Abstract

Towards a Bose-Einstein Condensate of SrF molecules

Varun Jorapur

2024

In recent years, there has been growing interest in methods for producing gases of ultracold polar molecules, driven by proposals to employ ultracold molecules in applications in ultracold chemistry, quantum information and quantum simulation, and precision measurements. There has been tremendous progress in producing ultracold polar molecules using pre-cooled and assembled alkali atoms, however, this faces limits in the molecule variety and interaction regimes that are accessible. Nonetheless, these techniques have successfully produced the first degenerate gases of polar molecules. Simultaneously, there has been a strong experimental and theoretical focus towards developing techniques to directly laser cool molecules which would enable accessing a diverse species of molecules in different interaction regimes.

Here, we review recent advances in the laser cooling and trapping of Strontium Monofluoride (SrF). Building upon previous work, we utilize velocity-selective coherent population trapping to substantially lower the temperature and truly reach the ultracold regime. We then describe how we can use this technique to load a conservative optical dipole trap, the first step towards observing collisions between the molecules. We also describe how we can use the interplay of the differential energy shifts caused by trap light polarizations and the cooling light itself to greatly enhance the loading and reduce the temperature in the trap, heralding large trap densities even with a low molecules number. We detail a new molecule source based on ultracold chemistry that is able to produce a larger molecule flux and a greatly increased experiment lifetime. Next, we detail a new and novel trapping technique that incorporates blue-detuned light, which can achieve large compression of the molecule cloud while cooling it at the same time. Using this new technique, we are able to achieve two orders of magnitude larger density, a big boost for loading the optical dipole trap. We demonstrate that, owing to the small size of this cloud, we are able to load an order of magnitude more molecules than before. With this high density, we demonstrate a measurement of the two-body inelastic collision rate in the trap, the first such demonstration in a bulk gas of directly cooled molecules. We describe ongoing work to achieve quantum control of the molecules in the trap, and the current experimental push to prepare the molecules in the rovibrational ground state. We briefly describe the design of a next generation apparatus which will enable the future experiments towards a BEC. We then describe our efforts towards implementing microwave shielding to lower the inelastic loss rate while also enhancing the elastic collision rate. This is a prerequisite to implementing evaporative cooling, and we describe some estimates of how well this will work in our system. We outline a path to efficient evaporative cooling and find that a BEC is within reach.

Towards a Bose-Einstein Condensate of SrF molecules

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Varun Jorapur

Dissertation Directors: Steve Lamoreaux and David P. DeMille
May 2024

Acknowledgements

First and foremost, I would like to thank my advisor, Prof. Dave DeMille, for offering me the opportunity to work on the SrF experiment. His endless patience for answering any and every question, his technical insights about problems in the lab, and his physical insights about the science, made this work possible. Dave, you are a truly fantastic advisor, and I think all your students really lucked out. I would also like to thank the members of my committee for their support over the years. Special thanks to Prof. Steve Lamoreaux for agreeing to be my nominal advisor when we moved to UChicago. Also thanks to Profs. Nir Navon, Dave Moore, and Danny McCarron for their support and advice.

I owe great thanks to the past and present SrF crew. Thanks to John Barry and Eric Norrgard for their outstanding theses that we still consult to this day, and that have set the precedent for this experiment, and many others in different universities. Thanks to Matt Steinecker for imparting his lab wisdom to me, and for teaching me what not to do, among other things. Also thanks to Matt for being a great friend and mentor, and for attending my wedding and dancing in it. Thanks to Yuqi Zhu for helping me with the ODT calculations, and for showing me how to continue to persevere even when the experiment does not cooperate. Thanks to Tom Langin for being with me throughout my PhD journey, and for imbibing in me the scientific mindset, and the technique of asking the right experimental questions. His method of combining physical insights with simulations really propelled this experiment forward, and I have learned a lot from him over the years. Finally, thanks to Qian Wang and Geoffrey Zheng, the next generation of the SrF crew, for their enthusiasm and willingness to get new things set up quickly for the next result, and for always being available for impromptu group meetings and discussions. Special thanks to Qian for building the new experiment control software, finally getting rid of the nightmare from Yale, and which made getting the results at UChicago a breeze. Also thanks to the Michelle, and Rohan, undergrads who worked with me over the years and built many lasers for our lab.

I would also like to thank the other DeMille group members for their help over the years, especially my cohort-mate, Jakob Kastelic, with whom I spent countless hours discussing problem sets, physics, and everything in life. I also owe a debt of gratitude to Sid Cahn, whose boundless enthusiasm and willingness to get things done no matter the time of day, made the experiment run smoothly. He was the backbone of the group at Yale, who knew where everything was, and how things could be done.

I would also like to thank my Yale Physics crew for undergoing the PhD journey together, and for being fantastic friends. From "solving" problem sets together in the common room, to the many werewolf/mafia game nights, to the GPSCY karaoke nights, they are what kept me sane throughout the initial years of the PhD.

I would also like to thank the many teachers and professors who have inspired me. Special

thanks to Prof. Manoj Harbola for being a great supporter of IITK physics students, and for giving us the opportunity to go to the IPhO as graders. Thanks to Prof. Kai Dieckmann for allowing me to be an intern in his lab at CQT, where I learned a lot about laser physics and AMO, which ultimately swayed me towards experimental physics. Special thanks also to Dr. Sarthak Chandra, who I was fortunate enough to be friends with at a very young age, and who inspired me to take up physics through his enthusiasm and passion for astronomy, and science in general.

I would be remiss if I did not thank my IITK physics group, a.k.a. Pappu gang, for making 3 years of physics so fun and enjoyable. Rutvij, Vidul, Shreya, Abu, Aditya, Uddipan, et al. were the best physics cohort ever, and I have learned so much from them. In fact, some of our physics conversations are what drove me towards experiments in the first place. I also feel proud that all of us are now doing/have done their PhDs, a truly remarkable feat.

Special thanks are owed to my Indian family at Yale, consisting of Milind, Vidul, Soham, and Sateja, without whose support I cannot imagine this PhD being completed. My first two years at Yale would have been impossible without the support of my roommate Milind, and we developed a self-sustaining ecosystem at 295 Humphrey. Through COVID, the Chicago move, and so much more, they all have been with me, and I owe them a debt of gratitude for being my family at Yale.

I would also like to thank my family for their unwavering support and constant belief in whatever I wanted to do since the beginning. Mom and Dad, you both are absolute gems of people and the best parents I could ever have hoped for. Your constant enthusiasm for life, for daring to dream, for always learning, and for putting family first, are the ideals I strive to live up to every day and were what kept the ship steady during tough times. Viraj, my elder brother, has been the one who always pushed me to try harder, to not settle for anything less than what he believed was the best, and indeed, he taught me 95% of all the skills I have learned in my life. Thanks also to Leah, who has so perfectly completed Viraj's life, and has, in turn, made our lives so much more joyful and colorful. I have learned how to always keep a calm mind, and also how to deliver a fantastic thesis defense from you. Thank you, LV for always being there for me.

I would also like to thank the family I gained through Sateja for their support during the latter years of my PhD. Aai and Baba, you have both accepted me as your son, and I feel very lucky to get to be a part of your family. Your approach to life of combining a fantastic work ethic with a strong family foundation are truly inspiring. I am also delighted that you share my enthusiasm for astronomy, and I am sure we will spend many nights under the skies. During, and after Covid, Sarvesh, Sateja's younger brother, was a big support, with his passionate interest in attending all of our events and celebrating the small things in life, and I feel like I now have a younger brother.

Finally, I would like to thank my wife, Sateja, without whom none of this would have been possible. I think it was destiny that made me ask you for a dance that fateful night, and I have never looked back. Your passion for science, your determination to complete experiments no matter the time of day, your strong desire to excel at whatever you choose, and your ability to find humor in the toughest of situations are so inspiring. Thank you for your fervent enthusiasm and backing for all of my endeavors, and for always accepting me, my lab, and my science.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

There has been growing interest and tremendous progress over the past decade in methods for producing gases of ultracold polar molecules [1, 2]. This interest has been fueled because ultracold polar molecules provide a new and diverse platform for ultracold chemistry [3–6], quantum computation [7–16], and precision tests of fundamental physics [17–25]. Polar molecules further provide a permanent electric dipole moment, which can lead to much stronger dipole-dipole interactions between molecules as compared to the weak magnetic dipole or collisional interactions in neutral atoms. Control over the internal and external molecular states of the molecules is key to many of these experiments. This control requires low internal and external energy, motivating research into techniques for producing ultracold molecules.

Ultra-cold temperatures are also necessary for creating Bose-Einstein condensates. In a general non-interacting gas, there are two length scales of interest: the inter-molecular length scale $n^{1/3}$ (where n is the density) and the de Broglie wavelength λ_{DB} which is defined as

$$\lambda_{DB} = \sqrt{\frac{h^2}{2\pi m k_B T}} \tag{1.1}$$

We can define a unitless quantity from these two length scales $\Phi = n\lambda_{DB}^3$, called the phase space density (PSD). One can see that when the temperature is high, the de Broglie wavelength is quite small than any interparticle spacing, and similarly so when the density of the gas is small. On the other hand, when a sufficiently dense gas is cooled to a very low temperature, the de Broglie wavelength can become quite large, and in this case, the wavefunction of the individual molecules starts to overlap. This is the regime characteristic of quantum degeneracy, making either a Bose-Einstein condensate (BEC) if the gas is made of bosons, or a Fermi degenerate gas, when the gas is made of fermions. Formally, the transition to a BEC occurs when $\Phi \geq 2.612$. Thus, to access the ultracold regime, it is necessary, not only to reduce the temperature, but to also increase the PSD in the process. This problem of increasing the PSD has been solved in alkali and alkaline-earth atoms: BECs of H [26], Li [27], Na [28], K [29], Rb [30], Cs [31], Ca [32], Sr [33, 34], Er [35] and Yb [36, 37] have been reported. However, it remains a rather tough challenge in directly cooled molecules.

Historically, there have been two approaches towards the goal of making molecular quantum degenerate gases: direct and indirect cooling. Indirect cooling involves creating gases of ultracold molecules by combining one or more species of pre-cooled atoms using either photoassociation [38, 39] or magneto-association [40, 41]. This method has been successful in producing high PSD gases

and quite recently, has been used to produce the first ever quantum degenerate Fermi gas of a polar molecule [42, 43] as well as a BEC of a polar molecule [44]. However, this method is limited to producing molecules of atoms which can be pre-cooled to ultracold temperatures and are thus mostly restricted to molecules consisting of alkali and alkaline-earth elements.

Direct laser cooling offers the opportunity to cool a wider range of molecules, and there has been considerable development in this field. An arsenal of techniques for both slowing molecular beams, such as Stark deceleration and white light slowing [45–49], and subsequently cooling the molecules, such as magneto-optical traps and sympathetic cooling [50–52] have been developed in recent years. The coldest molecular temperatures ($T \sim 5\mu$K [53, 54]) achieved through direct cooling of molecules implement laser cooling [53, 55–57].

Our lab employs direct laser cooling methods on the molecule $^{88}Sr^{19}F$ which has an electron spin $S = 1/2$ and a nuclear spin $I = 1/2$. One of the main challenges in direct laser cooling of molecules is the suppression of off-diagonal scattering events (scattering events in which the molecule decays into rovibrational states that are out of resonance with all the lasers, and hence these molecules are lost). The strongly diagonal Frank-Condon factors for SrF [2] allow us to solve this problem partially and achieve quasi-cycling, thus allowing us to apply laser cooling techniques to SrF, using just a few lasers.

1.1 Review of the current state of the art

The progress in direct laser cooling of molecules has followed quite closely the trends set by laser cooling of atoms. Over the past decade, our lab has demonstrated the laser slowing of a beam of SrF from a cryogenic buffer gas source [45], followed by loading into a radio-frequency magneto-optical trap (RF-MOT) [58, 59] and compressing the molecule cloud into a compressed MOT (cMOT) [60]. By optimizing the trapping sequence for such a cMOT, molecular temperatures as low as 250μK, close to the Doppler limit, and phase space density $\Phi \approx 6 \times 10^{-14}$ were achieved. Alongside our group, concurrent efforts were underway to trap CaF [61, 62] and YO molecules [63, 64].

During this period, there was tremendous progress in understanding the level structure of molecules and their amenability for laser cooling and trapping. Because of the complexity of such systems, one needs to resort to numerical simulations of the Optical Bloch Equations [65–67] and Mike Tarbutt's group was the leading force in this area. These simulations provided the necessary insight to really help us understand molecule structure vis-a-vis laser cooling, and how to further optimize and cool beyond the Doppler limit, perhaps even towards degeneracy.

Following the atomic recipe, after trapping the molecules in a MOT, our lab attempted sub-Doppler cooling. The peculiar level structure of molecules makes them unsuitable to standard red-detuned molasses. Instead one needs to resort to the blue-detuned gray molasses techniques [65]. Using this mechanism, our lab was able to produce SrF molecules cooled down to 50μK, cold enough to load these molecules into a conservative trap [55]. Alongside us, CaF and YO were also able to achieve sub-Doppler cooling [53, 56].

As with atoms, once sub-Doppler temperatures have been achieved, the molecules are loaded into a conservative trap, either a magnetic quadrupole trap (MQT) or an optical dipole trap (ODT). Transfer to a conservative trap paves the way for further cooling below the photon recoil

limit by eliminating heating from photon scatters. The advantage of a MQT is that the trap depth and the trap volume can be quite high, thus allowing easy trapping, with the caveat that molecules that land up too close to the center can be lost by Majorana spin-flip transitions. The advantage of an ODT is that the trapping is state independent, and high power lasers are easily available to accomplish this; the disadvantage is the small trapping volume, leading to low transfer efficiencies from large clouds. Our lab reported an MQT of SrF molecules in 2018 [55], but the density was not high enough to observe any interactions between the molecules in the trap. Around this time, CaF and YO were also able to load conservative ODTs [53, 68].

This proved to be a pivotal moment in molecule laser cooling and trapping. Groups started to explore different avenues, with CaF molecules being trapped in tweezers [13], and expanding the arsenal of molecules to polyatomic species [69–71]. New and more robust cooling schemes emerged based on utilizing dark states [53, 54]. Interactions between two molecules in a tweezer trap [72], and between atoms and molecules [73, 74] were observed.

However, significant challenges still remain towards the goal of making a BEC of a polar molecule. The low molecule numbers in typical conservative traps makes the prospect of direct evaporation daunting. Furthermore, even if the high numbers needed to start evaporation were achieved, the collisional properties in a bulk gas of molecules at low temperatures are not well understood, and have not been measured before. Thus, the primary focus of our lab has been to achieve high enough numbers to observe collisions in a conservative trap and characterize them. Alongside this, another avenue to increasing the PSD is to co-trap an atom, and sympathetically cool the molecules using the atom as the coolant. We are pursuing this endeavor with Rb atoms.

1.2 Organization of this thesis

In the remainder of this thesis, we will report a number of advances towards achieving the above goals of increasing molecule numbers and observing collisions. In Chapter 2, we review the molecule structure of SrF and calculate some of the important physical quantities for this thesis. Next, in Chapter 3 we review the current experimental apparatus and some of the changes made during the course of my PhD. In Chapter 4, we detail the first result of my PhD, the implementation of deep laser cooling of SrF molecules, and in Chapter 5, we follow this up with optical dipole trapping of SrF molecules. In Chapter 6, we detail a new molecule source based on ultracold chemistry between Sr and SF_6 gas. In Chapter 7, we document a new MOT scheme that continues to apply sub-Doppler forces inside a MOT. In Chapter 8, we report the observation of collisions between SrF molecules in a bulk gas. Finally, in Chapter 9 we review current and future directions for the experiment.

Chapter 2

SrF properties

2.1 Introduction

Understanding the behavior of any system under the application of external fields and forces is crucial to being able to fully harness the potential of that system. In this chapter, we summarize the structure of the SrF molecule and derive some useful quantities that are relevant to this work. We use the isotope $^{88}Sr^{19}F$ which has one unpaired electron, making it a spin $S = 1/2$ system. ^{88}Sr has no nuclear spin while ^{19}F has a nuclear spin $I = 1/2$. Thus SrF has $S = 1/2, I = 1/2$. The SrF molecule constants are available in Refs. [75, 76].

2.2 Hund's cases and spectroscopic notation

In any molecule, the couplings of the different angular momenta within the molecule determines the energy structure. In a diatomic molecule in particular, there is a single well-defined axis along the internuclear direction, which also defines a preferred direction for the coupling of angular momenta. Due to the Coulomb interaction, the electron angular momentum $\mathbf{L}$ is usually strongly coupled to the internuclear axis, and it rapidly precesses around the internuclear axis with a projection Λ. Similarly, due to the spin-orbit interaction, the electron spin $\mathbf{S}$ may be coupled to the internuclear axis and rapidly precess with a projection Σ. In some cases, as in heavy molecules, the spin-orbit interaction may by larger than the Coulomb interaction, resulting in $\mathbf{L}$ coupling to $\mathbf{S}$ first [1]. The total electronic projection along the internuclear axis is then $\Omega = \Lambda + \Sigma$. Furthermore, the molecule can rotate, with rotational angular momentum $\mathbf{R}$ which is perpendicular to the internuclear axis. In principle, these different angular momenta can couple in any order. However, to perform calculations, one must choose a basis to work in. The relative strengths of these couplings give rise to the idealized coupling cases, called the Hund's cases, see Figure 2.1. For SrF, only the case a and case b are relevant, though there are many different cases [77].

In Hund's case a, as is relevant for the Π states in SrF, $\mathbf{L}$ is strongly coupled to the internuclear axis with projection Λ, and $\mathbf{S}$ is strongly coupled to $\mathbf{L}$ and the internuclear axis, with projection Σ. Their sum is $\Omega = \Lambda + \Sigma$, and $\boldsymbol{\Omega}$ couples to the rotational quantum number $\mathbf{R}$ to form the total angular momentum $\mathbf{J}$. Hund's case a is appropriate when the spin-orbit interaction is large compared to the rotational interaction. In addition, the nuclear spin couples to $\mathbf{J}$ to form $\mathbf{F = J + I}$. A Hund's case a basis is expressed as $|\Lambda, S, \Sigma, \Omega, J, I, F, m_F\rangle$. Furthermore, both

[1] This is the so called Hund's case c.

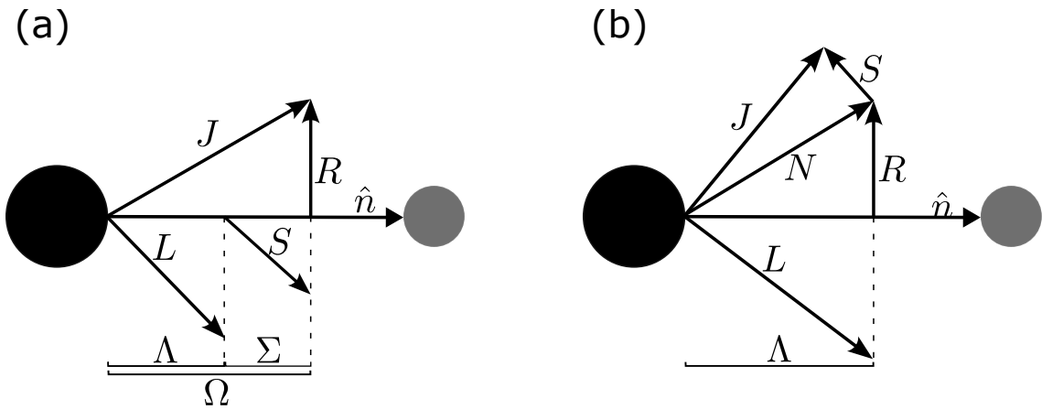

Figure 2.1: Vector coupling diagram for Hund's cases relevant to SrF. (a) Hund's case a and (b) Hund's case b.

L and **S** can have two opposite directions in which they can precess, leading to two projections $\pm\Lambda, \pm\Sigma$ and thus $\pm\Omega$. This is called Λ-doubling or Ω-doubling and we will see later how this affects the molecular structure.

In Hund's case b, as is relevant for the Σ states in SrF, $\Lambda = 0$ while $S \neq 0$, and so **L** is still coupled to the internuclear axis, but due to the weak spin-orbit interaction, **S** is no longer coupled to the internuclear axis. In this case, Λ is still a good quantum number, but Σ and Ω are not well defined. Λ and **R** combine to form **N**, the angular momentum excluding spin. The spin-rotation interaction is the next largest interaction, leading to **S** coupling to **N** to form the total angular momentum **J**. Hund's case b is appropriate when the spin-orbit interaction is small compared to the rotational interaction. In addition, the nuclear spin couples to **J** to form $\mathbf{F = J + I}$. A Hund's case b basis is expressed as $|\Lambda, S, N, J, I, F, m_F\rangle$.

Here, we have assumed that each successive interaction is much smaller than the previous one. These are idealized cases as stated before, and most states are well described by some intermediate case between a and b. For our purposes, it suffices to work in a single basis. Furthermore, other Hund's cases exist, corresponding to different couplings between the angular momenta, and these are documented elsewhere [77].

With the knowledge of "good" quantum numbers for molecules, we may go about spectroscopically labeling them. This is done using the following convention. The ground electronic state is labeled X, and excited states with the same spin multiplicity are labeled by capital Roman letters $A, B, C...$ in order of increasing energy. SrF has a single unpaired electron, and so $S = 1/2$ for all states. If different spin multiplicities exist, they are labeled by lower case Roman letters $a, b, c...$ in order of increasing energy.

This label is followed by a spectroscopic label of the form ${}^{2S+1}\Lambda^{\pm}_{\Omega}$. The symbol Λ denotes the projection of **L** on the internuclear axis as: $\Sigma : |\Lambda| = 0, \Pi : |\Lambda| = 1, \Delta : |\Lambda| = 2$ and so on. For states where Ω is a good quantum number, as for example case a states, this label is added as a subscript to the notation, otherwise it is omitted (for case b states). The $\pm$ superscript refers to symmtery or anti-symmetry upon reflection through a plane containing the internuclear axis - this is not to be confused with the parity of a state which we inspect later. With this notation, the ground electronic state for SrF is the $X^2\Sigma^+_{1/2}$ state, and most often we shall just write this as the $X^2\Sigma$ state. The first excited electronic state for SrF is the $A^2\Pi_{1/2}$ state, along with its spin-orbit partner with higher energy, the $A^2\Pi_{3/2}$ state.

