} \frac{-1}{3} \tag{5.9}$$

The rotational constant B_v is defined in [5] as

$$B_v = \frac{h}{4\pi\mu c}\int \Psi^* \frac{1}{R^2}\Psi dR \tag{5.10}$$

μ in this constant is the reduced mass of the molecule rather than the electric dipole moment. The electric dipole moment in the energy shift is found in a similar method as the polarizability. Data points for the electric dipole moment vs internuclear separation [15], shown as blue dots in Figure 5.4, are splined and fitted to create a curve. The long range fit is also to an exponential, but in this case, the data must also be fitted to extrapolate at short range. The data for short range is fitted to a second order polynomial, and this fit is used to extend it to 4 Bohr. The complete curve is shown as an orange line in Figure 5.4.

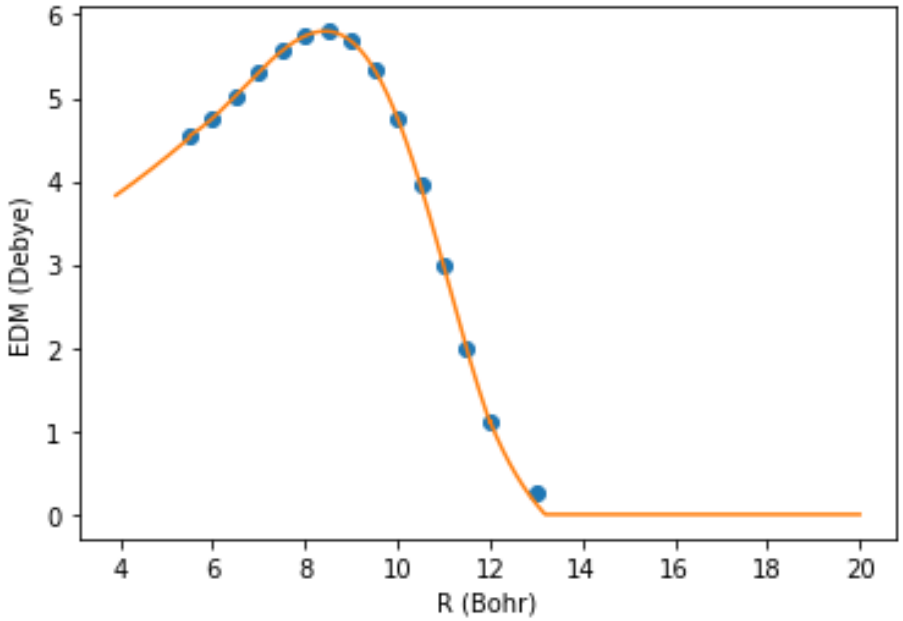


Figure 5.4: Electric dipole moment data (blue dots) and fit (orange line) for the $X^1\Sigma^+$ state of LiCs.

Using the wavefunctions calculated for each vibrational level and converting to atomic units, the second order Stark shift can then be calculated as

$$\Delta U = \frac{-\mu E^2}{3}\frac{(\int \Psi^* \mu^{el}\Psi dR)^2}{\int \Psi^* \frac{1}{R^2}\Psi dR} \tag{5.11}$$

This shift (divided by E^2) is shown in the lower panel of Figure 5.5. As shown in the figure, it is three orders of magnitude larger than the shift due to the polarizability, and as such this is the shift we should focus on minimizing.