2.3 Level structure

The general Hamiltonian for our molecule is as follows:

$$H = H_e + H_{\text{vib}} + H_{\text{SO}} + H_{\text{rot}} + H_{\text{SR}} + H_{\text{HFS}} + H_{\Lambda d} \tag{2.1}$$

where H_e, H_{vib}, H_{SO}, H_{rot}, H_{SR}, H_{HFS}, $H_{\Lambda d}$ are the electronic, vibrational, spin-orbit, rotational, spin-rotational, hyperfine, and Λ-doubling components of the Hamiltonian respectively. The order of these terms also roughly corresponds to the decreasing energy scales in the molecule.

2.3.1 Electronic structure

The electronic states are typically separated by 100s of THz. In SrF, the ground electronic state is the $X^2\Sigma$ state, and the first few excited electronic states are the $A^2\Pi_{1/2}, A^2\Pi_{3/2},$ and $B^2\Sigma$ states, with energies $2\pi \times 452$ THz, $2\pi \times 460$ THz, and $2\pi \times 518$ THz respectively. The separation between the $A^2\Pi_{1/2}$ and $A^2\Pi_{3/2}$ states also shows the scale of the spin-orbit coupling in SrF.

2.3.2 Vibrational structure and branching

Within a given electronic state, the energy due to the vibration of the nuclei is given by that of an anharmonic oscillator [77]:

$$G(v) = \omega_e(v + 1/2) - \omega_e x_e(v + 1/2)^2 + \omega_e y_e(v + 1/2)^3 + ... \tag{2.2}$$

where ω_e is the vibrational constant, and $\omega_e x_e$ and $\omega_e y_e$ are higher order terms that account for deviantions from the ideal harmonic oscillator potential. In most cases, the $\omega_e y_e$ term is very small and is ignored. The typical spacing between adjacent vibrational levels is on the few THz scale.

Due to the imperfect nature of the harmonic oscillator, there is often significant overlap between vibrational wavefunctions with different v. This overlap is quantified by the Franck-Condon factor (FCF) which is the overlap of an excited vibrational level (in an excited electronic state) $|v'\rangle$ with a lower vibrational level (in a lower electronic state) $|v\rangle$:

$$q_{v'v} = |\langle v'|v\rangle|^2 \tag{2.3}$$

For any random molecule in an excited state v', $q_{v'v}$ may be significant for ~ 10 or more values of v. This was long deemed to be the most significant hurdle in direct laser cooling of molecules - namely the requirement of a semi-infinite number of lasers to "plug" all these leaks. However, there exist a class of molecules [2], like SrF, where, the excitation of the electronic wavefunction barely perturbs the vibrational wavefunction, i.e. these are wavefunctions are well separated. These molecules are said to have nearly diagonal FCFs, and are now the workhorses of direct laser cooling of molecules. For a given electronic transition, one may also calculate the branching ratios:

$$b_{v'v} = \frac{q_{v'v}\omega_{v'v}^3}{\sum_{k=0}^{\infty} q_{v'k}\omega_{v'k}^3} \tag{2.4}$$

where $\omega_{v'v}$ is the transition energy. These ratios define the relative decay probability from an excited state into the ground vibrational manifold. For the transition $\left|X^2\Sigma(v = 0, N = 1)\right\rangle \rightarrow$

$\left|A^2\Pi_{1/2}(v=0, J=1/2^-)\right\rangle$ in SrF, a cycling scheme closed to $\sim 10^5$ photon scatters can be achieved using 2 repump lasers, or a scheme up to $\sim 5\times10^6$ photon scatters can be achieved using 3 repump lasers.

2.3.3 Rotational structure and branching

The next lower energy scale is the rotational structure, which spans the $\sim$ 10 GHz range. The rotational energy levels within a given vibrational level v are given by [77]:

$$F_v(J) = B_v J(J+1) - D_v J^2(J+1)^2 + ... \tag{2.5}$$

where B_v is the rotational constant, and D_v etc are higher order distortions. These constants depend on the vibrational level the molecule is in, and thus the rotational splitting is different for each vibrational level. For the ground state of SrF, the energy difference between the $N=0$ and $N=1$ states is around $2\pi \times 15$ GHz.

Another potential problem in optical cycling of molecules is the possible decay channels from an excited rotational state to multiple ground rotational states, each of which would need to be repumped to continue cycling. Rotational transitions are governed by the electric dipole selection rules, i.e. only transitions with $\Delta J = 0, \pm 1$, and between states of opposite parity, are allowed. Thus a typical excited state will decay to two rotational manifolds, with the exception being the lowest rotational excited state (for SrF the $A^2\Pi_{1/2}, J = 1/2^+$ state) which can only decay to a single rotational ground state. In SrF, this transition is the $J = 1/2^+ \rightarrow N = 1$ and is rotationally closed. We use this special case for optical cycling in SrF. However, this is a type-II transition, and has some negative consequences for cycling such as low trapping forces and sub-Doppler heating, as we shall see in Chapter 4.

2.3.4 Parity and Λ-doubling

The states within a Σ electronic manifold are parity eigenstates. For Σ^+ states, as in SrF, the parity is given by $(-1)^N$. However, Π states are composed of two nominally degenerate components of opposite parity with $\Lambda = \pm 1$. This degeneracy is lifted in second-order perturbation theory (see [77]) and separates the two states of opposite parity by:

$$\hbar\omega_{\Lambda d} = \mp(-1)^{J-1/2}(p+2q)(J+1/2)/2 \tag{2.6}$$

For the $A^2\Pi_{1/2}$ state, $p = -0.13269(22)$ cm^{-1} and $q = -1.1(7)\cdot10^{-4}$ cm^{-1} [76], which results in the $J = \pm 1/2$ states being split by $\approx 2\pi \times 4$ GHz.

2.3.5 Spin-rotation and hyperfine structure

The hyperfine levels are the lowest energy scales in the molecule, spanning the $\sim$ 100 MHz range. In the ground state of SrF, the relevant spin-rotation and hyperfine hamiltonian is:

$$H_{\mathrm{SR/HF}} = \gamma_{vN}\mathbf{S}\cdot\mathbf{N} + b_{vN}\mathbf{I}\cdot\mathbf{S} + c_{vN}(\mathbf{I}\cdot\hat{z})(\mathbf{S}\cdot\hat{z}) + C_{1vN}\mathbf{I}\cdot\mathbf{N} \tag{2.7}$$

where γ_{vN} is the spin-rotation constant, b_{vN} is the hyperfine constant, c_{vN} is the dipole-dipole constant and C_{1vN} is the nuclear spin-rotation constant. Here, $\mathbf{N}$ and $\mathbf{S}$ combine to produce $\mathbf{J}$, which in turn combines with $\mathbf{I}$ to produce $\mathbf{F}$, the grand total angular momentum. Since $S = 1/2$ and $I = 1/2$ for SrF, all the rotational levels are split into 4 hyperfine levels, except for the $N = 0$ level which has only 2 hyperfine levels. To continue cycling photons, it is necessary to address all these levels, as will be seen in Chapter 3. The hyperfine hamiltonian mixes states with the same F but different J, which is called J-mixing. On diagonalizing the full Hamiltonian, the J-mixing coefficients may be found and are expressed in Table 2.1.

For the $A^2\Pi_{1/2}(J = 1/2)$ state, the hyperfine structure is unresolved (split by < 3 MHz) and we will not go into the details here.

2.4 Matrix elements

Next, we look at calculating some of the quantities of interest for SrF, and for that we first start with calculating the matrix elements for the dipole allowed transitions in SrF. We want to calculate something of the form $\langle e|\vec{r}|g\rangle$. First, for such problems, it is best to work in the circular co-ordinate system and in this case, the 3 different polarizations are related to the Cartesian polarizations as:

$$\hat{e}_{+1} = -\frac{1}{\sqrt{2}}(\hat{e}_x + i\hat{e}_y), \quad \hat{e}_0 = \hat{e}_z, \quad \hat{e}_{-1} = \frac{1}{\sqrt{2}}(\hat{e}_x - i\hat{e}_y) \tag{2.8}$$

These are all rank 1 tensors and hence it is easy to calculate matrix elements. Note also that the $p = \pm 1, 0$ correspond to right/left circular and linear polarizations respectively.

In order to evaluate the matrix elements, we also need to know how to go from the lab-fixed frame to the molecule-fixed frame and this procedure is given in [77] §5.5.5 as well as in [78] §2.4. We briefly summarize this here. We follow the notation in [77] and use p for the lab-fixed components and q for the molecule-fixed components. Then, an operator in the lab-fixed frame $T^k_p(A)$ can be decomposed as

$$T^k_p(A) = \sum_q \mathcal{D}^{(k)}_{pq}(\omega)^* T^k_q(A) \quad ; \quad \mathcal{D}^{(k)}_{pq}(\omega)^* = (-1)^{p-q}\mathcal{D}^{(k)}_{-p,-q}(\omega) \tag{2.9}$$

where $\mathcal{D}^{(k)}_{pq}(\omega)^*$ is the complex conjugate of the pq element of the kth rank rotation matrix $D^{(k)}(\omega)$. Now, considering that a diatomic molecule behaves like a symmetric top molecule, we can write the rotational part of the wavefunction as ([77] eq 5.145):

$$|J, \Omega, M\rangle = [(2J + 1)/8\pi^2]^{1/2}\mathcal{D}^{(J)}_{M\Omega}(\omega)^* \tag{2.10}$$

Hence, we have the identity:

$$\langle J, \Omega, M|\mathcal{D}^{(k)}_{pq}(\omega)^*|J', \Omega', M'\rangle = (-1)^{M-\Omega}[(2J + 1)(2J' + 1)]^{1/2} \begin{pmatrix} J & k & J' \\ -\Omega & q & \Omega' \end{pmatrix} \begin{pmatrix} J & k & J' \\ -M & p & M' \end{pmatrix} \tag{2.11}$$

Now, we see that since M is just the projection of J in the lab-fixed frame, we can use the Wigner-Eckhart theorem on the left hand side of the above equation to get rid of p and find what the

reduced matrix element is in the lab-frame with respect to the molecule-frame:

$$\langle J,\Omega,M|\mathcal{D}^{(k)}_{pq}(\omega)^*|J',\Omega',M'\rangle = (-1)^{J-M}\begin{pmatrix} J & k & J' \\ -M & p & M' \end{pmatrix}\langle J,\Omega||\mathcal{D}^{(k)}_{.q}(\omega)^*||J',\Omega'\rangle \tag{2.12}$$

and comparing these expressions, we get that

$$\langle J,\Omega||\mathcal{D}^{(k)}_{.q}(\omega)^*||J',\Omega'\rangle = (-1)^{J-\Omega}[(2J+1)(2J'+1)]^{1/2}\begin{pmatrix} J & k & J' \\ -\Omega & q & \Omega' \end{pmatrix} \tag{2.13}$$

where the . indicates that the matrix element is reduced as far as the lab-frame is concerned but is not reduced in the molecule-frame. Now our operator r^1_p can be written as:

$$r^1_p = \sum_q \mathcal{D}^1_{pq}(\omega)^* r^1_q \tag{2.14}$$

and so, we get that the reduced matrix element of r in the lab-frame will be

$$\langle \Lambda',S',\Sigma',J',\Omega'||r^1||\Lambda,S,\Sigma,J,\Omega\rangle = \sum_q \langle \Lambda',S',\Sigma',J',\Omega'||\mathcal{D}^1_{.q}(\omega)^* r^1_q||\Lambda,S,\Sigma,J,\Omega\rangle \tag{2.15}$$

Now, the molecule-fixed frame $\mathcal{D}^1_{.q}$ acts only on J,Ω part of the wavefunction, and the molecule-fixed frame r^1_q acts on the rest. In this way, we can separate out the two, to get:

$$\langle \Lambda',S',\Sigma',J',\Omega'||r^1||\Lambda,S,\Sigma,J,\Omega\rangle = \sum_q \langle J',\Omega'||\mathcal{D}^1_{.q}(\omega)^*||J,\Omega\rangle\ \langle \Lambda',S',\Sigma'|r^1_q|\Lambda,S,\Sigma\rangle \tag{2.16}$$

We can replace the angular part with what we calculated before and further simplify

$$\langle \Lambda',S',\Sigma',J',\Omega'||r^1||\Lambda,S,\Sigma,J,\Omega\rangle = \sum_q (-1)^{J'-\Omega'}[(2J+1)(2J'+1)]^{1/2}\begin{pmatrix} J' & k & J \\ -\Omega' & q & \Omega \end{pmatrix} \langle \Lambda',S',\Sigma'|r^1_q|\Lambda,S,\Sigma\rangle \tag{2.17}$$

This equation is of great importance to us because we will be using it to calculate the matrix elements in the next part.

Next, we look at the general procedure for calculating the matrix element between a final state (we are working in the case a basis here) $|\{\eta'\}\rangle = |\Lambda',S,\Sigma',\Omega',J',F',M'_F\rangle$ and an initial state $|\{\eta\}\rangle = |\Lambda,S,\Sigma,\Omega,J,F,M_F\rangle$ for a p dipole transition. Here, we begin with $\langle r^1_p\rangle$ where r is in the lab-fixed co-ordinate system. We successively apply the Wigner-Eckhart theorm and the Spectator theorem to first reduce $F,M_F \to F \to J$ and then we have to transform to the molecule-fixed co-ordinate system to further simplify the matrix element. The procedure is outlined below:

$$\langle \Lambda',S,\Sigma',\Omega',J',F',M'_F|r^1_p|\Lambda,S,\Sigma,\Omega,J,F,M_F\rangle = (-1)^{F'-M'_F}\begin{pmatrix} F' & 1 & F \\ -M'_F & p & M_F \end{pmatrix} \langle \Lambda',S,\Sigma',\Omega',J',F'||r^1||\Lambda,S,\Sigma,\Omega,J,F\rangle \tag{2.18}$$

$$\langle \Lambda', S, \Sigma', \Omega', J', F' || r^1 || \Lambda, S, \Sigma, \Omega, J, F \rangle = (-1)^{J'+I+F'+1} \begin{Bmatrix} J' & F' & I \\ F & J & 1 \end{Bmatrix} \sqrt{2F+1}\sqrt{2F'+1} \langle \Lambda', S, \Sigma', \Omega', J' || r^1 || \Lambda, S, \Sigma, \Omega, J \rangle \tag{2.19}$$

In order to simplify this, we need to go from the lab-fixed frame to the molecule-fixed frame:

$$\langle \Lambda', S, \Sigma', \Omega', J' || r^1 || \Lambda, S, \Sigma, \Omega, J \rangle = \sum_{q=-1}^{q=+1} (-1)^{J'-\Omega'} \sqrt{2J+1}\sqrt{2J'+1} \begin{pmatrix} J' & 1 & J \\ -\Omega' & q & \Omega \end{pmatrix} \langle \Lambda', S, \Sigma' | r^1_q | \Lambda, S, \Sigma \rangle \tag{2.20}$$

$$= (-1)^{J'-\Omega'} \sqrt{2J+1}\sqrt{2J'+1} \begin{pmatrix} J' & 1 & J \\ -\Omega' & \Omega'-\Omega & \Omega \end{pmatrix} \langle \Lambda', S, \Sigma' | r^1_{\Omega'-\Omega} | \Lambda, S, \Sigma \rangle \tag{2.21}$$

Since the dipole operator cannot change spin, we require that $\Sigma = \Sigma'$ and S is always the same, so we can eliminate the spin part of the wavefunction in this manner to leave, in the end:

$$\begin{aligned} \langle \Lambda', S, \Sigma', \Omega', J', F', M'_F | r^1_p | \Lambda, S, \Sigma, \Omega, J, F, M_F \rangle &= (-1)^{F'-M'_F} \begin{pmatrix} F' & 1 & F \\ -M'_F & p & M_F \end{pmatrix} \\ &\quad (-1)^{J'+I+F'+1} \begin{Bmatrix} J' & F' & I \\ F & J & 1 \end{Bmatrix} \sqrt{2F+1}\sqrt{2F'+1} \\ &\quad (-1)^{J'-\Omega'} \sqrt{2J+1}\sqrt{2J'+1} \begin{pmatrix} J' & 1 & J \\ -\Omega' & \Omega'-\Omega & \Omega \end{pmatrix} \\ &\quad \langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle \, \delta_{\Sigma,\Sigma'} \\ &= C_{\eta'\eta} \langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle \, \delta_{\Sigma,\Sigma'} \end{aligned} \tag{2.22}$$

This is the end product for the matrix element for a p dipole transition between any two states expressed in the Hund's case a basis. The term $C_{\eta'\eta}$ is just the product of all the angular factors, and the last term $\langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle$ is what remains after all the angular momenta have been decoupled and it is the same for all the states for any given pair Λ, Λ'. Since $\Omega' - \Omega = \Lambda' - \Lambda = \pm 1, 0$, we have 3 terms here and are typically written as $\langle \Lambda' | r^1_{\pm 1} | \Lambda \rangle = D_\perp$ and $\langle \Lambda' | r^1_0 | \Lambda \rangle = D_\parallel$. These are usually hard to calculate but can be determined experimentally by relating it to the linewidth, which will be described in the next part.

2.4.1 The term $\langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle$

We cannot proceed to calculate the trap depth for our ODT without knowing this last term. However, the upside about calculating this term is that it needs to be done only once per Λ, Λ'. This term can be calculated by relating it to the excited state linewidths and the way to do this for SrF is described below. Let us look at a sample calculation for the excited state $A^2\Pi_{1/2}$ with $|\Lambda| = 1$ and the ground state $X^2\Sigma_{1/2}$ with $\Lambda = 0$ i.e. we want $\left\langle A^2\Pi_{1/2}; \Lambda' = 1 \middle| r^1_1 \middle| X^2\Sigma; \Lambda = 0 \right\rangle = \left\langle A^2\Pi_{1/2}; \Lambda' = -1 \middle| r^1_{-1} \middle| X^2\Sigma; \Lambda = 0 \right\rangle$. Here we take $\mathbf{I} = \mathbf{0}$ for simplicity.

- First, we need to find a specific state in the Λ' manifold that can only decay to states in the

Λ manifold. For this calculation, we look at the state

$\left|A^2\Pi_{1/2}; \Lambda' = -1, S = 1/2, \Sigma' = 1/2, \Omega' = -1/2, J' = 1/2^+, M'_J = 1/2\right\rangle$

- Then, we identify all the decay channels, and here we note that this can only decay to the states
  - $X^2\Sigma, N = 1, J = 1/2, M_J = \pm 1/2$
  - $X^2\Sigma, N = 1, J = 3/2, M_J = \pm 1/2$
  - $X^2\Sigma, N = 1, J = 3/2, M_J = 3/2$
- For each of these decay channels, we calculate the matrix element

$$|M_{\eta'\eta}|^2 = |C_{\eta'\eta}|^2 \times |\left\langle A^2\Pi_{1/2}; \Lambda' = -1\right|r^1_{-1}\left|X^2\Sigma; \Lambda = 0\right\rangle|^2$$

  The table for the coefficients $|C_{\eta'\eta}|^2$ is given below:

Decay Channel	$\lvert C_{\eta'\eta}\rvert^2$
$\lvert J' = 1/2^+, M'_J = 1/2\rangle \to \lvert J = 1/2, M_J = 1/2\rangle$	$\frac{2}{9}$
$\lvert J' = 1/2^+, M'_J = 1/2\rangle \to \lvert J = 1/2, M_J = -1/2\rangle$	$\frac{4}{9}$
$\lvert J' = 1/2^+, M'_J = 1/2\rangle \to \lvert J = 3/2, M_J = 1/2\rangle$	$\frac{1}{18}$
$\lvert J' = 1/2^+, M'_J = 1/2\rangle \to \lvert J = 3/2, M_J = -1/2\rangle$	$\frac{1}{9}$
$\lvert J' = 1/2^+, M'_J = 1/2\rangle \to \lvert J = 3/2, M_J = 3/2\rangle$	$\frac{1}{6}$

- Now, we note that the partial linewidth for a transition from a given excited state $\{\eta'\}$ to a given ground state $\{\eta\}$ is given by

$$\begin{aligned}\Gamma_{\eta'\eta} &= \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}|\left\langle\eta'|\mu|\eta\right\rangle|^2 \\ &= \frac{e^2\omega_0^3}{3\pi\epsilon_0\hbar c^3}|M_{\eta'\eta}|^2 \\ &= \frac{e^2\omega_0^3}{3\pi\epsilon_0\hbar c^3}|C_{\eta'\eta}|^2 \times |\left\langle A^2\Pi_{1/2}; \Lambda' = -1\right|r^1_{-1}\left|X^2\Sigma; \Lambda = 0\right\rangle|^2 \end{aligned} \tag{2.23}$$

- We also know, for SrF that the total linewidth for each excited state is the same, and for the $A^2\Pi_{1/2}$ state it is $\Gamma_{\eta'} = 2\pi \times 6.6$ MHz.
- Hence we see that if we sum over $\{\eta\}$, we should get the total linewidth i.e.

$$\begin{aligned}\Gamma_{\eta'} &= \sum_\eta \frac{e^2\omega_0^3}{3\pi\epsilon_0\hbar c^3}|C_{\eta'\eta}|^2 \times |\left\langle A^2\Pi_{1/2}; \Lambda' = -1\right|r^1_{-1}\left|X^2\Sigma; \Lambda = 0\right\rangle|^2 \\ &= \frac{e^2\omega_0^3}{3\pi\epsilon_0\hbar c^3}|\left\langle A^2\Pi_{1/2}; \Lambda' = -1\right|r^1_{-1}\left|X^2\Sigma; \Lambda = 0\right\rangle|^2 \times \sum_\eta |C_{\eta'\eta}|^2 \end{aligned} \tag{2.24}$$

- Performing this sum (which turns out to equal 1) and using the known values for the energy (see §3.1 for values), We find that

$$\left\langle A^2\Pi_{1/2}; \Lambda' = -1\right|er^1_{-1}\left|X^2\Sigma; \Lambda = 0\right\rangle = 6.21676 \text{ D} \tag{2.25}$$

- Looking at the values quoted for the Transition Dipole Moment in [79] Table XI, we see that their dipole moment value is 6.67529 D and the experimental value is 6.2279 D. Since the

two are pretty close, we shall not be calculating this dipole moment for the other states in this thesis as the paper already quotes a value.

2.4.2 Calculation of matrix elements taking into account Hund's cases and parity

In the above example calculation for the term $\langle\Lambda'|r^1_{\Omega'-\Omega}|\Lambda\rangle$ we had made a simplification about identifying all the possible decay channels in that we had not decomposed the Σ states into their respective Hund's case a forms. We had also not decomposed the $A^2\Pi_{1/2}$ states into the eigenstates of parity. In the end, both of these simplifications work out to be correct for us and there is actually no need to decompose the states, and here We will give 2 explicit examples to show that that is indeed true.

2.4.2.1 Decomposing case b into case a

The procedure to convert from case b to case a has a typo in Brown and Carrington and the right way to convert them is given in [75] and it is as follows:

$$|\Lambda; N, S, J\rangle = \sum_{\Omega=-1/2}^{1/2}\sum_{\Sigma=-1/2}^{1/2}(-1)^{J+\Omega}\sqrt{2N+1}\begin{pmatrix} S & N & J \\ \Sigma & \Lambda & -\Omega\end{pmatrix}|\Lambda, S, \Sigma, \Omega, J\rangle \tag{2.26}$$

Here, We decompose all the Σ states that are of interest to us:

- $|0; 0, 1/2, 1/2\rangle = \frac{1}{\sqrt{2}}\big[\,|0, 1/2, -1/2, -1/2, 1/2\rangle + |0, 1/2, 1/2, 1/2, 1/2\rangle\,\big]$
- $|0; 1, 1/2, 1/2\rangle = \frac{1}{\sqrt{2}}\big[\,|0, 1/2, -1/2, -1/2, 1/2\rangle - |0, 1/2, 1/2, 1/2, 1/2\rangle\,\big]$
- $|0; 1, 1/2, 3/2\rangle = \frac{1}{\sqrt{2}}\big[\,|0, 1/2, -1/2, -1/2, 3/2\rangle + |0, 1/2, 1/2, 1/2, 3/2\rangle\,\big]$
- $|0; 2, 1/2, 3/2\rangle = \frac{1}{\sqrt{2}}\big[\,|0, 1/2, -1/2, -1/2, 3/2\rangle - |0, 1/2, 1/2, 1/2, 1/2\rangle\,\big]$
- $|0; 2, 1/2, 5/2\rangle = \frac{1}{\sqrt{2}}\big[\,|0, 1/2, -1/2, -1/2, 5/2\rangle + |0, 1/2, 1/2, 1/2, 5/2\rangle\,\big]$

2.4.2.2 Parity states for $A^2\Pi_{1/2}$

The eigenstates of correct parity are noted below using formula 6.234 from [77] and are written in the form $|\Lambda, S, \Sigma, \Omega, J\rangle$

- $\big|A^2\Pi_{1/2}, J = 1/2, +\big\rangle = \frac{1}{\sqrt{2}}\big[\,|1, 1/2, -1/2, 1/2, 1/2\rangle + |-1, 1/2, 1/2, -1/2, 1/2\rangle\,\big]$
- $\big|A^2\Pi_{1/2}, J = 3/2, +\big\rangle = \frac{1}{\sqrt{2}}\big[\,|1, 1/2, -1/2, 1/2, 3/2\rangle - |-1, 1/2, 1/2, -1/2, 1/2\rangle\,\big]$

2.4.2.3 Matrix elements between $X\Sigma$ and $B\Sigma$

Let us look at 1 specific example of this. We will calculate the matrix elements between the states $\big|X^2\Sigma; N = 1, J = 1/2, M_J = 1/2\big\rangle$ and $\big|B^2\Sigma; N = 0, J = 1/2, M_J = 1/2\big\rangle$ since these 2 states are of opposite parity and hence have non-zero matrix elements. Then, the matrix element between them is:

$$\begin{aligned} M_{\eta'\eta} = \frac{1}{2}\big[\,\langle 0, 1/2, -1/2, -1/2, 1/2, 1/2| + \langle 0, 1/2, 1/2, 1/2, 1/2, 1/2|\,\big]|r^1_0| \\ \big[\,|0, 1/2, -1/2, -1/2, 1/2, 1/2\rangle - |0, 1/2, 1/2, 1/2, 1/2, 1/2\rangle\,\big] \end{aligned} \tag{2.27}$$

Now, we note that only the terms that have the same spin (Σ) have a non-zero matrix element. Hence, the 4 terms are reduced to 2:

$$\begin{aligned} M_{\eta'\eta} &= \frac{1}{2}\big[\,\langle 0,1/2,-1/2,-1/2,1/2,1/2|r_0^1|0,1/2,-1/2,-1/2,1/2,1/2\rangle \\ &- \langle 0,1/2,1/2,1/2,1/2,1/2|r_0^1|0,1/2,1/2,1/2,1/2,1/2\rangle\,\big] \\ &= \frac{1}{2}\left[\frac{1}{3}\langle\Lambda'=0|r_0^1|\Lambda=0\rangle - \frac{-1}{3}\langle\Lambda'=0|r_0^1|\Lambda=0\rangle\right] \\ &= \frac{1}{3}\langle\Lambda'=0|r_0^1|\Lambda=0\rangle \end{aligned} \tag{2.28}$$

which is exactly the same as the first term

$$\langle 0,1/2,-1/2,-1/2,1/2,1/2|r_0^1|0,1/2,-1/2,-1/2,1/2,1/2\rangle \tag{2.29}$$

This is true for the other states as well between the $X\Sigma$ and the $B\Sigma$ manifolds. Hence, we see here that even though the correct way is to decompose the states into a superposition of Hund's case a states, in the end it boils down to not making a difference.

2.4.2.4 Matrix elements between $X\Sigma$ and $A\Pi$

Let us look at 1 specific example of this as well. We will calculate the matrix elements between the states $\left|X^2\Sigma; N=1, J=1/2, M_J=1/2\right\rangle$ and $\left|A^2\Pi_{1/2}, J=1/2^+, M_J=1/2\right\rangle$ since these 2 states are of opposite parity and hence have non-zero matrix elements. Then, the matrix element between them is:

$$\begin{aligned} M_{\eta'\eta} = \frac{1}{2}\big[\,\langle 1,1/2,-1/2,1/2,1/2,1/2| + \langle -1,1/2,1/2,-1/2,1/2,1/2|\,\big]|r_0^1| \\ \big[\,|0,1/2,-1/2,-1/2,1/2,1/2\rangle - |0,1/2,1/2,1/2,1/2,1/2\rangle\,\big] \end{aligned} \tag{2.30}$$

Again, we see that only states with the same spin have a non-zero matrix element, and so the 4 terms are reduced to 2:

$$\begin{aligned} M_{\eta'\eta} &= \frac{1}{2}\big[\,\langle 1,1/2,-1/2,1/2,1/2,1/2|r_0^1|0,1/2,-1/2,-1/2,1/2,1/2\rangle \\ &- \langle -1,1/2,1/2,-1/2,1/2,1/2|r_0^1|0,1/2,1/2,1/2,1/2,1/2\rangle\,\big] \\ &= \frac{1}{2}\left[\frac{-\sqrt{2}}{3}\langle\Lambda'=1|r_{+1}^1|\Lambda=0\rangle - \frac{\sqrt{2}}{3}\langle\Lambda'=-1|r_{-1}^1|\Lambda=0\rangle\right] \\ &= -\frac{\sqrt{2}}{3}\langle\Lambda'=-1|r_{-1}^1|\Lambda=0\rangle \end{aligned} \tag{2.31}$$

which is exactly the same as the term

$$\langle -1,1/2,1/2,-1/2,1/2,1/2|r_0^1|0,1/2,1/2,1/2,1/2,1/2\rangle \tag{2.32}$$

Again, this is true for other states in the $X\Sigma$ and $A\Pi$ manifolds. Hence, we see that we also do not need to decompose the $A\Pi$ states into their states of correct parity.

We will use this formalism for all the calculations in this thesis - first convert the operator from the lab frame to the molecule frame, then use the Wigner-Eckhart and spectator theorems to calculate the angular part, and use the spectroscopy data to calculate the radial part.

2.5 Zeeman Shift

Magnetic fields are produced by any simple current loop, and are thus very easy to generate in the lab. The ability to apply magnetic fields easily makes this a highly versatile tool. For example, in a traditional MOT, the combination of a Doppler cooling force (with red detuned light) and the application of a magnetic field gradient, produces a position dependent restoring force, allowing trapping and cooling. One can use magnetic fields to remix dark Zeeman sublevels, as is done for white light slowing. In other cases, magnetic fields can be harmful - for example, if one is trying to achieve good coherence between two states (such as the $|N=0, F=1\rangle$ and $|N=1, F=0\rangle$ states in SrF for instance), then any stray magnetic field can cause the energy difference between these states to shift, reducing the coherence times drastically. To put concrete numbers to this, if one wants to achieve a coherence time ~ 1 s (as one would want for a decent EDM experiment), then, the fluctuation of the energy difference between the states needs to be less than ~ 1 Hz, which would in turn determine what level of field stability is needed based on the relative behavior of the states under a magnetic field.

It is thus crucial to understand the behavior of the energy levels under the influence of a magnetic field, in order to make full use of this tool. The Zeeman shift details the behavior of a system under the effect of an applied magnetic field, with the Hamiltonian [77]:

$$H_Z = g_S \mu_B \mathbf{S} \cdot \mathbf{B} + g_L \mu_B \mathbf{L} \cdot \mathbf{B} - g_I \mu_N \mathbf{I} \cdot \mathbf{B} \tag{2.33}$$

where $g_S \approx 2, g_L \approx 1$ and $g_I \approx 5.585$ are the electron, electron orbital, and the nuclear g-factors respectively, μ_B is the Bohr magneton, and μ_N is the nuclear magneton. $\frac{\mu_N}{\mu_B} \approx \frac{1}{1836}$, and so the effect of the last term is very small; however, for cases where the energy shift is needed to be known to less than the % level, one needs to account for this term as well.

For molecular Σ states, $\mathbf{L} = 0$, and thus for the ground $X^2\Sigma_{1/2}$ state of SrF, we need only look at the first and third terms. For a field along the z-direction, the Hamiltonian is:

$$H_Z = (g_S \mu_B S_z - g_I \mu_N I_z) B_z \tag{2.34}$$

For low fields, one may write an equivalent shift:

$$H_Z = g_F \mu_B F_z B_z \tag{2.35}$$

where g_F is the Landé g-factor, given by:

$$g_F = g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} + g_I \left(\frac{\mu_N}{\mu_B}\right) \frac{F(F+1) + I(I+1) - J(J+1)}{2F(F+1)} \tag{2.36}$$

$$g_J = g_S \frac{J(J+1) - N(N+1) + S(S+1)}{2J(J+1)} \tag{2.37}$$

The g-factors for the hyperfine manifold of SrF are in Table 2.1, where, the g-factors of the two $F = 1$ states are modified due to J-mixing.

Nominal label	J composition	No J mixing	With J mixing
$\|J = 3/2, F = 2\rangle$	$\|J = 3/2, F = 2\rangle$	0.5	0.5
$\|J = 3/2, F = 1\rangle$	$0.888\,\|J = 3/2, F = 1\rangle + 0.4598\,\|J = 1/2, F = 1\rangle$	0.83	0.97
$\|J = 1/2, F = 0\rangle$	$\|J = 1/2, F = 0\rangle$	0	0
$\|J = 1/2, F = 1\rangle$	$0.888\,\|J = 1/2, F = 1\rangle - 0.4598\,\|J = 3/2, F = 1\rangle$	-0.33	-0.47

Table 2.1: g-factors of the $X^2\Sigma(v = 0, N = 1)$ states. The table shows the g-factors both with and without J-mixing, and can be used for small B-fields.

Beyond a few 10s of G field, the energy shift becomes comparable to the hyperfine splitting. Thus, to fully account for the Zeeman shift, one needs to construct the full Hamiltonian including the Zeeman term, and simultaneously diagonalize it for the applied field. The expression for these terms are given in [75, 77], and can be calculated based on the formalism developed before. The Zeeman shift for the $N = 1$ hyperfine manifold is shown in Figure 2.2. Beyond a field of around 100 G, the spin states are completely separated, i.e. the electron and nuclear spins are decoupled from the axis of rotation. In this case, we enter the linear Zeeman regime, where the energy shift is now dominated by the terms in eq. 2.34, i.e. only due to the electronic and nuclear spins. Here, the F is no longer a good quantum number, and one must resort to the bare spin states.

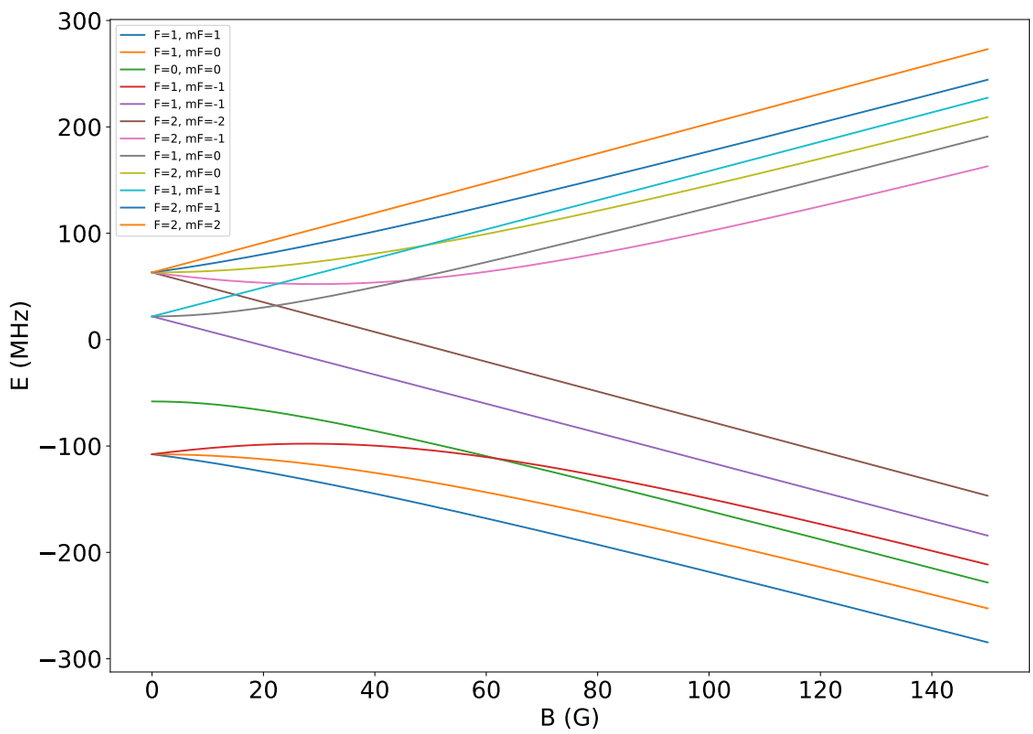

Figure 2.2: Energy shifts in the presence of a B-field. Beyond around 100 G field, the spin states are completely separated, meaning the electron and nuclear spins are decoupled from the axis of rotation.

The story is quite different for the first excited $A^2\Pi_{1/2}$ state. Here, $\mathbf{S}$ and $\mathbf{L}$ are both pinned to the internuclear axis, as in a Hund's case a basis. For a $\Omega = 1/2$ state, Λ and Σ point in opposite directions, with $\Lambda = 2\Sigma$. However, since $g_S \approx 2g_L$, both of these angular momenta almost exactly cancel each other, leaving the $A^2\Pi_{1/2}$ state with a nearly zero g-factor. There is still some mixing of the $A^2\Pi_{1/2}$ with the neighboring $^2\Sigma$ states, which gives them some non-zero character, but it is still quite small. For example, for SrF, the g-factor for the $A^2\Pi_{1/2}$ state is ≈ -0.088. This has adverse effects in the trapping forces that can be generated for molecular MOTs, as that force is $\propto g_e/g_g$, i.e. it strongly depends on the ratio of the excited state to ground state g-factors [65].

2.6 DC Stark Shift

As with magnetic fields, one can somewhat easily apply modest electric fields in the lab using charged electrodes. The biggest draw of working with a polar molecule is its polar nature, i.e. the ability of the molecule to respond to an applied electric field. For instance, being able to tap into the dipole-dipole interaction, which scales as $\propto 1/r^3$ instead of the van der Waals interaction, which scales as $\propto 1/r^6$ makes molecules very interesting for quantum simulation applications. One can also apply an electric field that polarizes the molecule, and engineers a repulsive force between the molecules in a trap by tuning the C_6 coefficient, which in turn allows one to shield these molecules from inelastic collisions that would kick them out of the trap [80–82]. As with magnetic fields, electric fields can be bad - for example, electric fields can mix states of opposite parity, which can cause problems when trying to scatter photons. In the case of the RF MOT, because of the fast switching of large currents, the residual field generated in the center of the trap caused the mixing of opposite parity states in the $A^2\Pi_{1/2}$ manifold, which in turned caused molecules to decay to the $N = 0, 2$ states [59].

It is thus also important to understand the effect of an applied electric field on the molecules. This is given by the Stark Hamiltonian:

$$H_S = -\mathbf{d} \cdot \mathbf{E} \tag{2.38}$$

where $\mathbf{d} = e\mathbf{r}$ is the dipole operator, and $\mathbf{E}$ is the applied electric field. It is worthwhile to do this exercise analytically for a simple case to build intuition. Let us consider a set of rotational states $|\Omega, J, M_J\rangle$. For a field along the z direction, the matrix elements of the dipole operator $T^1_p(d)$ in the Hund's case a basis are given by [77]:

$$\begin{aligned} &- \langle \Omega, J, M_J | \mathcal{E} T^1_0(d) \mathcal{D}^1_{00}(\omega)^* | \Omega', J', M'_J \rangle \\ &\qquad = -\mathcal{E} d (-1)^{M_J - \Omega} \sqrt{(2J+1)(2J'+1)} \begin{pmatrix} J & 1 & J' \\ -\Omega & 0 & \Omega \end{pmatrix} \begin{pmatrix} J & 1 & J' \\ -M_J & 0 & M'_J \end{pmatrix} \end{aligned} \tag{2.39}$$

This requires that $M_J = M'_J$ and $J' = J \pm 1$, i.e. it couples together rotational states of opposite parity (here J is assumed to be integer). Thus, there is no first order energy shift, i.e. the Stark shift is a quadratic effect. In particular, for $J = 0$, i.e for a rotational ground state, we have

$$H_S = -\frac{d\mathcal{E}}{\sqrt{3}} \tag{2.40}$$

which leads to a second order energy shift:

$$\Delta E^{(2)} = -\frac{d^2 \mathcal{E}^2}{6B} \tag{2.41}$$

where $2B$ is the rotational energy difference between the $J = 0$ and $J = 1$ states. This also tells us that at low fields, the polarizability of the molecule is actually proportional to the electric field. Thus the electric field does two things - it mixes adjacent rotational states of opposite parity, and as a result, it induces a dipole moment along the direction of the electric field.

The situation gets more complicated when we involve the electron and nuclear spins, where

it is necessary to account for the full Hamiltonian including the Stark term. However, we have calculated the dipole matrix elements already, and the bare Hamiltonian terms are available in [75, 77], and we use this here. Here it is useful to look at two cases, first at low electric fields, where F, m_F are still somewhat good quantum numbers. This is shown in Figure 2.3 for the hyperfine manifold in the $N = 1$ rotational level. Here, for modest fields of up to 2 kV/cm, one can manipulate the different hyperfine sublevels with an electric field.

As the field is ramped up, the molecules lose their hyperfine character and the energy shift is almost entirely dominated by the rotational levels. Because the Stark term couples $N \to N' = N \pm 1$, one needs to account for as many rotational states as possible to get the full effect. This is shown for the rotational states up to $N = 2$ in Figure 2.4, where the calculation accounts for states up to $N' = 20$. For molecules like SrF, with largely spaced rotational levels, a large field is needed to polarize the molecules, and it is very hard to fully polarize them. However, EDM sensitive molecules have closely spaced opposite parity states called Λ or Ω doublets, which can fully polarize the molecules even at modest fields [78].

There are some other interesting things that can be seen from Figure 2.4. The $|N = 0, m_N = 0\rangle$ state is pushed down in energy since it only couples to the $N = 1, m_N = 0$ state, which is higher in energy. Similarly, the $|N = 1, m_N = \pm 1\rangle$ states are pushed down in energy since they only couple to the $|N = 2, m_N = \pm 1\rangle$ states, which are higher in energy. However, the $|N = 1, m_N = 0\rangle$ state is initially pushed higher in energy because of the closer $|N = 0, m_N = 0\rangle$ state, before it is eventually pushed down when the $|N = 0, m_N = 0\rangle$ state gets further away than the $|N = 2, m_N = 0\rangle$. This has the effect that, for small fields, the $|N = 1, m_N = 0\rangle$ state has a negative dipole moment, i.e. it is anti-aligned with the electric field.