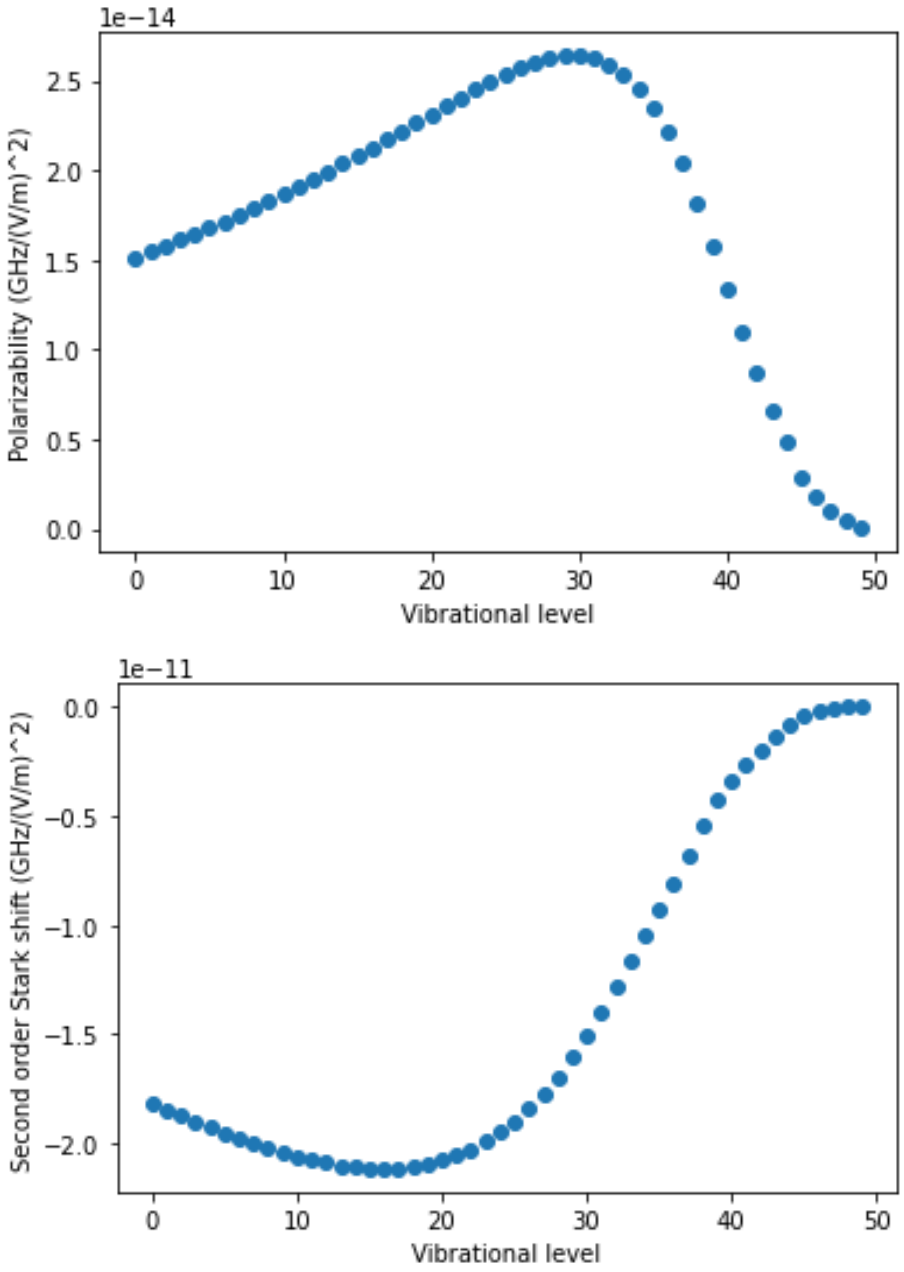


Figure 5.5: Polarizabilities (top) and second order Stark shifts (bottom) for the $X^1\Sigma^+$ state of LiCs.

Aside from producing the plots shown in this section, the final addition I made to the code was a function to find the two states with the smallest difference in Stark shift, checking not just adjacent states but all combinations of states. For the DC electric field we will likely use the two highest vibrational states regardless of this result as their interaction strength will be the weakest, but for the circularly polarized microwave method we have

many more choices. The two highest vibrational levels, 48 and 49, have a difference of $1.9144x10^{-14}GHz/(\frac{V}{m})^2$, but the two states closest in shift are actually states 3 and 25, with a difference of only $9.5857x10^{-16}GHz/(\frac{V}{m})^2$, which is roughly an order of magnitude smaller. While the exact shifts will be different for the circularly polarized microwave method as the calculation is different, it is worth keeping these two states in mind as storage state candidates for the future. This method is briefly explained in the next section.

5.3.2 Stray electric fields and noise

For a stray electric field of 50 V/m, which is easy to achieve in the lab, the energy shifts of the two highest vibrational states would be -0.526 Hz for state 48 and -0.047 Hz for state 49, resulting in a difference in second order Stark shift of 0.479 Hz. More focus on removing stray fields in the lab could result in fields as small as 1 V/m, for which the energy shifts would be -2.102x10^{-4} Hz for state 48 and -1.882x10^{-5} Hz for state 49. The difference in Stark shift would then be -1.914x10^{-4} Hz for these two states.