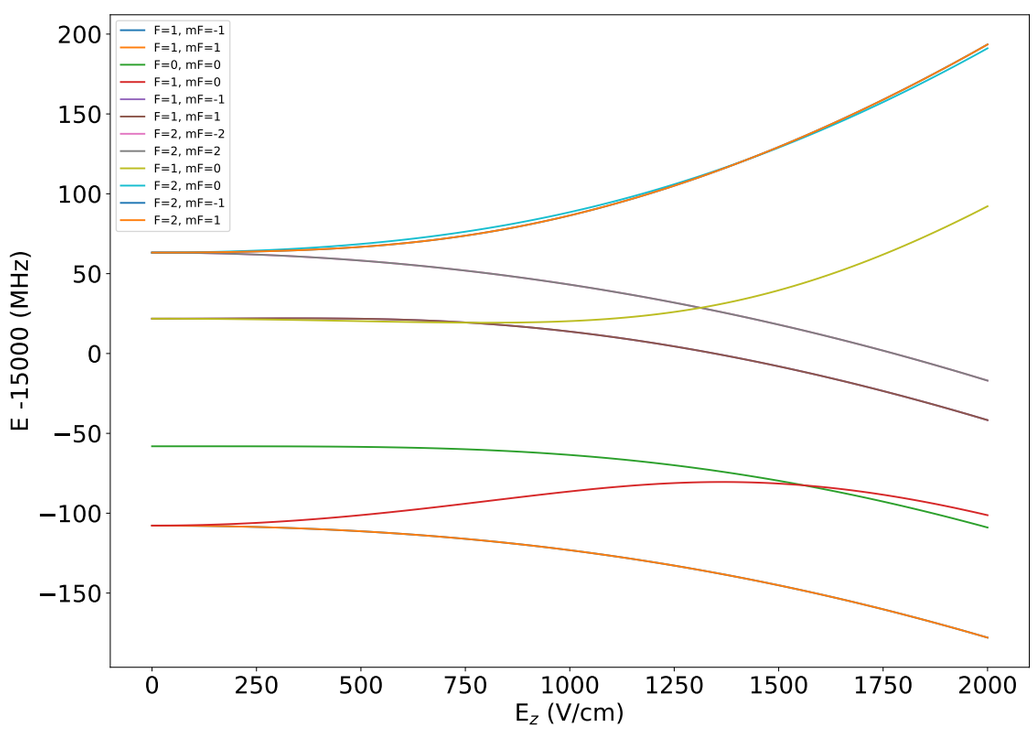

Figure 2.3: Energy shifts in the presence of a small E-field in the $N = 1$ hyperfine manifold. For modest fields of up to 2 kV/cm, the different hyperfine sublevels can be manipulated.

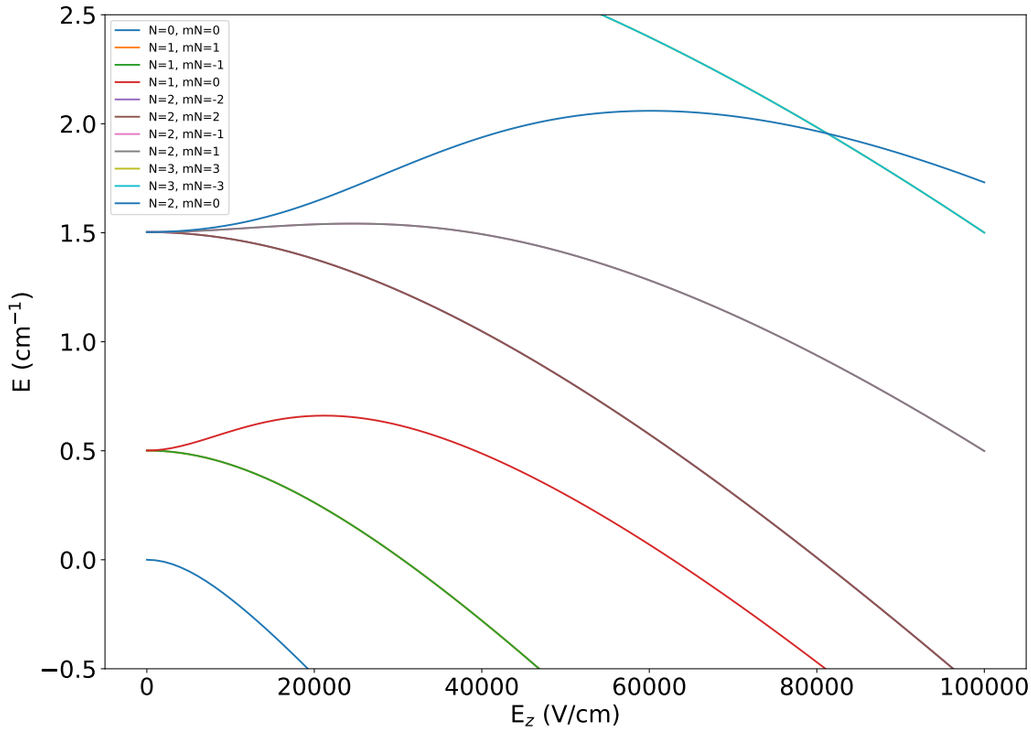

Figure 2.4: Energy shifts for large E-fields. The $N = 0$ state is pushed down by the $N = 1$ state. However, the $N = 1, m_N = 0$ state is actually pushed up initially by the closer $N = 0, m_N = 0$ state, inducing an opposite dipole moment.

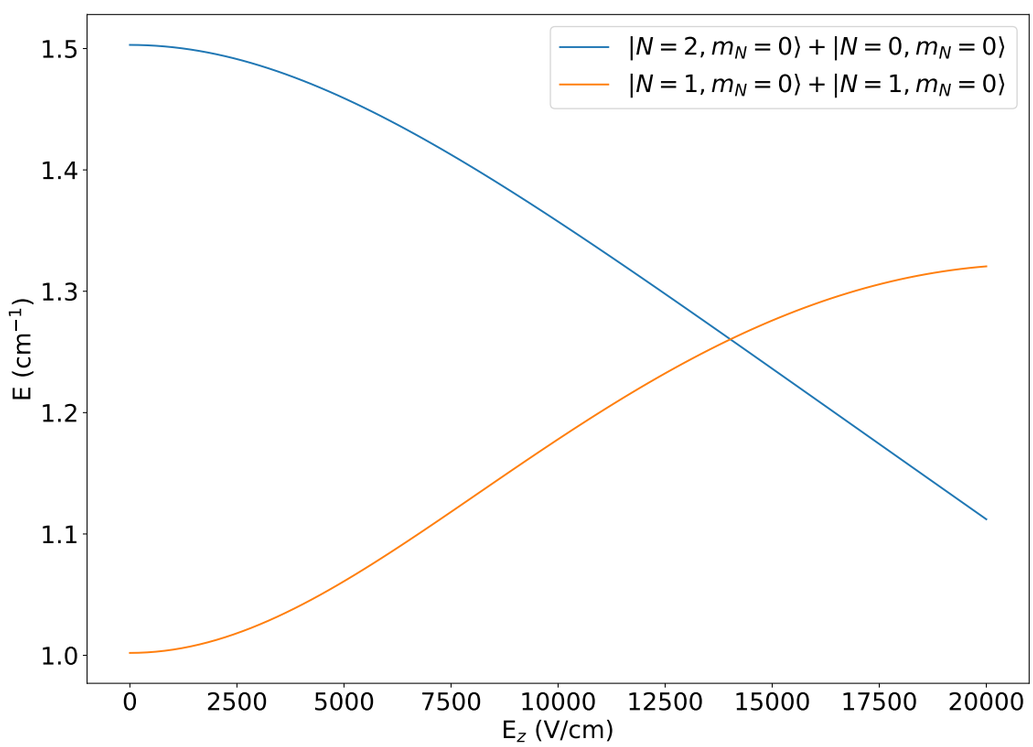

Figure 2.5: Energy shifts of two colliding molecules in the presence of an E-field. Two molecules initially prepared in the $|N = 1, m_N = 0\rangle$ state, have a crossing with the neighboring $|N = 2, m_N = 0\rangle$ and $|N = 0, m_N = 0\rangle$ states at an electric field ~ 15 kV/cm. This can be used to tune the scattering behavior of the molecules, and make the potential attractive or repulsive using the electric field.

The anti-alignment of the dipole for the $|N = 1, m_N = 0\rangle$ has another important ramification for the case of two colliding molecules. As shown in Figure 2.5, the energy of a pair of molecules each prepared in the $|N = 1, m_N = 0\rangle$ state, matches the energy of a pair of molecules where one is in the $|N = 2, m_N = 0\rangle$ and the other is in the $|N = 0, m_N = 0\rangle$ state, when an applied electric field is around 15 kV/cm. Around this crossing then, one can use the electric field to tune the nature

of the van der Waals interaction between the $|N = 1, m_N = 0\rangle$ molecules. Below the crossing, the interaction is attractive, but above the crossing, the interaction becomes repulsive [82]. This has been proposed to lead to a drastic reduction in inelastic scattering processes for molecules [82], and was recently demonstrated as well [83].

2.7 AC Stark shifts

Light is useful to not only drive transitions between electronic states, but to also move the energy levels of the molecule around. These energy shifts due to the molecule-light interaction are known as AC Stark shifts due to the AC nature of the light field. The AC Stark shift is an important tool, since that constitutes the basis of optical dipole trapping (ODT) [84], which is really a pre-requisite to making BECs, and which will be used to trap SrF molecules in Chapter 5. Furthermore, for highly precise experiments such as atomic clocks, controlling and accounting for these light shifts is critical [85–87]. These shifts can also be used to implement quantum gates [88, 89]. Thus, it is vitally important to understand the nature of these shifts, as well as their magnitude. We develop the general formalism to calculate these shifts below.

The effect of far detuned laser light can be treated via an interaction Hamiltonian of the form:

$$H_1 = -\vec{\mu} \cdot \vec{\mathcal{E}}_L(\omega, \vec{r}) \tag{2.42}$$

where $\vec{\mu} = -e\vec{r}$ is the dipole moment operator and $\vec{\mathcal{E}}_L$ is the laser light field. As a general result of second-order time-independent perturbation theory for non-degenerate states, this interaction leads to an energy shift of the i^{th} state (unperturbed energy $\mathcal{E}_i$) that is given by:

$$\Delta E_i = \sum_{j \neq i} \frac{|\langle j|H_1|i\rangle|^2}{\mathcal{E}_i - \mathcal{E}_j} \tag{2.43}$$

When we consider an atom interacting with laser light, we need to treat the energy to be the combined energy of the atom and the laser field. Thus, in the ground state, the atom has zero energy while the laser field has energy $n\hbar\omega$ for n photons, leading to a total energy $\mathcal{E}_i = n\hbar\omega$. When the atom is put into an excited state by absorbing a photon, the sum of its internal energy $\hbar\omega_0$ and the field energy $(n-1)\hbar\omega$ becomes $\mathcal{E}_j = \hbar\omega_0 + (n-1)\hbar\omega = -\hbar\Delta_{ij} + n\hbar\omega$. Hence, we have $\mathcal{E}_i - \mathcal{E}_j = \hbar\Delta_{ij}$. So, for a two level system, interacting with a field of the form $\vec{\mathcal{E}}(\omega, \vec{r}) = \vec{\mathcal{E}}_L(\vec{r})e^{-i\omega t} + c.c$, we have that

$$\Delta E = \pm \frac{|\langle e|\mu|g\rangle|^2}{\Delta} |\mathcal{E}_L|^2 = \pm \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma_{eg}}{\Delta} I \tag{2.44}$$

where the $\pm$ is for the ground and excited states respectively. Here Γ_{eg} is the excited state linewidth for the state $|e\rangle$ to decay to the state $|g\rangle$ and is given by

$$\Gamma_{eg} = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} |\langle e|\mu|g\rangle|^2 \tag{2.45}$$

So we see that for $\Delta < 0$ the ground state is shifted down in energy while the excited state is shifted up in energy, and we can exploit this to trap atoms.

Doing this same calculation for multiple states and using time-dependent perturbation yields a similar result, and we have summarized the result below. For a laser field of the form $\vec{E} =$

$\mathcal{E}_0\hat{z}\cos(\omega t)$, one has the expression for the AC Stark shift for a state n:

$$\Delta E_{\text{AC Stark,n}} = \frac{e^2\mathcal{E}_0^2}{4}\sum_{s\neq n} |\langle z\rangle_{ns}|^2 \frac{1}{\hbar}\left(\frac{1}{\omega_{ns}-\omega}+\frac{1}{\omega_{ns}+\omega}\right) \tag{2.46}$$

Relating this to the light intensity (I_0), one gets:

$$\Delta E_{\text{AC Stark,n}} = \frac{e^2 I_0}{2c\hbar\epsilon_0}\sum_{s\neq n} |\langle z\rangle_{ns}|^2 \left(\frac{1}{\omega_{ns}-\omega}+\frac{1}{\omega_{ns}+\omega}\right) \tag{2.47}$$

This energy shift is precisely what constitutes the trap for atoms and molecules when the light is red detuned and hence this energy shift is conventionally called as the trap depth. The main challenge lies in calculating the matrix elements and in determining the energy shifts accurately. We have already developed the formalism for calculating the matrix elements in § 2.4 and the spectroscopic data for SrF is available [75, 76, 79, 90, 91]. We have developed code to calculate the AC Stark shifts [92] for SrF and the results will be presented in Chapter 5 where they are most relevant.

Chapter 3

Apparatus overview

3.1 Introduction

This chapter details the experiment apparatus as used for the results in the rest of this thesis. Figure 3.1 shows the CAD of the current apparatus used in the experiment. A beam of SrF molecules is produced in the source region, using a cryogenic buffer gas beam source (CBGB). These molecules are then slowed using white-light slowing over a length of around 1.4 m. At the end of the slowing region, the molecules are captured in a magneto-optical trap (MOT), where further experiments are performed.

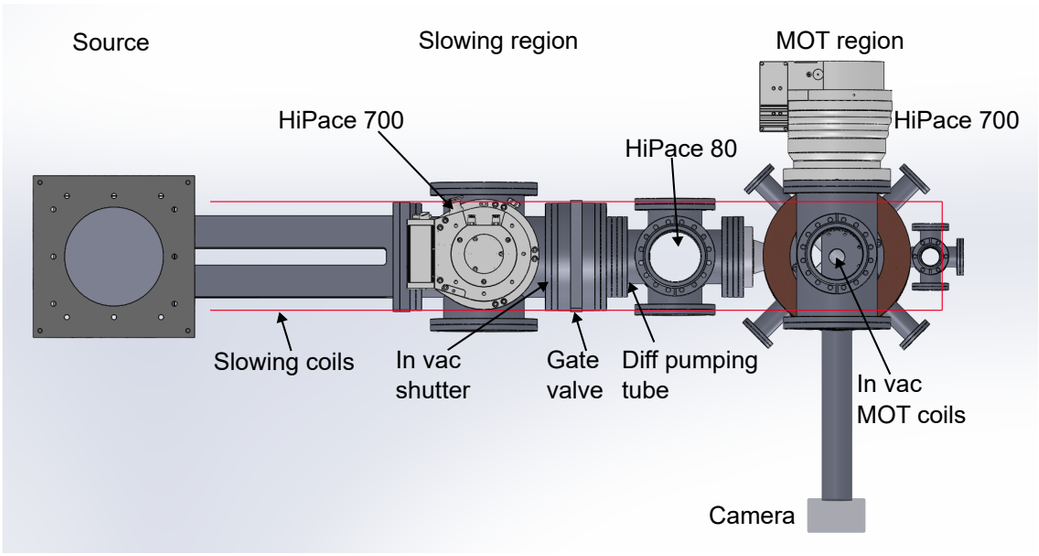

Figure 3.1: Current apparatus used in the experiment. The molecules are produces in the source, then slowed in the slowing region, and finally captured in the MOT region. The different turbos, valves and magnetic field coils used in the experiment are indicated.

3.2 Vacuum system

The vacuum system is divided into three parts - the source region (everything to the left of the gate valve), the UHV region (everything to the right of the gate valve), and the backing region. A lot of the details of this system can be found in previous theses [75, 76, 93] and we summarize the important parts here, as well as any changes made to the system.

3.2.1 Source region

The first component in the source region is the molecule source, i.e. the CBGB. This consists of a room temperature shield which houses a pulse tube refrigerator (Cryomech PT415). The pulse tube has two stages, the first of which is held at ≈ 30 K and the second is held at ≈ 4 K. An Al heat shield is attached to the first stage, and a OFHC Cu heat shield is attached to the second stage. Inside the second stage is the cell which holds a target that is laser ablated to produce the molecules. The initial design used sintered SrF_2 targets, while the new design uses solid Sr targets, see Chapter 6. Cold He is flowed in from the back of the cell to cool the ablated products and extract them from the cell. All the shields, as well as the cell, have windows in the back and holes in front that allow a laser beam to pass through in both directions. The windows are held in place using Indium seals.

The CBGB is followed by a long rectangular slowing region. This rectangular region was previously used for a transverse cooling attempt, but is currently empty. After the rectangular region, there is a 6-way 8" cross, to which a HiPace 700 turbo is attached that pumps on the source region. The other end of the 8" cross is connected to the gate valve that separates the source from the UHV region. The square block on top of the pulse tube has a pirani gauge and an ion gauge attached to it which monitors the pressure. Further downstream, a glass ion gauge is attached to the bottom of the 8" cross, but this is rarely monitored. We further use an in-vacuum shutter (Uniblitz VS14S1S0L-ECE-HS) in the source region to block ballistic He from entering the UHV region and reducing the lifetime of the trap. We have used two iterations, the first one had the shutter right at the beginning of the rectangular block [93]. Here, there were still some gaps around the shutter, and our guess was that Helium could potentially flow around it. The current iteration has the shutter attached immediately before the gate valve, on a solid flange with a central bore. However, we have not noticed any difference in performance between these two iterations.

The flanges in the source region are a mix of O-ring seals and KF flanges, which limits how low the pressure gets at room temperature - normally, the pressure is $\approx 5 \times 10^{-7}$ Torr. However, under normal operation, the source is at 4 K, and thus acts as a cryopump, bringing the pressure to as low as $\approx 10^{-8}$ Torr. With He flowing, this pressure can increase up to $\approx 5 \times 10^{-7}$ Torr.

3.2.2 UHV region

The first component in the UHV region is the 6-way 6" cross immediately after the gate valve. In earlier iterations of the experiment, this was used for the slowing detection, with an in-vacuum spherical mirror and a PMT for light collection. Currently, this chamber is empty, except for the 0.5" ID, 5" long OFHC differential pumping tube that is in place immediately after the gate valve. A small HiPace 80 turbo is connected to the bottom of this chamber for some additional pumping speed.

The other end of this chamber has a zero-length 6"-4.5" reducer and a 4.5"-2.75" conical reducer that connects this chamber to the MOT chamber (Kimball MCF600-SphOct-F2C8). The MOT chamber has 8 2.75" flanges, of which 4 are used for the MOT beam, 2 constitute the beam line, and the final 2 are for imaging, see Figure 3.1. The MOT optics are at a 45° angle to the beam line. For these flanges, we use additional 2.75" full nipples to extend the distance of the viewports from the center in order to reduce scattered light. Additionally, these nipples have blackened inserts in

them to further reduce scatter. The viewports themselves are home made, and consist of a high quality optical window (10-5 scratch-dig, anti-reflection V-coating at 663 nm, CVI PW1-2037-UV) set on a CF 2.75" - KF40 adapter (Kurt J. Lesker F0275XQF40) with UHV epoxy (Epotek 302-3M), see [75] for construction details. One of the two imaging flanges houses the camera, and the opposite end is currently blanked off, with a blackened solid copper gasket to reduce scattered light. The top of the MOT chamber has a 4-way 8" cross - this houses another HiPace 700 turbo, an ion gauge, and the up arm of the MOT optics. The bottom of the MOT chamber has a multiport flange from MDC (MDC 409006) that has a 2.75" port in the center along with 6 1.33" ports. The 2.75" port constitutes the down arm of the MOT optics, whereas the other 1.33" ports are used for making thermal and electrical connections for the in-vacuum coils, see [76] for construction details.

The other end of the beam line port is connected to a 6-way 2.75" cross. This is used for the slowing detection scheme as outlined in [76]. The two horizontal arms of this chamber have standard 2.75" viewports, and the up-down arm houses a PMT that collects any scattered photons (there are no focusing optics in this chamber). The other end of this chamber also has a 2.75" viewport that is used for the slowing light.

All the parts in the UHV region are separately pre-baked for Hydrogen, and then water baked after assembly at around 70-80°C. The temperature is limited because of the Indium solder used in the in-vacuum coils - it melts at around 140°C, which would be disastrous. Future iterations can change the design to instead use mechanical seals, in which case the baking temperature would be limited by other things like the epoxy breakdown point or the temperature rating of the viewports. Under normal operating conditions, the UHV pressure is $\approx 5 \times 10^{-10}$ Torr with no Helium, and rises to $\approx 3 \times 10^{-9}$ Torr with He flowing. If the in-vacuum shutter is used, the pressure only rises to $\approx 1 \times 10^{-9}$ Torr with Helium flowing.

3.2.3 Backing region

The need to flow Helium constantly also creates the need for active turbo-pumping, which is why the experiment uses the big HiPace 700 turbos. These turbos need to have a low backing pressure to maintain the UHV pressures needed in the experiment. These turbos typically have compression ratios $\approx 10^7$ for He, so a UHV pressure of 10^{-10} Torr demands a backing pressure of 10^{-3} Torr.

The backing region houses a HiCube 80 Eco pumping station, with a MVP015-2 diaphragm pump. This single turbo is sufficient to back all other turbos in the experiment, and the vacuum connections are shown in Figure 3.2. The backing connections of both the UHV and the source regions have valves that, in combination with the gate valve, allows each region to be completely isolated from any other region, making diagnostics very easy. Under normal conditions, the pressure in the backing region is around 10^{-4} Torr. The slow backing speed does limit the pump down time after a target change, with a typical pump down taking ≈ 30 minutes to reach 10^{-4} Torr in the source region, and over 6-7 hours to reach the final pressure. However, by planning the opening and closing such that the system is pumping overnight, this is barely a hassle.

Furthermore, there are multiple ion gauges spread throughout the experiment to monitor the pressure at different points. A SRS RGA is connected to the backing region to monitor for any leaks. Since any Helium injected into the system must ultimately exit through the backing region, leak checking using this RGA is straightforward.

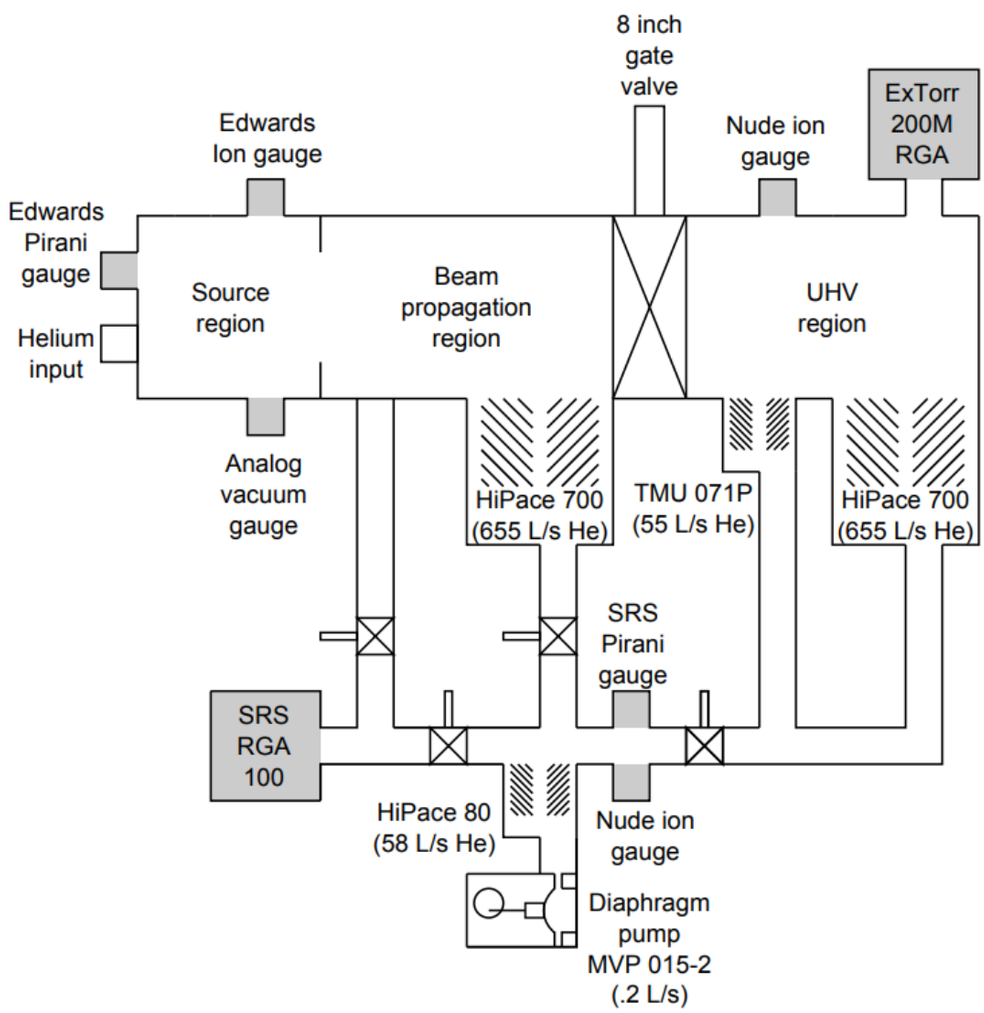

Figure 3.2: Current vacuum system used in the experiment. Figure adapted from [75].

3.3 Lasers and cycling

The optical cycling scheme for SrF is shown in Figure 3.3. The main cycling transition is the $\left|X^2\Sigma_{1/2}, v=0, N=1^-\right\rangle \rightarrow \left|A^2\Pi_{1/2}, v'=0, J=1/2^+\right\rangle$ transition. This transition, however, is not closed, and it has several vibrational decay channels as shown in Figure 3.3(a). This necessitates the use of three additional repump lasers, from the $v=1,2,3$ vibrational manifolds, to continue scattering photons in the MOT.

There are several additional decay channels that need to be addressed to varying degrees. Initial iterations of the MOT saw very short lifetimes, $\tau_{MOT} \lesssim 50$ ms, due to off-resonant excitation to the $\left|A^2\Pi, v'=0, J=3/2^+\right\rangle$ state, which in turn can decay to the $\left|X^2\Sigma, v=0, N=3\right\rangle$, caused by amplified spontaneous emission (ASE) from the tapered amplifiers (TAs) used in the experiment. An additional laser, $\mathcal{L}_{00}^{N=3}$ was added to mitigate this. Due to the fast switching of the magnetic field in the RF MOT, there is an induced electric field at the MOT which causes the $\left|A^2\Pi, v'=0, J=1/2^\pm\right\rangle$ states to mix. The $|J=1/2^-\rangle$ state can decay to the $N=0,2$ states which need to be repumped. An additional laser was introduced, $\mathcal{L}_{00}^{N=2}$ to repump molecules from $N=2$ down to $N=0$, where microwaves were used to transfer population back to $N=1$ and continue cycling.

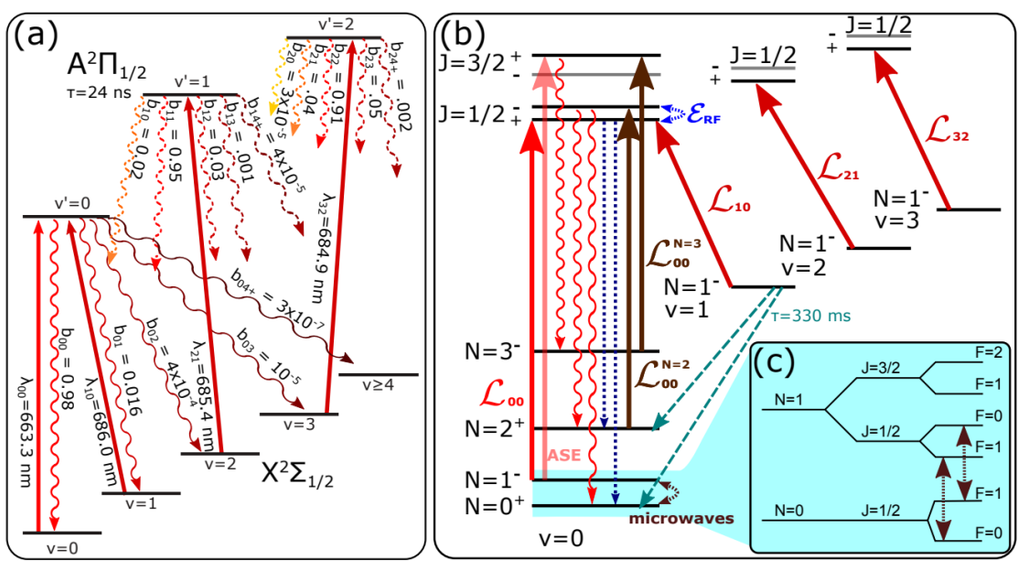

Figure 3.3: Optical cycling scheme used for SrF. (a) The vibrational repumping lasers (solid lines with upward arrows) along with the various vibrational decay channels from the excited states (dashed and solid squiggly lines) are shown. The $v=0 \rightarrow v'=0$ is the cycling transition, and three additional repump lasers are sufficient to achieve vibrational closure for all the experiments in this thesis. (b) Additional loss channels present in different iterations of the MOT, addressed by different rotational repump lasers. Excitation to the $J=3/2^+$ state by ASE from TAs leads to decay into $N=3$, addressed by $\mathcal{L}_{00}^{N=3}$. Electric field induced mixing of the $J=1/2^\pm$ states in the RF MOT leads to decays into the $N=0,2$ states. Decays to $N=2$ are pumped back to the excited state by $\mathcal{L}_{00}^{N=2}$, while decays to $N=0$ are pumped to $N=1$ using microwaves. M1 decays from the cycling state (dashed blue) and vibrational decay from $v=1$ (dashed teal) land molecules in repumped states. (c) Microwave scheme for repumping. Figure adapted from [76].

However, not all of these lasers are necessary to operate the MOT. For most purposes, and while trying to explore new avenues, only the $\mathcal{L}_{00}^{N=3}$ laser is used along with the other vibrational repumps, trading off the slightly lower signal for ease of operation. After the switch to the DC MOT (see Chapter 7), the need for repumping due to electric field induced decays was nullified. A list of lasers currently used in the experiment, as well as their function is detailed Table 3.1. We further note that after the move to UChicago, the laser slow-

ing scheme was changed from the $\left|X^2\Sigma_{1/2}, v=0, N=1^-\right\rangle \rightarrow \left|A^2\Pi_{1/2}, v'=0, J=1/2^+\right\rangle$ to the $\left|X^2\Sigma_{1/2}, v=0, N=1^-\right\rangle \rightarrow \left|B^2\Sigma_{1/2}, v'=0, N=0^+\right\rangle$ transition, as detailed in Chapter 8. This breaks the Λ-system caused in the previous slowing scheme, allowing for a larger scattering rate, thus increasing the MOT signal substantially. We have included this laser in Table 3.1 for completeness. The laser naming convention is $V^x_{v,v'}$ where v, v' are the vibrational states coupled by the laser, and x denotes the function of the laser (M: MOT, s: slowing, $N=3$: repump out of $N=3$, d: laser of opposite polarization used for the MOT, $X \rightarrow B$: address the $X \rightarrow B$ transition, abs: absorption laser outside the cell).

Laser	Frequency (THz)	Current status
1, V_{00}^{M}	451.95729	Was used for RF MOT as MST, replaced by Precilaser
2, V_{10}^{M}	437.04030	Was used for RF MOT as MST, replaced by Precilaser
3, V_{21}^{M}	437.39520	In use as MS
4, V_{32}^{M}	437.74578	In use as M
5, $V_{00}^{N=3}$	451.89950	In use as M
6, V_{00}^{s}	451.95729	Was used for RF MOT as MST, replaced by Precilaser
7, V_{10}^{s}	437.04072	Was used for RF MOT as MST, replaced by Precilaser
8, V_{00}^{M}	451.95729	Precilaser used currently as MST for V_{00}^{d} & V_{00}^{M}
9, V_{00}^{abs}	451.96845	In use as M
10, V_{00}^{PO}	451.96845	In use as M
11, V_{00}^{d}	451.95743	Was used for RF MOT as MST, replaced by Precilaser
12, $V_{10}^{M,s}$	437.04020	Precilaser used currently as MST for V_{10}^{M}, M for V_{10}^{s}
$V_{00}^{X\rightarrow B,s}$	258.77630	MOGLABS 1158 nm seed laser for $X \rightarrow B$ slowing

Table 3.1: All the lasers used in the experiment over the course of this thesis. Status in red indicates the lasers are currently unused, while green indicates in use. Possible configurations for the status are: M: master laser, S: slave laser, T: tapered amplifier.

3.3.1 Sideband generation

Fortuitously for SrF, all the transitions of interest in the $X \rightarrow A$ manifold are at convenient wavelengths (663 nm and 685 nm) where laser diodes were readily available. Thus, until the move to UChicago, we used home-built external cavity diode lasers (ECDLs) in the Littrow configuration to address all the transitions. The design and construction of these can be found in [75, 76]. These diodes are easily tunable over a few nm which is handy to address the rovibrational structure of SrF.

SrF has a well resolved hyperfine structure in the ground state, which means that we somehow need to address these additional states. In the case of SrF, the hyperfine structure spans ~ 170 MHz, and so this is easily achieved through the use of electro-optic modulators (EOMs) that create frequency sidebands on the laser, thus allowing us to address multiple hyperfine states at once [75, 76]. For the repump lasers, where there is no need to precisely control the sideband frequencies/amplitudes, we use home-built resonant EOMs. These EOMs are driven by frequency synthesizers (Novatech 409B) followed by RF amplification. However, for the MOT lasers, where precise frequency and amplitude control is needed to achieve the results presented in this thesis, we use fiber EOMs. To drive these, we use a series of RF switches and Novatechs to allow the changing of the frequency/amplitude of the sidebands in the experiment.

3.3.2 High power generation

The need for high power ($\sim$ 100 mW) and frequency manipulation for each cycling and first vibrational repump laser ($\mathcal{L}_{00}$ and $\mathcal{L}_{10}$) means an ECDL alone is not enough to drive these. We therefore use a three tier system consisting of a master ECDL, a slave diode laser, and a TA to achieve the necessary power and frequencies. The master ECDL is typically operated at $\approx$ 50 mW laser power to maximize the running lifetime. All the frequency manipulation is done on the master - this includes not only the sideband generation, but also any frequency jumps using double pass acousto-optic modulators (AOMs), as needed for the different stages in the experiment. After these manipulations, and accounting for fiber coupling losses, typically 10 mW (less if a fiber EOM is used, $\sim$ 0.5 mW) laser power is available, which is used to seed a slave laser.

The slave laser is chosen such that the free running wavelength is close to the wavelength of interest. Typically, these are the same diodes as the master, although in some cases, may be diodes with higher power rating. As the name implies, the slave laser follows the master laser frequency spectrum, allowing it to be run at higher currents - typical slave lasers are run at $\sim$ 80-90 mW output power. Part of the output of the slave is picked off to monitor the injection, while the rest is fiber coupled and sent to the TAs. For proper injection of the slave, care has to be taken to maximize the forward coupling of the seed light, and the backward coupling of the leaked light through the isolator back into the seed fiber. Only on optimizing these two can the slave be reliably injected. Furthermore, if the slave needs to perform frequency jumps, e.g. in the case of changing from the MOT configuration to the Λ-cooling configuration, then care needs to be taken that the forward coupled seed power does not change too much between the frequencies. These jumps are performed by double-passed AOs which can have slightly different efficiencies for different frequencies - this needs to be accounted for carefully with the corresponding RF power and/or alignment.

After fiber coupling, we are typically left with $\sim$ 40 mW coupled to the input of the TAs. This light is then cleaned up using a telescope and an anamorphic prism pair to mode match the input beam shape of the TAs. The TAs can typically output $\sim$ 350 mW if well coupled. The output of the TA is again beam shaped with a cylindrical lens and a telescope to optimize fiber coupling. With good alignment, $\sim$ 50% fiber coupling efficiency can be achieved. Care must be taken to optimize the backward coupling of the TA light into the seed fiber before running the TA at full power - otherwise the TA may be destroyed. Similarly, care must be taken to optimize the forward coupling of the TA light into the fiber at low power to avoid damage.

Ultimately, we are able to achieve $\sim 100 - 150$ mW laser power coupled into the fiber for each TA used in the experiment.

3.3.3 Cavity locking

The lasers are frequency stabilized using a software monitored transfer cavity locking technique. We use home built Fabry-Perot cavities (free spectral range, FSR, 500 or 750 MHz) enclosed in a KF flange (to prevent air currents) to monitor the spectrum of the laser. The length of these cavities are modulated at 400 Hz using a piezoelectric element. Alongside the lasers of interest, a frequency stabilized HeNe laser (stabilized to $\approx$ 1 MHz) is also monitored as a reference, and the length of the cavity is locked to this HeNe laser by changing the offset of the modulation. The

amplitude of the modulation is set such that two HeNe peaks are visible, which defines the FSR. Then, the laser frequency is determined (modulo the FSR) by comparing the location of the laser peak to the two HeNe peaks, which creates an error signal. This is fed back to the laser's piezo driver, thus locking the laser. Linewidths of ~ 1 MHz are routinely achieved using this technique. Before the move to UChicago, the software was written in LabView, but we have switched to a python version written by Q. Wang. Currently, each instance of the program can lock 3 lasers along with a HeNe, limited by the number of analog outputs available on NI DAQs.

To monitor the absolute frequency of the lasers, a wavemeter (Bristol Instruments 671) is used along with a high speed 16:1 fiber switch. Some of the reject light from each laser of interest is coupled into this fiber switch. A computer program written by T. Langin controls an arduino that in turn controls the logic for the fiber switch, allowing us to choose the laser being monitored. The program can perform statistics on the laser frequency to monitor stability and improve precision if needed. A custom table of pre-defined frequencies allows easy comparison to determine if the laser is off-resonant. This is used to get the lasers to roughly the correct frequency, after which they are locked using the cavities.

3.3.4 Light delivery

The laser light around the experiment is divided into two parts - the slowing light and the MOT light.

For slowing, only three lasers are used - V_{00}^s, V_{10}^s and V_{21}^s. The light for V_{00}^s, V_{10}^s is made using TAs and individually coupled to the slowing breadboard, whereas V_{21}^s is made using only a slave laser. Typically, we achieve $\sim (150, 150, 30)$ mW powers for $(V_{00}^s, V_{10}^s, V_{21}^s)$ respectively. On the breadboard, V_{00}^s and V_{21}^s are combined on a dichroic mirror with the same polarization, and then V_{10}^s is combined with these two using a PBS cube. The beams are then expanded to ≈ 15 mm beam diameter and focused along the beam direction such that the beam shape matched the cell exit aperture downstream. To optimize the alignment of these lasers, the overlap of the beams is checked and optimized on cameras after the first PBS cube and after the beam is focused.

The MOT light geometry is different - for this, all the beams are combined on the MOT optics table, before being fiber coupled and brought to the apparatus table. On the MOT optics table, V_{00}^M and V_{00}^d are combined on a PBS cube. V_{21}^M and V_{32}^M are combined on a 70:30 BS cube, and then combined with V_{10}^M on a PBS cube. The repumps and cycling light are then combined on a dichroic mirror and fiber coupled to the apparatus. Typically, we achieve $(100, 45, 100, 15, 3)$ mW for $(V_{00}^M, V_{00}^d, V_{10}^M, V_{21}^M, V_{32}^M)$ respectively. On the apparatus table, the beam is expanded to ≈ 14 mm diameter using a telescope. The second lens of the telescope is on a translation stage to fine tune the beam convergence, as needed for the results in this thesis. This MOT beam is then circulated over the first three passes of the experiment before it is retro-reflected on itself to constitute the other three passes. This retro mirror is piezo actuated (Newport 8816-6) to allow for fine control and alignment.

3.3.5 Precilasers systems

After moving to UChicago, we decided to overhaul the laser system and make it more modern. To that end, we purchased three precilasers systems capable of producing 2 W each - at 663 nm,

685 nm, and 579 nm. The 663 nm and 685 nm lasers are sum-frequency generators - 663 nm is generated by summing 1018 nm and 1901 nm, whereas 685 nm is generated by summing 1067 nm and 1920 nm. The 579 nm is generated using a 1158 nm seed laser, that is fiber amplified and doubled. All the lasers match their specified laser power and stability.

Each of the red lasers have monitor output fibers that allow us to sample a fraction of the seeds to monitor the frequency and spectrum. We lock the 1000 nm seeds to a cavity that has good Q for both a reference HeNe and this wavelength. The 1900 nm seeds are locked to a cavity that has good Q for a 1550 nm reference laser and this wavelength. The 1550 nm light is picked off from the neighboring Ag-Fr lab. Each red laser also has a fiber EOM on the 1000 nm arm that allows us to generate sidebands on the laser as needed. The 663 nm laser output is split into two arms and is used to seed the MOT cycling lasers, see Chapter 7 for more details. Similarly, the 685 nm laser output is also split into two arms, one seeding the V_{10}^{M} TA, and the other being used for white light slowing instead of V_{10}^{s} TA. Ultimately, the goal is to use the high power from these lasers to implement a 3-beam MOT instead of the single beam MOT we use currently.

The seed light for the 579 nm laser comes from a MOGLABs cat eye ECDL. This has some leaked light that is used to monitor the frequency and the spectrum. We lock this laser to a cavity that has good Q for a reference HeNe and this wavelength. The seed light goes through a fiber EOM before being sent to the fiber amplifier and doubler - the fiber EOM allows us to add sidebands to this laser as needed.

For both modules, care must be taken to ensure there is the right amount of seed light going to the fiber amplifiers before turning up the current. Failure to do this will destroy the amplifier and the laser will need to be shipped back for repairs. Other than this, the lasers are very simple to use, and are essentially turnkey systems.

3.4 Magnetic field generation

3.4.1 Slowing coils

During the slowing stage, molecules are pumped into hyperfine Zeeman sublevels that are dark to the laser polarization. One way to bypass this problem is to periodically switch the polarization of the laser using a pockels cell, as we do in the RF MOT. However, it is a bit inconvenient to do this for all lasers, hence we use a magnetic field of ≈ 5 G oriented at 45° to the slowing polarization to remix the dark states into bright states. This is accomplished with the slowing coils that run along the length of the slowing region. To avoid the slowing field disrupting MOT loading, a pair of anti-coils is also wound near the MOT that cancels the slowing field in this region. The details of the construction can be found in [76]. The slowing coils are controlled using an IGBT to rapidly switch the current on/off. For $X \rightarrow A$ slowing, the coils were turned on for 35 ms, whereas for $X \rightarrow B$ slowing, the coils are turned on for 17 ms.

3.4.2 MOT coils

The MOT quadrupole magnetic field is generated by a set of in-vacuum coils placed inside the UHV chamber. The construction and details of the design are documented in [76], and we summarize the salient features here. The coils have an inductance $L \approx 30\,\mu$H and $R \approx 0.7\ \Omega$, and produce a

gradient $dB_z/dz \approx 2.57$ G/cm/A. The coils are connected in series in vacuum, which restricts us to only produce quadrupole fields. The central bore size has $\phi \approx 1.181$" allowing the MOT beams to pass through.

These coils are designed for both RF and DC operation. For RF operation, the in-vacuum coils are connected to a tank circuit built outside vacuum to impedance match the coils to 50 Ω as needed to transfer maximum power from the amplifiers. Typically, the RF MOT is operated at $f_{MOT} = 1.27$ MHz. A frequency generator produces 1.27 MHz, which is fed into a voltage variable atteunator (VVA) and an RF switch, before being amplified (up to a maximum of 100 W) and sent to the coils. The VVA allows for easy control of the gradient by changing the RF amplitude. For DC operation, the tank circuit is removed, and the coils are connected to the power supply directly. In this case, we use a MOSFET (IXFN230N10) in series with the coil to control the current by controlling the gate voltage. In both cases, a lookup table is made of the magnetic field gradient vs control voltage - this allows for easy creation of magnetic field waveforms in the experiment by only specifying the field one desires.