We also need to consider the noise present in the polarizing DC electric field. If we express this electric field as

$$E = E_0 + \varepsilon \tag{5.12}$$

the shift in energy from the Stark effect will be

$$\Delta U = \alpha(E_0 + \varepsilon)^2 = \alpha(E_0^2 + 2E_0\varepsilon + \varepsilon^2) \tag{5.13}$$

where α is the difference between the state dependent constants found for the storage states in the previous section. The magnitude of the electric field E_0 would be around 10^5 V/m, and for a coherence time of 1 second, we want the second term to be less than 1 Hz, i.e.

$$\alpha\varepsilon \leq \frac{1x10^{-9}}{2E_0} \tag{5.14}$$

Since E_0 is very large compared to ε, we neglect the third term, $\alpha\varepsilon^2$. For storage states 48 and 49, the difference in Stark shift was found in the previous section to be 1.9144x$10^{-14}GHz/(\frac{V}{m})^2$. Plugging this in shows that, for a coherence time of 1s, our polarizing field must be stable to within 0.26 V/m, which would be very difficult to achieve in the lab. A longer term goal would be a coherence time of 10 seconds, for which we would need an electric field stable to within 0.026 V/m. This would be nearly impossible. Therefore, if we want coherence times as large as 10 seconds, it will be necessary to consider other methods such as the circularly polarized microwave method. That method uses circularly polarized microwaves rather than a large electric field to polarize the ground state by causing the molecules to rotate in sync. While this results in a lower dipole-dipole interaction energy, it has the

advantage that there is no polarization of states other than the ground state. This method is not analyzed in detail here, but it will be considered in the future.

5.4 Summary and future steps

In conclusion, some of the real world limiting factors for the DC electric field dipole blockade method have been investigated for $X^1\Sigma$ LiCs. I inherited code to calculate wavefunctions and energies for the vibrational levels of NaCs, and modified it to work with LiCs on Python 3. I added a function to dynamically determine the integration limit beyond the classical turning point to avoid overflows in the wavefunction calculation. I also added functions to take in and fit polarizability and electric dipole moment data and use this data along with the wavefunctions to calculate the polarizabilities and second order Stark shifts for each vibrational state and find the two states with the smallest difference in shift. Using these values, I then calculated what the difference in energy shift would be for our potential storage states from various stray electric fields, and determined how stable our large polarizing electric field would need to be to achieve the desired coherence times.

In the future, I will expand this analysis to the triplet state of LiCS as well as several other molecules for comparison. I will also consider the effects of laser noise and intensity limitations on the fidelity. After that, I will repeat the whole process for a different method that uses circularly polarized microwaves to polarize the ground state. The calculations will be different due to the nature of the polarization, but the comparison of noise limitations and shift differences will be useful to determine which of the two methods we will use in the long term. Based on these calculations and the vacuum work done in chapter 4, I will set up an experimental apparatus for the chosen method and attempt to perform a CNOT gate with ultracold LiCs molecules.

Chapter 6

Conclusion

In conclusion, ultracold molecules are a promising platform for quantum computing that will be further investigated in the future, and the work done in this thesis is applicable not only to this goal but to other fields as well. In chapter 3, I discussed a control system using DE10 Nano FPGAs, DACs, and a NI clock that allows remote, synchronized control of electronic components in a lab, such as laser shutters and magnetic field coils. Any lab that uses similar components could make use of a similar control system.

In chapter 4, which discussed vacuum technology, I wrote an extensive guide to vacuum work, documenting the tips and tricks I learned and the mistakes I made. Much of what I know about practical implementation of vacuum technology in the lab was either passed down verbally by previous students or learned the hard way through my own mistakes. There is no written guide to this in our group, and I could not find any written guide anywhere else either. This guide can therefore be very useful to any group using vacuum technology, especially if there are no students with previous experience to pass down the tips and tricks. The guide goes over the basic rules of vacuum systems for both assembly and baking, lists tips and tricks for troubleshooting, and goes over several mistakes that are easily avoidable if you know about them but easy to make if not. It will allow future students to learn from the mistakes I have made rather than learning what works through trial and error as I sometimes did. In addition, this chapter discusses issues with the new oven designed in [9] and the solutions we implemented. Any future lab interested in a similar oven design can learn from the issues we encountered and the resulting changes to the initial design.

Chapter 5 discusses the effect of external electric field noise on the coherence time of an ultracold molecular qubit. The code written to calculate the vibrational state energies, polarizabilities, and Stark shifts can be very easily modified to work for other bi-alkali molecules, and possibly other diatomic molecules in general. The most notable result here, however, is the near-impossible requirement for the noise level of the polarizing DC electric field. A field on the order 10^5 V/m that is stable to within 0.26 V/m would be so difficult to achieve in the lab that this method of performing a dipole blockade is not actually

practical for our purposes. Therefore, a different method should be considered for this dipole blockade option. Future work will evaluate the potential of other methods, and once the best one is determined, attempt to implement it in the lab.

Bibliography

[1] Analog Devices. *EVAL-AD5791SDZ User Guide*. Analog Devices. Norwood, MA, USA, 2019.

[2] Kishor Bharti et al. "Noisy intermediate-scale quantum algorithms". In: *Reviews of Modern Physics* 94.1 (Feb. 2022). ISSN: 1539-0756. DOI: 10.1103/revmodphys.94.015004. URL: http://dx.doi.org/10.1103/RevModPhys.94.015004.

[3] Claude Cohen-Tannoudji and David Guery-Odelin. *Advances in Atomic Physics, An Overview*. World Scientific Publishing, 2011.

[4] D. DeMille. "Quantum Computation with Trapped Polar Molecules". In: *Physical Review Letters* 88.6 (Jan. 2002). ISSN: 1079-7114. DOI: 10.1103/physrevlett.88.067901. URL: http://dx.doi.org/10.1103/PhysRevLett.88.067901.

[5] Wolfgang Demtroder. *Molecular Physics: Theoretical Principles and Experimental Methods*. WILEY-VCH Verlag GmbH and Co. KGaA, 2005.

[6] Christopher J. Foot. *Atomic Physics*. Oxford University Press, 2005.

[7] Andrew T. Grier et al. "Lambda-enhanced sub-Doppler cooling of lithium atoms in D1 gray molasses". In: *Physical Review A* 87.6 (June 2013). ISSN: 1094-1622. DOI: 10.1103/physreva.87.063411. URL: http://dx.doi.org/10.1103/PhysRevA.87.063411.

[8] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. *Optical dipole traps for neutral atoms*. 1999. arXiv: physics/9902072 [physics.atom-ph]. URL: https://arxiv.org/abs/physics/9902072.

[9] Zach Hinkle. "Ultra-Cold CsLi Molecules as a Platform for Quantum Information Processing". MA thesis. University of Waterloo, 2022.

[10] L. Isenhower et al. "Demonstration of a Neutral Atom Controlled-NOT Quantum Gate". In: *Physical Review Letters* 104.1 (Jan. 2010). ISSN: 1079-7114. DOI: 10.1103/physrevlett.104.010503. URL: http://dx.doi.org/10.1103/PhysRevLett.104.010503.

[11] D. Jaksch et al. "Fast Quantum Gates for Neutral Atoms". In: *Physical Review Letters* 85.10 (Sept. 2000), pp. 2208–2211. ISSN: 1079-7114. DOI: 10.1103/physrevlett.85.2208. URL: http://dx.doi.org/10.1103/PhysRevLett.85.2208.

[12] Aviv Keshet. *Cicero Word Generator Technical and User Manual*. 2008. URL: https://akeshet.github.io/Cicero-Word-Generator/Cicero%20Technical%20and%20User%20Manual.pdf.

[13] Wolfgang Ketterle and N.J. Van Druten. "Evaporative Cooling of Trapped Atoms". In: ed. by Benjamin Bederson and Herbert Walther. Vol. 37. Advances In Atomic, Molecular, and Optical Physics. Academic Press, 1996, pp. 181–236. DOI: https://doi.org/10.1016/S1049-250X(08)60101-9. URL: https://www.sciencedirect.com/science/article/pii/S1049250X08601019.

[14] C. Kokail, C. Maier, and R. van Bijnen. "Self-verifying variational quantum simulation of lattice models". In: *Nature* 569 (2019), pp. 355–360. DOI: https://doi.org/10.1038/s41586-019-1177-4.