The coils are connected via thermal feedthroughs to two copper blocks for thermal dissipation. Under normal operation, a thermo-electric cooling (TEC) element is attached to the copper blocks to actively cool the coils and avoid overheating - this is particularly important because the in-vacuum electrical connections for the current set of coils are made using Indium solder, which melts at 140°C.

3.4.3 Shim coils

In most experiments, it is necessary to cancel out the Earth's field at the molecules. This is typically accomplished using 3 sets of shim coils on the 3 axes that create uniform fields at the molecules. Each set of shim coils usually consists of a pair of coils on either end of the chamber that carry currents in the same direction that produces a uniform field, while producing negligible gradients in a small area around the molecules.

In our case, there is only one shim coil along each axis [93]. The three shims are labeled shims 1, 2, and 3 and they produce $(0.228, 0.212, 0.538)$ G/A respectively. This has the problem that it necessarily generates a gradient at the molecules, however, because the molecule cloud is typically small (≈ 1 mm) compared to the variation of the field, it is good enough for most of the results presented here. Earlier iterations of the experiment used a MOSFET (IXFN230N10) in series with each coil to control the shim currents. As with the MOT coils, a lookup table was made of the applied gate voltage vs the field, allowing ease of generating waveforms by specifying the field desired. However, this method has the drawback that there is no feedback mechanism, and if the MOSFET behavior changes, for example due to heating during operation, there is no way to correct for it. Ongoing work is addressing this issue by implementing a feedback circuit.

3.4.4 Magnetic quadrupole coils

A large pair of field coils consisting of a square cross-section wire is wound around the top and bottom of the MOT chamber. In earlier iterations of the experiment, these coils were used for producing the large gradients needed to magnetically trap the molecules. These coils have $R \approx 0.1\Omega$ and $L \approx 1.8$ mH per coil. The coils also have a hollow core to allow for active water cooling. These

coils are currently not in use, and are disconnected from the experiment. See [76] for further details about these coils.

3.5 Imaging

3.5.1 Camera

The main signal in our experiment is light-induced fluorescence (LIF) images of the molecules. To that end, we need an imaging system good enough to resolve the molecule cloud size well. The current imaging system consists of an Apogee Ascent A285 CCD camera, with a 50-mm focal length f/0.95 camera lens attached to it. The light from the molecules is collected by a 2"-diameter, 150-mm focal length spherical singlet lens placed immediately outside the vacuum chamber. An additional one-to-one telescope is placed after this lens, before being attached to the camera lens. An interference filter that only transmits 663 nm light is placed before the light hits the sensor. The one-to-one telescope was installed to move the camera further away from the large MQT coils, as the fast switching of those coils caused problems in the camera.

Currently, the camera has a field of view of 17.75 mm by 13.25 mm, with a pixel size of 12.8 μm per pixel. This is sufficient to resolve the MOT, as well as the long axis of the ODT, but not enough to resolve the transverse axis. Under typical operation, the camera is used in 10x10 binning mode, which means each binned pixel size is 0.128 mm. The different efficiencies of the imaging elements, as well as their construction can be found in [76, 93].

The camera takes an image and then immediately reads it out to the computer - this limits the maximum frame rate the camera can achieve. It was found through manual testing that the maximum fps achievable is around 20 Hz, meaning that there needs to be at least a 50 ms gap between consecutive images. This is somewhat limiting as we would like to, in a single experimental cycle, take an image of the MOT as well as whatever quantity we are measuring, so as to normalize out number fluctuations in the MOT. Currently, a single image is taken each cycle.

3.5.2 LIF based Time-Of-Flight imaging

To measure physical properties (such as temperature, number or density) of the molecular cloud, we use the time-of-flight (TOF) imaging technique. Here, molecules are released from the MOT by simultaneously turning off the magnetic field gradient and the V_{00}^{M} and $V_{00}^{\dagger}$ light. The molecules then freely expand under gravity and due to the temperature of the molecular cloud. After a variable period of free expansion, the molecule cloud is imaged at the maximum light intensity for a very short period - 2 ms if the temperature is higher than $\sim 100\,\mu$K and 1 ms otherwise. During the short exposure, the expansion of the cloud is negligible, and the light does not perturb the physical properties significantly. The LIF emitted during this time is recorded on the CCD camera. For each TOF, a number of images are averaged and then fit to 2D or 1D (only in the case of the ODT) Gaussians to extract the r.m.s width. By tracking the expansion of the cloud versus time, one may derive the temperature of the sample via the relation:

$$\sigma_x^2 = \sigma_{0,x}^2 + \frac{k_B T_x}{m_{\mathrm{SrF}}}\tau^2 \tag{3.1}$$

where σ_x is the r.m.s width of the Gaussian cloud after a TOF of τ, $x = \rho, z$ denotes the radial or axial axis of the image, $\sigma_{0,x}$ is the r.m.s width at $\tau = 0$, k_B is the Boltzmann constant, T_x is the temperature of the cloud and $m_{\mathrm{SrF}} = 107$ amu is the mass of SrF. We are only able to observe the z and one of the radial axes, as a result we assume that both the radial temperatures are the same by symmetry.

These average images allow us to extract different quantities. The temperature is obtained as described above. We can also determine the cloud sizes along all directions by extracting σ_0 from the fits. Together, these allow us to extract the trap motional frequencies by assuming a harmonic potential and with the equipartition theorem. Furthermore, by integrating over the fluorescence counts in the image and by measuring the scattering rate, we can obtain the molecule number - this then allows us to infer the density of the cloud.

3.6 Microwaves

Microwaves are an important tool for controlling the rotational degree of freedom of SrF. The rotational splitting between the $N = 0$ and $N = 1$ states is around 15 GHz - well within reach of conventional microwave synthesizers. Initial designs for the microwave circuit were focused on using the microwaves to repump decays into $N = 0$, or to deplete molecules from $N = 1$ - this did not place any stringent requirements on how pure the microwaves needed to be. This first microwave circuit is described in [76] and it consists of an Agilent E8267C that is mixed with 4 SynthUSB II modules and then amplified to produce the different microwave components as needed, see Figure 3.4.

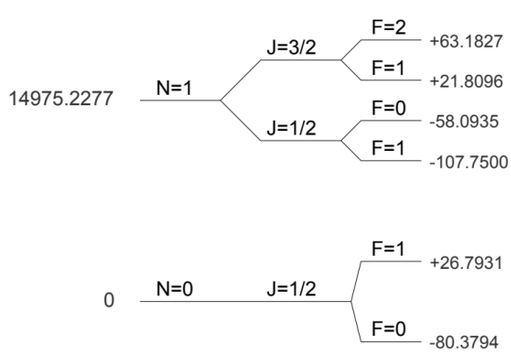

Figure 3.4: Rotational energy levels (in MHz) in the ground state of SrF. Figure adapted from [76].

However, as the experiment developed, and we got to the stage where we now wanted to drive coherent oscillations between the two states, the phase noise of the microwave source started becoming an issue. This is especially true for the case of microwave shielding [80] which places very stringent requirements on the phase noise as well as the polarization of the microwaves inside the chamber. For the Rabi oscillation work described later in the thesis, we stopped mixing the Agilent with the Synths, and instead just used the Agilent. This is enough for the first steps, but eventually better phase noise will be needed, as we describe later.

Chapter 4

Deep laser cooling

4.1 Introduction

There are several important physical limits that one must consider in laser cooling. Typical laser cooling is performed using red-detuning and relies on the Doppler effect to cool [94]. There is a corresponding physical limit, called the Doppler temperature ($T_D = \hbar\Gamma/2k_B = 160\,\mu$K for SrF) where each photon scatter, on average, does nothing. For typical atomic systems (where the excited state angular momentum $F_e > F_g$, the ground state angular momentum in a so called type-I transition), sub-Doppler cooling is achieved using a sisyphus technique, relying on the interplay of Stark shifts and optical pumping caused by the light polarization and intensity gradients. Here, atoms climb up potential hills created by the positive light shift experienced by the bright states (thereby losing kinetic energy), and are optically pumped to the bottom of the hills at the top. Over a number of these events, the atom on average, loses kinetic energy and cools down. This technique faces a corresponding limit called the photon recoil temperature ($T_{\text{recoil}} = \hbar^2k^2/mk_B = 400$ nK for SrF) where one is limited by the recoil energy of the photon. Further cooling is possible using two-photon processes, as we shall see in this chapter. It is important to consider how these limits apply to molecules, and how one may go about cooling them, due to their different level structure.

Rotational closure in molecules requires driving transitions where the ground state angular momentum $F_g \geq F_e$, the excited state angular momentum, i.e. what are called type-II transitions. A consequence of this is that the red-detuned sisyphus mechanism that naturally cools in a type-I transition, actually leads to heating in the system. This effect has been characterized by simulations [65], where red-detuning leads to sub-Doppler heating in type-II systems because of the negative Stark shift seen by the bright states. Under red-detuning, the molecules run down the potential hill (thereby gaining kinetic energy), and are optically pumped to the top of the hills. In order to get sub-Doppler cooling in type-II systems, one has to switch to blue-detuning. The cooling mechanism here is facilitated by the positive light shifts for bright states, which causes molecules to lose kinetic energy as they climb the potential hills when they are in the bright states (Figure 4.1). The polarization gradient and detuning is such that the bright states are most likely to be optically pumped into the dark states at the top of the hill, i.e. at the intensity maxima. The Zeeman composition of the dark and bright states changes as the molecule moves through the polarization gradient. Thus, there is a chance that a non-adiabatic transition to a bright state can occur. This is most likely to happen when the two states are close in energy, i.e. at the intensity minima. Once the molecular state changes to a bright state, the molecule can again climb the

potential hill and undergo this same process many times, thereby losing kinetic energy and cooling down. Using this mechanism, our lab was able to produce SrF molecules cooled down to 50μK (Figure 4.1 (b)), cold enough to load these molecules into a conservative trap [55]. Why this is limited to 50 μK and how to get colder are what we try to answer in this chapter.

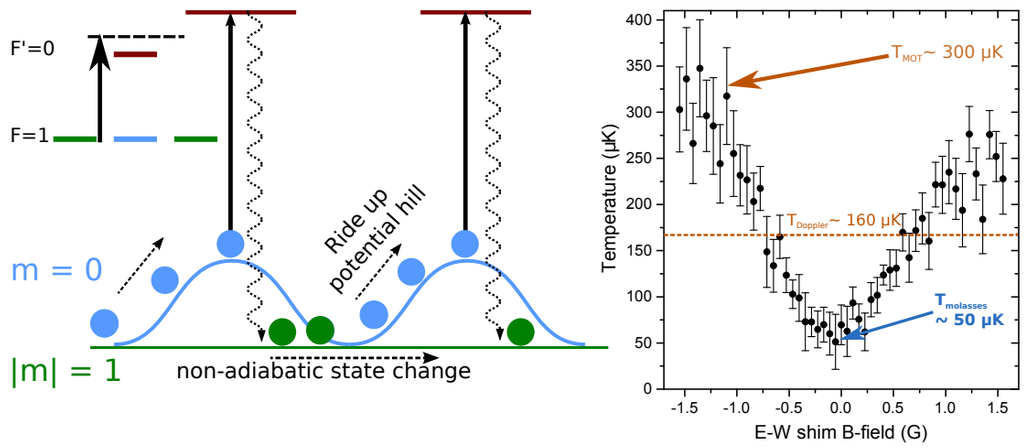

Figure 4.1: (a) The blue-detuned molasses scheme for a $F = 1 \rightarrow F' = 0$ transition. The $|m| = 1$ sublevels are always dark to the light but the blue detuning causes molecules in the $m = 0$ states to lose energy as they climb the potential hill (blue) and are optically pumped to the $|m| = 1$ states at the top, thus losing kinetic energy. (b) Data for sub-doppler cooling of SrF showing 50μK minimum temperature achieved [55]. Figure obtained from M. Steinecker.

4.2 Λ-enhanced gray molasses principle

First, we will build some intuition for the Λ-cooling phenomenon and understand how dark states are actually beneficial for cooling below the Doppler limit. A lot of the treatment will be followed from Refs. [95–97], although there are numerous other references that treat coherences in atomic systems.

4.2.1 Dark states in a simple 3-level system

Let us start with the most basic case, that of a three-level system, where two degenerate ground states are coupled to a single excited state. For such a system, the bare Hamiltonian is

$$H_0 = \hbar\omega_0 \left|e\right\rangle \left\langle e\right| \tag{4.1}$$

where the two ground states are taken to be at zero energy, and the light resonantly couples the two ground states via the Hamiltonian

$$H_1(t) = -\frac{\hbar}{2}\left(\Omega_1 e^{-i\omega_0 t}\left|e\right\rangle\!\left\langle g_1\right| + \Omega_2 e^{-i\omega_0 t}\left|e\right\rangle\!\left\langle g_2\right|\right) + h.c. \tag{4.2}$$

We will treat all the Rabi frequencies to be real, i.e. our Hamiltonian is real. From here, we can use the rotating wave approximation (RWA) to get rid of the fast oscillating terms, to get the

time-independent interaction Hamiltonian:

$$H_{\text{int}} = -\frac{\hbar}{2}\begin{pmatrix} 0 & 0 & \Omega_1 \\ 0 & 0 & \Omega_2 \\ \Omega_1 & \Omega_2 & 0 \end{pmatrix} \tag{4.3}$$

where the basis ordering is $\{|g_1\rangle, |g_2\rangle, |e\rangle\}$. In this case, there is an eigenstate of zero energy, i.e. a stationary eigenstate, which is a superposition of the two ground states, given by

$$|D\rangle = \frac{1}{\sqrt{\Omega_1^2 + \Omega_2^2}}\left[-\Omega_2 |g_1\rangle + \Omega_1 |g_2\rangle\right] \tag{4.4}$$

which does not couple to the excited state at all, i.e. if a atom starts out in this state, it will not scatter any photons. This is the simplest form of a dark state, arising out of coherence between the two ground states.

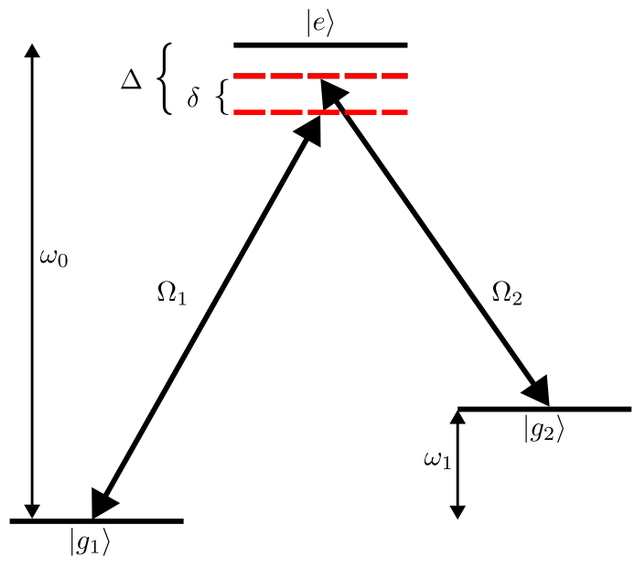

Figure 4.2: Simple 3 level system. Two ground states are coupled to the same excited state. In this case, a dark state arises which is a superposition of the two ground states, when $\delta_R = 0$ is satisfied.

But wait, what if the two ground states are not degenerate? In this case, the bare Hamiltonian is changed to

$$H_0 = \hbar\left(\omega_1 |g_2\rangle\langle g_2| + \omega_0 |e\rangle\langle e|\right) \tag{4.5}$$

as in Figure 4.2. In this case, we can again make use of the RWA to transform into the interaction picture and find

$$H_{\text{int}} = -\frac{\hbar}{2}\begin{pmatrix} 0 & 0 & \Omega_1 \\ 0 & \delta_R & \Omega_2 \\ \Omega_1 & \Omega_2 & \Delta \end{pmatrix} \tag{4.6}$$

This Hamiltonian in general does not have a stationary dark state. However, in the special case of two-photon resonance, i.e. when the Raman detuning $\delta_R = 0$, we again get back our dark state:

$$|D\rangle = \frac{1}{\sqrt{\Omega_1^2 + \Omega_2^2}}\left[-\Omega_2 |g_1\rangle + \Omega_1 |g_2\rangle\right] \tag{4.7}$$

Now, this was all considered for an atom at rest. If our atom is moving in a light field, the two photon detuning will be modified due to the Doppler effect as $\delta_R = (\mathbf{k_1} - \mathbf{k_2}) \cdot \mathbf{v}$. Thus, we can now

set our dark state to appear at any arbitrary velocity by choosing the relative detuning between the two lasers. This is called the velocity selective coherent population trapping (VSCPT) method, because, if an atom at velocity $\mathbf{v}$ interacts with this light field, it will not scatter any photons, and will be "trapped" at that velocity.

4.2.2 Velocity Selective Coherent Population Trapping

The simple three level system we looked at in the previous section helps to build intuition about the origin of VSCPT. However, it is not the complete picture, since it does not tell us why there should be any cooling, it only tells us that there is a dark state at a certain velocity decided by the two-photon detuning. For the complete picture, we also need to consider the momentum eigenstates, and the momentum kick imparted to the atom during each photon scatter. This is beautifully described in [97], which was also the first demonstration of cooling below the photon recoil energy. We briefly recap the treatment here.

Consider a $F = 1 \rightarrow F' = 1$ system under the influence of two counter-propagating beams with σ^+ and σ^- polarizations respectively. In this case, the system is reduced to an effective 3-level system, with the two ground states being $|g_-\rangle = |F = 1, m_F = -1\rangle$ and $|g_+\rangle = |F = 1, m_F = +1\rangle$, and a single excited state $|e\rangle = |F = 1, m_F = 0\rangle$. Now, we also need to consider the atomic momentum operator $\mathbf{P}$, with eigenvalue for $P_z = p$. Then, an excited state with composition $|e, p\rangle$ is only coupled to two ground states $|g_-, p - \hbar k\rangle$ and $|g_+, p + \hbar k\rangle$ via stimulated emission of a σ^+ or σ^- photon with momentum $\pm\hbar k$. Thus, for each value of p, we have a family of three states that are coupled by the laser, with coupling strength:

$$\langle g_\pm, p \pm \hbar k|V|e, p\rangle = \mp\frac{\hbar\Omega}{2}e^{i\omega_L t} \tag{4.8}$$

where the $\mp$ signs come from the Clesbch-Gordon coefficients for the two states. The main thing to note here is that for non-zero p, the two ground states differ in kinetic energy by $2\hbar kp/M$.

Lets first look at the special case of $p = 0$. As before, a dark state exists that is a superposition of the two ground states, given by

$$|D(0)\rangle = \frac{1}{\sqrt{2}}\left[|g_-, -\hbar k\rangle + |g_+, +\hbar k\rangle\right] \tag{4.9}$$

Crucially, this is the only dark state that is stationary. The two ground states both have the same internal energy, as well as the same kinetic energy of $(\hbar k)^2/2M$, and thus an atom pumped here will remain trapped indefinitely. There exist dark and bright states for non-zero momentum as well, given by the same superposition

$$|D(p)\rangle = \frac{1}{\sqrt{2}}\left[|g_-, p - \hbar k\rangle + |g_+, p + \hbar k\rangle\right] \tag{4.10}$$

$$|B(p)\rangle = \frac{1}{\sqrt{2}}\left[|g_-, p - \hbar k\rangle - |g_+, p + \hbar k\rangle\right] \tag{4.11}$$

However, because the two ground states have an energy difference of $2\hbar kp/M$, the atom will oscillate between these two states at a frequency $2kp/M$, and thus the dark state for non-zero p is not trapping. If the lasers are close to resonance, the absorption rate for the bright state is $\Gamma' \approx \Omega^2/\Gamma$ where Γ is the natural linewidth of the transition. The atom takes roughly half a cycle

to precess from the dark to the bright state, i.e. the frequency at which the atom has a chance to absorb a photon from the bright state is $\omega_{abs} \approx kp/M$. For small p, i.e. when $p \ll \Gamma'$, the atom also spends roughly $\tau \approx M/kp$ time in the bright state, and the chance of absorbing a photon during the first cycle itself is high; thus the absorption rate is $\Gamma'' \approx (kp/M)/(\Gamma\tau) = (kp/M)^2/\Gamma'$. This means that for an atom in the dark state, the linewidth is reduced from Γ to $\Gamma'' = \Gamma(kp/M\Omega)^2$. The smaller the momentum, the longer the atom spends in the dark state - precisely the VSCPT mechanism at play.

Adding spontaneous emission is the key for this cooling, since it introduces the dissipation needed to accumulate atoms at lower velocities. During a spontaneous emission event, there is a random momentum kick, anywhere between $\pm\hbar k$ that can change the momentum state of the atom. In this new momentum state, the atom can undergo this process again, and on average lose momentum. The lower the momentum gets, the less the atom scatters due to the lower scattering rate, whereas at higher momenta, the atom has a larger chance of undergoing scattering events. Thus over time, the atoms accumulate in lower velocity states, until they are distributed among the true stationary dark states with $p = \pm\hbar k$ as in eq. 4.9. Furthermore, because the linewidth depends on p and Ω, there is in principle no limit to how low we can cool the atom. It will only be limited by how long one can maintain the coherence between the two states, i.e. on environmental factors that can disrupt the dark state and induce additional scattering. Cooling below the photon recoil limit was demonstrated for the first time using this method in [97], and is a beautiful illustration of the power of dark states.