[15] Hela Ladjimi and Michał Tomza. "Diatomic molecules of alkali-metal and alkaline-earth-metal atoms: Interaction potentials, dipole moments, and polarizabilities". In: *Physical Review A* 109.5 (May 2024). ISSN: 2469-9934. DOI: 10.1103/physreva.109.052814. URL: http://dx.doi.org/10.1103/PhysRevA.109.052814.

[16] William Lunden. "Bose-Einstein condensation of 162Dy for spin-mixture experiments". PhD thesis. Massachusetts Institute of Technology, 2020.

[17] M. Marinescu and H.R. Sadeghpour. "Long-range potentials for two-species alkali-metal atoms". In: *Physical Review A* 59 (1999).

[18] *Molecular structure and spectra*. URL: https://titan.physx.u-szeged.hu/~dpiroska/atmolfiz/molecular_structure_and_spectra.html.

[19] John O'Hanlon. *A User's Guide to Vacuum Technology*. John Wiley and Sons, Inc., 2003.

[20] D. Rio Fernandes et al. "Sub-Doppler laser cooling of fermionic 40 K atoms in three-dimensional gray optical molasses". In: *EPL (Europhysics Letters)* 100.6 (Dec. 2012), p. 63001. ISSN: 1286-4854. DOI: 10.1209/0295-5075/100/63001. URL: http://dx.doi.org/10.1209/0295-5075/100/63001.

[21] Timur M. Rvachov et al. "Long-Lived Ultracold Molecules with Electric and Magnetic Dipole Moments". In: *Physical Review Letters* 119.14 (Oct. 2017). ISSN: 1079-7114. DOI: 10.1103/physrevlett.119.143001. URL: http://dx.doi.org/10.1103/PhysRevLett.119.143001.

[22] Timur M. Rvachov et al. "Two-photon spectroscopy of the NaLi triplet ground state". In: *Physical Chemistry Chemical Physics* 20.7 (2018), pp. 4739–4745. ISSN: 1463-9084. DOI: 10.1039/c7cp08481a. URL: http://dx.doi.org/10.1039/C7CP08481A.

[23] SAES Getters. *NEXTorr Z Pumps: Technical Specifications*. SAES Group. Lainate, Italy, 2021.

[24] M. Saffman and T. G. Walker. "Analysis of a quantum logic device based on dipole-dipole interactions of optically trapped Rydberg atoms". In: *Physical Review A* 72.2 (Aug. 2005). ISSN: 1094-1622. DOI: 10.1103/physreva.72.022347. URL: http://dx.doi.org/10.1103/PhysRevA.72.022347.

[25] Andreas Schindewolf et al. "Evaporation of microwave-shielded polar molecules to quantum degeneracy". In: *Nature* 607.7920 (July 2022), pp. 677–681. ISSN: 1476-4687. DOI: 10.1038/s41586-022-04900-0. URL: http://dx.doi.org/10.1038/s41586-022-04900-0.

[26] *Sisyphus cooling*. 2024. URL: https://en.wikipedia.org/wiki/Sisyphus_cooling.

[27] Terasic Technologies. *DE0-Nano User Manual*. Terasic Technologies. Hukou, Taiwan, 2021.

[28] E. Urban et al. "Observation of Rydberg blockade between two atoms". In: *Nature Physics* 5.2 (Jan. 2009), pp. 110–114. ISSN: 1745-2481. DOI: 10.1038/nphys1178. URL: http://dx.doi.org/10.1038/nphys1178.

[29] Gamma Vacuum. *Ion Pumps*. Tech. rep. Shakopee, MN, USA: Gamma Vacuum, Sept. 2021. URL: https://www.gammavacuum.com/media/pdf/c7/53/95/3601-7151-01_Ion-Pumps.pdf.

[30] Nikolay V. Vitanov et al. "Stimulated Raman adiabatic passage in physics, chemistry, and beyond". In: *Reviews of Modern Physics* 89.1 (Mar. 2017). ISSN: 1539-0756. DOI: 10.1103/revmodphys.89.015006. URL: http://dx.doi.org/10.1103/RevModPhys.89.015006.

[31] Rachel Wang. *Machine-learning Online Optimisation for Evaporative Cooling in Cold-atom Experiments*. URL: https://spectra.mathpix.com/article/2021.09.00018/online-optimisation-for-cold-atom-experiments.