4.2.3 Λ-enhanced gray molasses

The picture presented above works very well, and is fantastic for atoms with simple level structures. However, for most atoms and molecules, it is not possible to create such a pristine 3-level system. For these cases, one has to rely on combining the robust dark states found in a Λ-system with a traditional gray molasses technique, to form the Λ-enhanced gray molasses scheme. This was first demonstrated for Li and K atoms, where traditional red molasses is hard to implement because of the hyperfine structure [98–100], but has now become a widespread tool in the atomic community because of the low scattering rates, the low excited state population, and the low final temperatures achieved using this method [101–109].

This works as follows. In a typical alkali system, conventional gray molasses cooling was implemented by using a single laser that was blue detuned of one of the hyperfine transitions. This is very efficient up to a certain velocity range, after which point the off-resonant scattering rate serves to more or less cancel out the cooling you can get via the dark states. For a typical blue molasses, the off-resonant dipolar coupling scales as $V_{\text{off}} \simeq \hbar\Gamma(\Gamma/\Delta)I/I_{\text{sat}}$ [99]. This is around $\simeq 50\,\mu$K for SrF molecules, as was demonstrated in the group [55]. Thus, if we can introduce robust dark states in this approach, we should be able to cool much more by reducing the scattering rate. In the first atomic papers, this was accomplished by adding a sideband to the main cooling laser that addressed the lower hyperfine level in the ground state, albeit with a much lower intensity. This additional laser, when tuned to the Raman resonance, enabled the creation of robust dark states composed mainly of the lower hyperfine level, and introduced the strong VSCPT mechanism. This allowed the creation of atomic gases with temperatures as low as $\sim 5\,\mu$K very easily, and with

high capture fraction from the MOTs.

For molecular systems, this is more challenging because of the complex hyperfine manifold in the ground state. For molecules like SrF/CaF, the ground state hyperfine manifold is separated by only a few Γ, thus off-resonant scattering can be an issue if not handled properly. However, in CaF, the Doyle group showed that it is still possible to engineer VSCPT by using two lasers tuned blue of the $|J = 1/2, F = 1\rangle$ and $|J = 3/2, F = 2\rangle$ states [53]. The laser that is blue of the $|F = 2\rangle$ state is also slightly blue of the $|J = 3/2, F = 1\rangle$ state, and provides simultaneous gray molasses. The Raman condition between the two beams produce the VSCPT dark states, and optimal cooling is observed close to Raman resonance, as expected. At large laser intensities, the capture velocity of this mechanism is quite high, and most molecules from the MOT are captured and cooled efficiently. Here, the power ratio between the sidebands is also important, because of the closely spaced hyperfine levels. Ultimately, they were able to cool the CaF molecules to around $T \sim 5\,\mu$K, much colder than the standard gray molasses. In fact, these dark states are so robust, that this cooling mechanism even works in the presence of an optical dipole trapping (ODT) beam. The ODT beam induces AC Stark shifts that can, in principle, ruin the dark states and reduce the robustness. However, by matching the two-photon resonance with the ODT Stark shift, they showed that they can use Λ-cooling to not only load in the ODT, but also take in-situ images of the ODT by collecting the photons scattered during Λ-cooling, which is very encouraging.

In this type of Λ-cooling for molecules, the scattering rate is still somewhat high. The Tarbutt group came up with another scheme, called single-frequency gray molasses, that can engineer even more robust dark states, and further reduce the scattering rate [54]. In this scheme, there is a single frequency of light, blue detuned from all the transitions (closest to the $|J = 1/2, F = 1\rangle$ state) which provides the gray molasses cooling. There is a robust dark state of the complete Hamiltonian in the $|F = 2\rangle$ state, that drastically reduces the scattering rate, and has the potential to cool even deeper. In their paper, they were able to achieve $T \sim 5\,\mu$K, very similar to the Doyle group result. The prospects of loading an ODT with this scheme are a bit less clear, since the $|F = 2\rangle$ state also has a tensor shift, which can ruin the dark states.

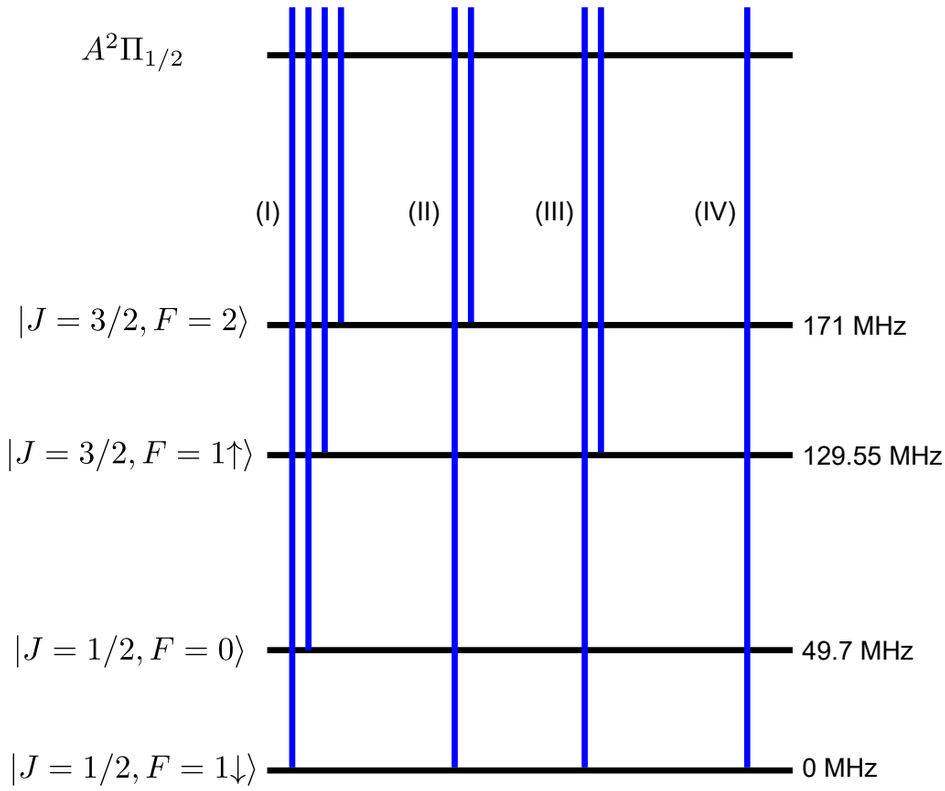

Figure 4.3: Level diagram for different sub-Doppler cooling schemes. (I) Four-frequency molasses, (II) Λ-cooling on $|F = 2\rangle$ and $|F = 1\downarrow\rangle$, (III) Λ-cooling on $|F = 1\uparrow\rangle$ and $|F = 1\downarrow\rangle$, (IV) Single frequency cooling.

In the rest of this chapter, we will detail our attempts at implementing these cooling schemes for SrF (Figure 4.3), and how ultimately we were able to load an ODT using Λ-cooling.

4.3 Experimental setup and results

All the work in this chapter was done with the RF MOT as the starting point of the experiment. Briefly, following laser ablation, molecules are slowed using white light slowing on the $X \rightarrow A$ transition, then captured in a low gradient RF MOT [59]. The MOT is subsequently compressed by simultaneously increasing the gradient and lowering the intensity - this compressed MOT is the starting point of our experiment, see Appendix C. The RF MOT traps around $N \approx 4000$ molecules, at a temperature $T \approx 1$ mK.

Following the RF MOT, the gradient is switched off and the polarization switching is stopped. The light is briefly extinguished, for 0.1 ms, while the trapping laser frequencies are jumped to the molasses frequencies using double-passed AOMs. Simultaneously, the EOMs are also jumped to the cooling configuration using RF switches. In the case of the final Λ-cooling, as will be described later, the V_{00}^M light is extinguished using a physical shutter for the cooling duration. The shim coils are turned on to approximately cancel earth's field during the cooling.

After the cooling is complete, the cooling light is extinguished, the laser frequencies are jumped back to the MOT configuration, polarization switching is turned back on, and a TOF series is taken to determine the temperature. In the cases where any variable is being scanned for optimization, the TOF time is kept the same, usually a large value such as 17 ms, and a series of images for different scan parameters are taken. Following eq. 3.1, the parameter value that corresponds to the least width for a fixed TOF minimizes the temperature. The important variables that can affect the cooling are $B_x, B_y, B_z, \Delta, \delta_R, I$, and t.

4.3.1 Initial attempts using $F = 1 \downarrow$ and $F = 2$

The laser system used to implement the RF MOT and the four-frequency molasses in Ref. [55] is shown in Figure 4.4. The V_{00}^d laser is needed to provide the opposite polarization as needed for the MOT, owing to the opposite g-factor of the $|F = 1 \downarrow\rangle$ state. In the sub-Doppler cooling reported in [55], the V_{00}^M and V_{00}^d masters were simultaneously jumped blue by $\sim +4.2\Gamma$ while the sideband frequency was kept more or less the same to provide the molasses cooling, as shown in the right panel of Figure 4.4. The ability to generate arbitrary sidebands was somewhat limited, as V_{00}^M only had a resonant EOM at 42 MHz in its path, and V_{00}^d did not have any EOMs. Furthermore, V_{00}^d power was limited as we were using high power slave diodes (Mitsubishi LPC-840) that can provide ~ 120 mW power but have a short lifespan.

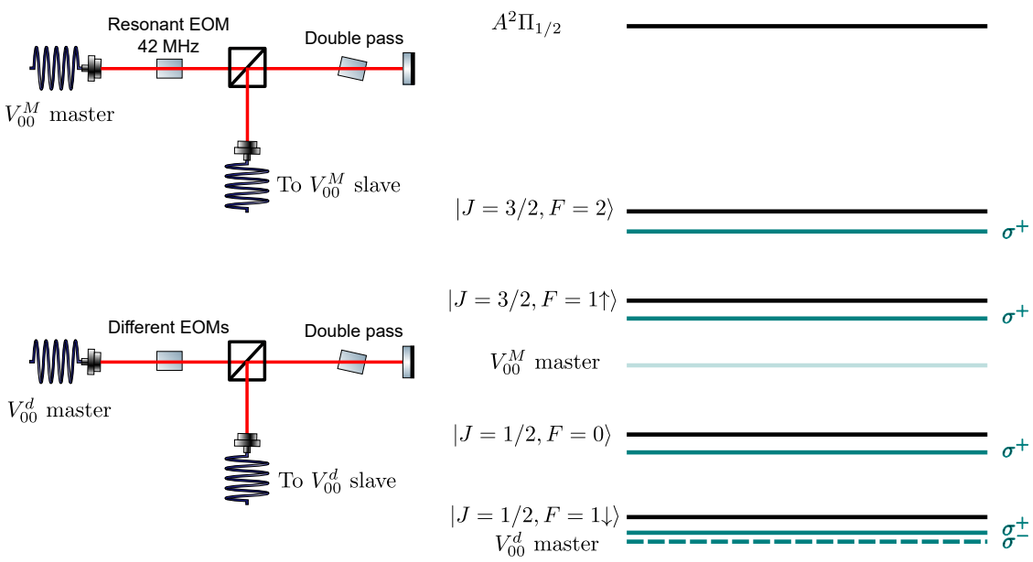

Figure 4.4: (left) Old laser scheme for the MOT and four-frequency molasses. Two separate cycling lasers, V_{00}^M and V_{00}^d with opposite polarizations were used to produce the cooling light. Ultimately, the fiber going to V_{00}^d slave was replaced with a fiber EOM. (right) The different sidebands produced by the laser configuration. V_{00}^M produced four sidebands of σ^+ light, and V_{00}^d produced one sideband of σ^- light as needed for the MOT.

Our first attempt at Λ-cooling was geared towards replicating the CaF result using the $|F = 2\rangle$ and $|F = 1 \downarrow\rangle$ states [53]. With the above restrictions, we first attempted this by jumping V_{00}^M to be blue of $|F = 2\rangle$ and turned off the EOM, while V_{00}^d was kept blue of $|F = 1 \downarrow\rangle$. However, this did not work, and we quickly realized that this was due to the two lasers not being phase coherent, which destroyed the dark states. The light was coming from two different master lasers, making this cooling impossible.

Once it was clear that the sidebands needed to be generated from the same laser, we added a fiber EOM to the V_{00}^d slave setup. This fiber EOM was added after the double pass to add sidebands after any frequency jumps. With this fiber EOM, we were easily able to generate any arbitrary sideband on the V_{00}^d laser. Now, since all the cooling was being done with V_{00}^d, we also added a normally open shutter (Uniblitz LS6T2[1]) after V_{00}^M TA to shut it off during the cooling phase. V_{00}^M still had a resonant EOM in its path generating the four sidebands for the MOT.

With this new setup, we again attempted Λ-cooling using the $|F = 2, 1 \downarrow\rangle$ states, this time by applying a sideband at 171 MHz on the V_{00}^d fiber EOM, and shuttering V_{00}^M. We started out with the Doyle group parameters, namely the ratio between the intensities addressing the

[1]Uniblitz sells normally open and normally closed shutters. We have empirically found that the NO shutters can close and open very quickly, with a minimum total close time ~ 1.5 ms. The NC shutters take around 10 ms to perform an open and close operation.

$|F=2\rangle$ and $|F=1\downarrow\rangle$ states being $R_{2,1} \approx 1$. However, with this sideband ratio, no matter what Δ or δ_R we used, there was always slight heating of the cloud compared to the four-frequency molasses technique. On scanning the sideband ratio, it was found that optimal cooling occurs when $R_{2,1} \sim 0.1$, i.e. basically no power addressing the $|F=2\rangle$ state. For any higher sideband ratio, heating would occur, presumably due to additional scattering from the $|F=2\rangle$ state. At $R_{2,1} = 0.1$ turning the sidebands on/off did not affect the temperature either, implying that it was actually only the laser addressing $|F=1\downarrow\rangle$ that was doing all the work - in other words, single frequency cooling was being implemented instead of Λ-cooling.

This was rather surprising, since prior simulations for SrF had shown that using the $|F=2,1\downarrow\rangle$ states should work, similar to CaF. Figure 4.5 shows Optical Bloch Equation (OBE) simulations of the cooling forces under the influence of Λ-cooling light on the indicated transitions. We will not go into the details of these simulations, as they were designed and implemented by T. Langin and are documented elsewhere [110, 111]. The results of these initial simulations though, were indicative of what cooling forces we could expect, and they showed that Λ-cooling using the $|2,1\downarrow\rangle$ is robust with large cooling forces. This OBE solver was also able to replicate the Doyle and Tarbutt results [53, 54, 67] which made us confident in these results, and we spent a long time trying to get this working. Ultimately however, with this setup, we were able to achieve around $T \sim 40\,\mu$K with $R_{2,1} \sim 0.1$, just barely better than the four frequency molasses.

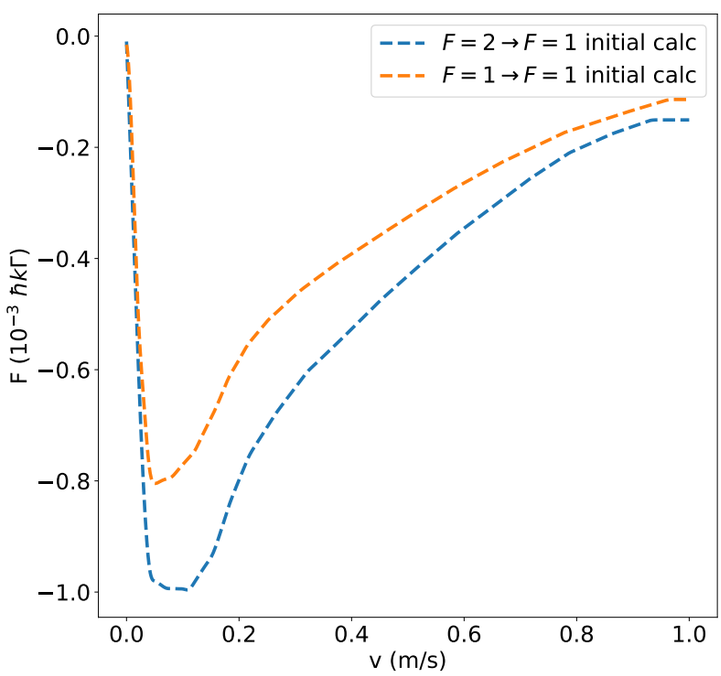

Figure 4.5: Initial OBE simulations for cooling on $|F=2,1\downarrow\rangle$ and $|F=1\uparrow,1\downarrow\rangle$. The cooling force vs velocity is shown, indicating that the force is large for both cooling schemes, which should lead to fast and robust cooling. Plot obtained from T. Langin.

4.3.2 Switch to single frequency cooling and using $|F=1\downarrow, F=1\uparrow\rangle$

With Λ-cooling not working as intended, and with hints of single frequency cooling, we decided to switch to single frequency cooling, since that was easier to understand and implement, with a much smaller parameter space. The only experimental change needed here was to turn off the sidebands on V_{00}^d, which can be done by turning off the RF power to the fiber EOM. The Tarbutt group had demonstrated single frequency cooling at a relatively high intensity of $I \sim 200$ mW/cm^2. They further found that there is a local optimum in cooling near resonance of $|F=1\downarrow\rangle$ and then

there is gentle reduction in temperature at very high detunings [54]. Our first attempt at single frequency cooling was with V_{00}^d at the MOT frequency, which is around -1.4Γ of $|F=1\downarrow\rangle$ but blue of every other state by a large amount - this worked quite well, and we were able to cool down to $T \sim 30\,\mu$K with single frequency cooling applied at maximum power, with a single beam intensity $I \sim 35$ mW/cm^2.

At this point, we were able to scan the one-photon detuning and figure out an optimum, see and found that it was optimized at $\Delta = +0.2\Gamma$, matching the Tarbutt results, see Figure 4.6. This makes sense since the molecules want to spend the most time in the dark states created from $|F=2\rangle$ and thus having a high scattering rate from $|F=1\downarrow\rangle$ enables faster optical pumping into the dark states. The slight blue detuning favors the molasses mechanism. As always, we found that this cooling works best when the shims cancel earth's field, as this preserves the dark states necessary for cooling. We then scanned the dependence of the cooling on the laser intensity, and found that even at the maximum laser intensity that we could get with the high power slave on V_{00}^d, the cooling had not saturated. At this point, it was decided to build a new TA setup for V_{00}^d as the higher power would be required not only for better cooling, but also for trying new MOTs/cooling techniques in the future. We purchased a MOGLABs 670 nm TA system, and calibrated and characterized it - it matched spec and we were now able to increase the intensity to $I \sim 50$ mW/cm^2. This allowed us to saturate the cooling efficiency, and we were able to cool down to $T \approx 10\,\mu$K using single frequency cooling alone. We were not quite sure what was limiting us to this temperature, since in principle we should be able to match the Tarbutt result of $T \approx 5\,\mu$K. However, it was decided to press ahead and try the other Λ-cooling scheme.

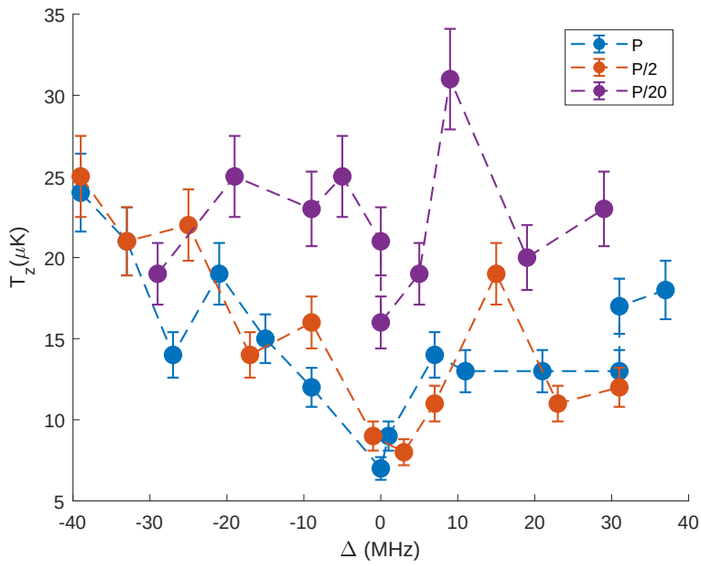

Figure 4.6: Detuning scan for single frequency cooling at different laser powers. Optimal cooling is observed for high laser intensities (here $P = 50$ mW/cm^2) and at $\Delta = +0.2$ Gamma, i.e. very close to resonance with $|F=1\downarrow\rangle$.

With the high power available for V_{00}^d, we revisited Λ-cooling using $|F=2,1\downarrow\rangle$ first to check if we were intensity limited. However, that still led to heating, which confirmed that this combination cannot work. We then switched schemes to attempt Λ-cooling on the $|F=1\uparrow,1\downarrow\rangle$ states - this was again easily accomplished by switching the fiber EOM frequency from 171 MHz to 129.55 MHz. We started out with $R_{1\uparrow,1\downarrow} \approx 0.5$, $\Delta \approx +1\Gamma$ and $\delta_R \approx 0$, with the maximum laser intensity. This was immediately much better than the first scheme, and we were able to obtain $T \approx 35\,\mu$K in our first pass. This was a good starting point to start exploring the parameter space.

We scanned δ_R first and found that the widths are minimized at $\delta_R \approx 0$ as expected, with a slight preference for blue detuning. We also saw that the widths were minimized for $\Delta \approx +1.4\Gamma$, although it is not very sensitive to this, which also makes sense, since ultimately, the dark states depend on δ_R, and having Δ blue detuned only provides the molasses cooling. The cooling is very sensitive to $R_{1\uparrow,1\downarrow}$, with optimal cooling found at $R_{1\uparrow,1\downarrow} \approx 0.2$. This is because the dark state composition depends on the coupling of these two levels to the excited state, and the right balance will ensure robust dark states. With these parameters optimized, we were able to achieve $T \approx 20\,\mu$K, not quite as cold as what the Doyle group demonstrated, but still colder than the other Λ-cooling scheme. We found that the Λ-cooling is much faster, most likely due to the larger scattering rate, and can cool from the MOT starting temperature of $T \approx 1000\,\mu$K to $T \approx 20\,\mu$K within 1.5 ms. This allows us to switch to single-frequency cooling after this point, to cool the rest of the way to $T \approx 10\,\mu$K, see Figure 4.7.

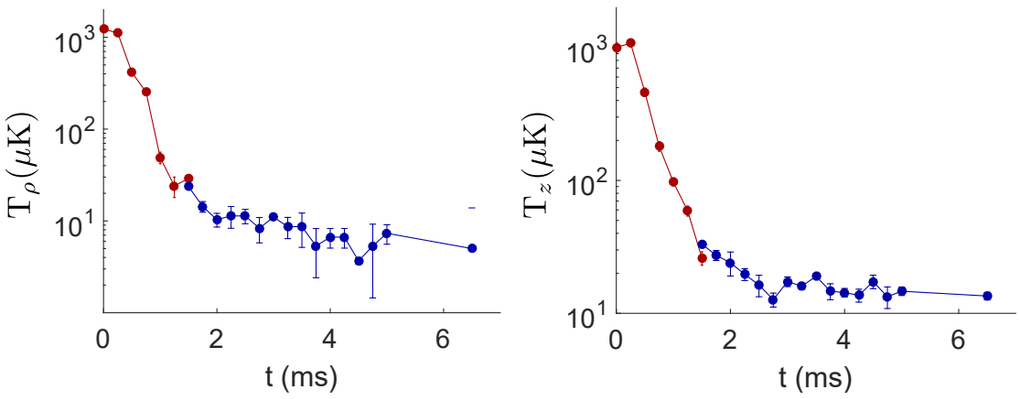

Figure 4.7: Combination of Λ-cooling and single frequency cooling to reach $10\,\mu$K. Points in red are Λ-cooling and in blue are single frequency cooling. (left) The radial temperature vs time and (right) The axial temperature vs time. The MOT temperature is around 1 mK, after which Λ-cooling is applied. This cools down to $\approx 20\,\mu$K after which single frequency cooling takes over and cools to $\approx 10\,\mu$K.

We thought this should be good enough to load an ODT, and so we also ordered a high-power 1064 nm ODT laser. While we were waiting for that laser to arrive, we still wanted to figure out why the Λ-cooling using the $|F = 1\uparrow, 1\downarrow\rangle$ states was limited to $T \approx 20\,\mu$K. Our best guess was that the high scattering rate that allows us to cool quickly from the MOT, may also be preventing us from cooling further. So, we decided to try a low power Λ-cooling stage after the first high power stage, which would result in lowering the scattering rate, and in turn lower the temperature, if our guess was right. We found that by applying an initial high power cooling pulse with maximum intensity for 1.5 ms, and then lowering the intensity to 20% of the maximum intensity resulted in fantastic cooling! Indeed, by applying a second low power stage for 1.5 ms resulted in lowering our temperature to $T \approx 10\,\mu$K, as good as the single frequency case. We again scanned all the Λ-cooling parameters with this approach, and found that they all more or less stayed the same, with the second stage cooling saturating after 1.5 ms. Figure 4.8 shows some of the parameter scans performed during this process, with optimum $\delta_R = +200$ kHz and the shims set to cancel earth's field as expected.

4.3.3 Beam alignment and other issues

It was found that the cooling is highly sensitive to the MOT beam alignment, and in particular that the beam is retro-reflected onto itself perfectly. In our experiment, the 6 MOT beams are made

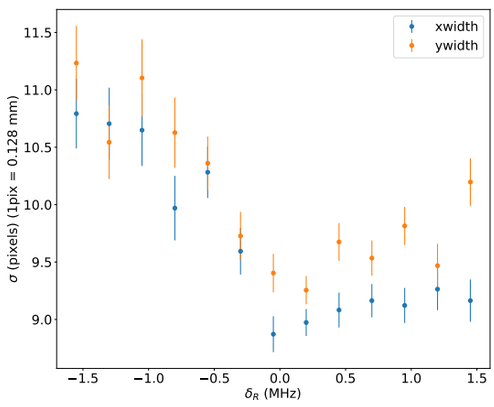

(a) Width vs δ_R

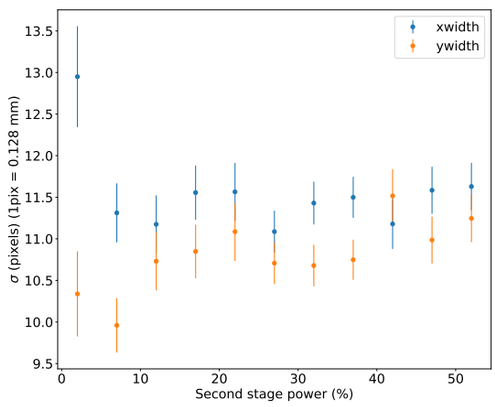

(b) Width vs second stage power

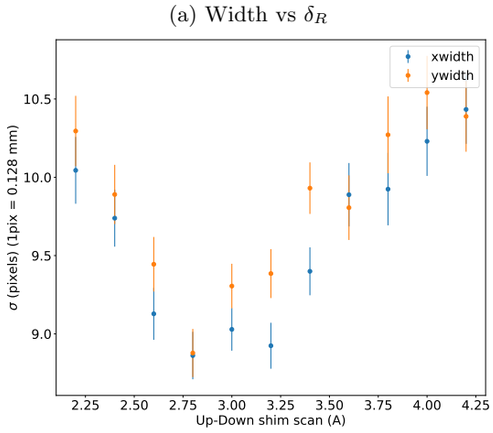

(c) Width vs Up-Down shim current

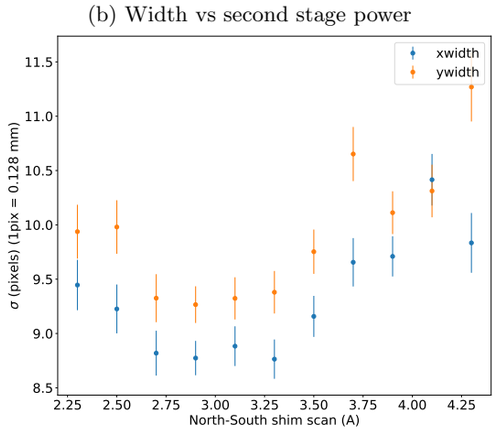

(d) Width vs North-South shim current

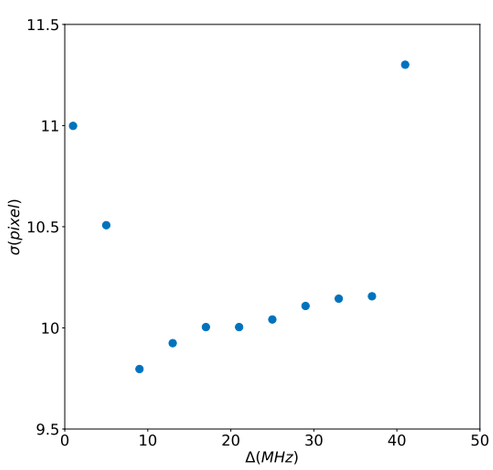

(e) Width vs Δ

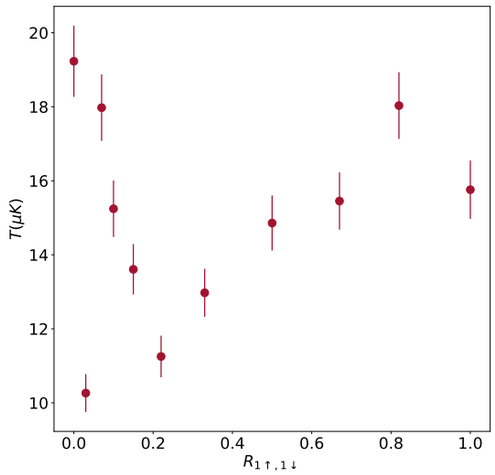

(f) Width vs sideband ratio $R_{1\uparrow,1\downarrow}$

Figure 4.8: Λ-cooling parameter scans. (a) δ_R scan showing optimum at +200 kHz. (b) Second stage power scan with optimum around 20%. (c) Up-Down shim current scan with optimum canceling earth's field. (d) North-South shim current scan with optimum canceling earth's field. (e) Δ scan showing optimum at $\Delta = 10$ MHz i.e. $\approx +1.4\Gamma$. (f) Sideband ratio, $R_{1\uparrow,1\downarrow}$ scan showing optimum sideband ration $R_{1\uparrow,1\downarrow} \approx 0.2$.

from a single beam that gets circulated around the experiment - this is because molecular MOTs and quite power hungry, and historically, sufficient power was not available at these wavelengths economically to make a traditional 3 beam setup. Thus as the beam passes through all six passes of the chamber, it loses a significant amount of power and intensity, and the sharp dependence on the retro alignment is presumably a consequence of finding the right polarization gradients at the center of the beam. To enable fine adjustment of the retro beam, we use a piezo actuated mirror on the final mirror that reduces errors in alignment.

We also saw that the camera exposure time matters a lot for low temperature measurements. In particular, exposing the molecules to resonant light for 2 ms (as was usual for the MOT) led to significant and artificial heating, which would inflate our temperature measurements by a few μK. On reducing the light exposure to 1 ms, most of this went away, and we were able to better resolve the actual cloud temperature.

Finally, we saw that the master laser for V_{00}^{d} had a large linewidth ($\sim$ 3 MHz) owing to a bad cavity. So, we switched the master-slave configuration such that now there was only 1 master laser called V_{00}^{M}. The master was now parked close to $|F = 2\rangle$. The master light was split into two arms, with the V_{00}^{M} arm containing the resonant 42 MHz EOM that generates the sidebands for the MOT. We added a 171 MHz AOM before the V_{00}^{d} double-pass setup to bring this laser close to $|F = 1 \downarrow\rangle$ and then it went through the same double-pass AOM + fiber EOM setup as before. Now, both V_{00}^{M} and V_{00}^{d} were being seeded by the same master laser, which had a narrow linewidth ($\sim$ 1 MHz), and this also helped improve the cooling somewhat. Since V_{00}^{d} now also had a TA in its path, we combined some $V_{00}^{N=3}$ light with the slave before sending it to the TA to reduce the detrimental effects of ASE. This helped improve the SNR of our images.

4.4 An explanation

We were puzzled as to why cooling on the $|F = 2, 1 \downarrow\rangle$ states did not work, but the $|F = 1 \uparrow, 1 \downarrow\rangle$ states worked so well, particularly since the simulations suggested that they should both be good, with the former being slightly better, see Figure 4.5. However, during this period, T. Langin noticed an error in the branching ratios that we had calculated for the $\left|A^2\Pi, J = 1/2^+ \rightarrow X^2\Sigma, N = 1\right\rangle$ transition manifold. This error first appeared in [75] and was the result of a relative sign error in calculating the matrix elements for the $|F = 1 \uparrow\rangle$ and the $|F = 1 \downarrow\rangle$ states. This is a really insidious error, as these are the two states that are mixed by J-mixing (see Table 2.1) and a sign error in any of the matrix element would affect the overall matrix element for either of these states. In particular, it was found that the sign of the matrix elements for $|F = 1 \uparrow\rangle$ should be of the opposite sign to the matrix element for the $|F = 1 \downarrow\rangle$ state. The old branching ratios are shown in Table 4.1 whereas the new and corrected ones are shown in Table 4.2. Of course, how do we know these calculations are correct, and do not have an error themselves? We cross-checked the way we calculate this with one of the Tarbutt group papers [112] that show the branching ratios for their $|F = 1\rangle$ states which includes the J-mixing coefficients for CaF. Indeed, the new branching ratios match this paper, lending further credence to this calculation.

This has the result that the coupling strengths, and thus the cooling forces are drastically different for the two cooling methods than what we had previously simulated. The new and corrected cooling forces are shown in Figure 4.9 where the difference is quite apparent. Now, the

			$F' = 0$	$F' = 1$		
$\tilde{J}$	F	m_F	$m'_F = 0$	$m'_F = -1$	$m'_F = 0$	$m'_F = 1$
3/2	1	−1	0.0063	0.133	0.133	0
3/2	1	0	0.0063	0.133	0	0.133
3/2	1	1	0.0063	0	0.133	0.133
1/2	1	−1	0.3271	0.117	0.117	0
1/2	1	0	0.3271	0.117	0	0.117
1/2	1	1	0.3271	0	0.117	0.117

Table 4.1: Old and incorrect branching ratios for SrF. These have the wrong relative sign for the matrix elements between the $|F = 1 \downarrow\rangle$ and $|F = 1 \uparrow\rangle$ leading to wrong coupling strengths. Somehow, this error was propagated from [75].

			$F' = 0$	$F' = 1$		
$\tilde{J}$	F	m_F	$m'_F = 0$	$m'_F = -1$	$m'_F = 0$	$m'_F = 1$
3/2	1	−1	0.2629	0.0047	0.0047	0
3/2	1	0	0.2629	0.0047	0	0.0047
3/2	1	1	0.2629	0	0.0047	0.0047
1/2	1	−1	0.0704	0.2453	0.2453	0
1/2	1	0	0.0704	0	0.2453	0.2453
1/2	1	1	0.0704	0	0.2453	0.2453

Table 4.2: New and correct branching ratios for SrF. These have the correct relative sign for the matrix elements between the $|F = 1 \downarrow\rangle$ and $|F = 1 \uparrow\rangle$ leading to different coupling strengths.

cooling force on the $|F = 2, 1 \downarrow\rangle$ states is drastically reduced by a factor ~ 10 whereas the cooling force for $|1 \uparrow, 1 \downarrow\rangle$ is barely affected.

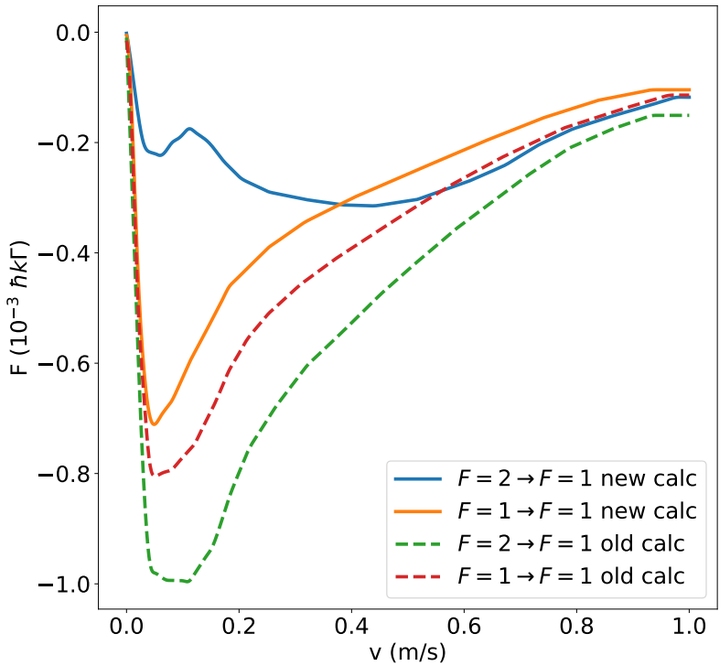

Figure 4.9: New OBE simulations for cooling on $|F = 2, 1 \downarrow\rangle$ and $|F = 1 \uparrow, 1 \downarrow\rangle$. The cooling force vs velocity is shown, indicating that the force is large only for $|F = 1 \uparrow, 1 \downarrow\rangle$, but is very low for $|F = 2, 1 \downarrow\rangle$. Plot obtained from T. Langin.

One might argue that there is still some cooling force for cooling on the $|F = 2, 1 \downarrow\rangle$ states, so why does Λ-cooling not work at all, and if anything, leads to heating? To answer this question, the OBE simulations are not enough, since they do not take into account the actual process of cooling, which involves scattering photons. For the full simulation, we need to turn to a Monte Carlo simulation, also known as a Quantum Trajectories simulation. This takes the previously produced force profile, and also takes into account the spatial profile of the cooling beams, the process of scattering photons and momentum diffusion during this. This was developed by T. Langin and the results are shown in Figure 4.10. We see that the case of CaF and SrF are actually

Cooling scheme	Δ	δ_R (MHz)	R_Λ	I (mW/cm^2)	Best T	Second stage?
Single frequency	$+0.2\Gamma$	N/A	N/A	50	10 μK	Yes/No
$\|F = 2, 1\downarrow\rangle$	$+1.4\Gamma$	$+0.2$	0.1	50	40 μK	N/A
$\|F = 1\downarrow, 1\downarrow\rangle$	$+1.4\Gamma$	$+0.2$	0.2	50	10 μK	Yes

Table 4.3: Optimal cooling parameters for each cooling method.

quite different. For CaF, cooling on the $|F = 2, 1\downarrow\rangle$ as well as the $|1\uparrow, 1\downarrow\rangle$ transitions works, for a wide range of laser intensities. In fact, while the Doyle group paper used the former cooling transition, it appears from the simulations that the latter should work even somewhat better for CaF. The story for SrF is radically different, where cooling on the $|F = 2, 1\downarrow\rangle$ does not work at all, and if left to go for long times, actually leads to heating, as we observed in the experiment. This is because the cooling force is so weak that it can no longer overcome the random scattering that accompanies the cooling process. Indeed, if the sideband addressing $|F = 2\rangle$ is large enough, then the scattering rate is much higher than the cooling rate, and we get runaway "heating". This is also why, when this transition appeared to work for us, it was only when $R_{2,1\downarrow}$ was very small. However, the strong cooling forces for the $|F = 1\downarrow, 1\uparrow\rangle$ allow us to cool very efficiently, as we verified experimentally.

4.5 Final cooling results

Ultimately, we were able to cool down to $T \approx 10\,\mu$K in a few ways: using single frequency cooling alone, high power Λ-cooling (on $|F = 1\uparrow, 1\downarrow\rangle$) with a second stage of single frequency cooling, and with two stages of Λ-cooling (on $|F = 1\uparrow, 1\downarrow\rangle$) alone. Figure 4.11 shows a comparison of the scans for δ_R for the different schemes we tried - no matter the δ_R, $|F = 2, 1\downarrow\rangle$ is always worse, as expected. However, single frequency cooling and Λ-cooling using the $|F = 1\uparrow, 1\downarrow\rangle$ work equally well, and we are able to cool down to $T \approx 10\,\mu$K using both. Figure 4.12 shows a closer comparison between single frequency cooling and Λ-cooling using $|F = 1\uparrow, 1\downarrow\rangle$. For $\delta_R = 200$ kHz, cooling is optimized. The best cooling parameters are listed in Table 4.3.

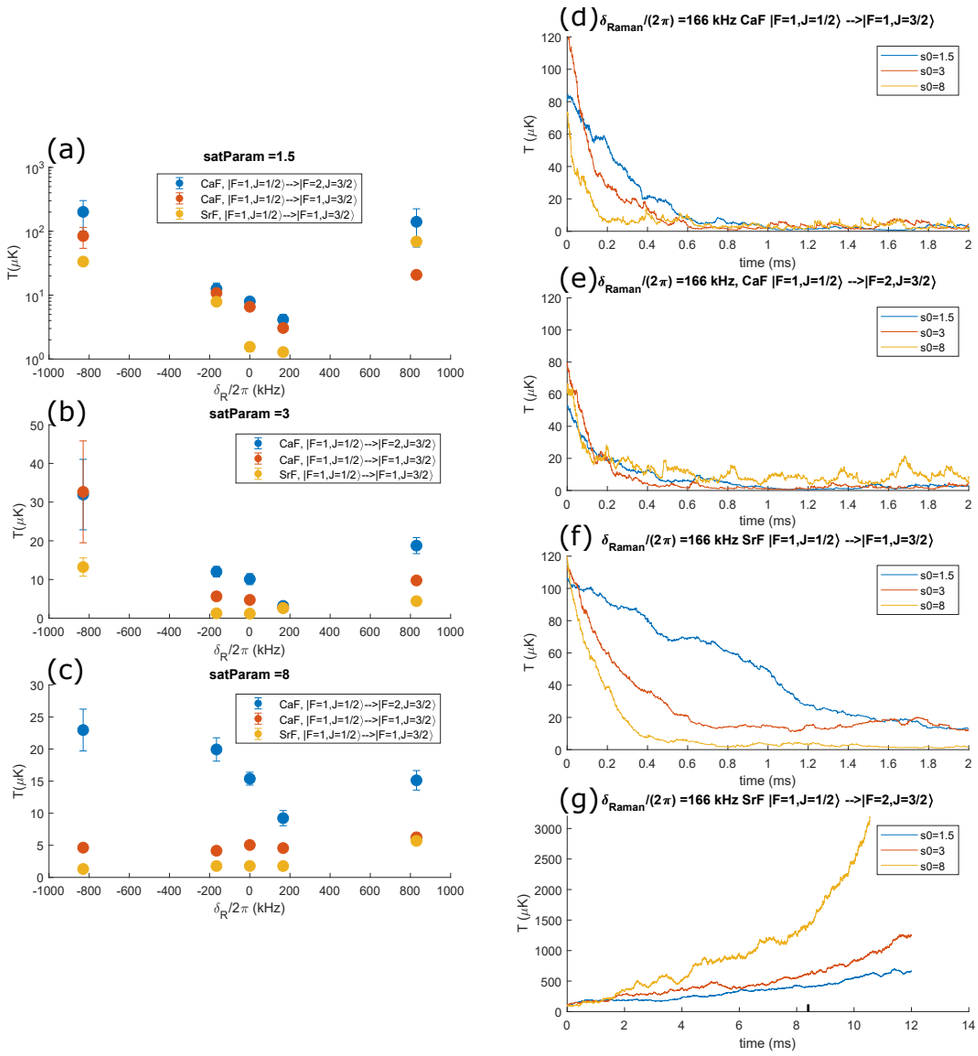

Figure 4.10: Quantum trajectories simulations of cooling forces for different configurations with the corrected branching ratios. **Left panel** (a)-(c) Comparison of cooling forces between CaF and SrF for different laser saturation parameters vs the Raman detuning δ_R. Cooling for CaF using the $|F = 2, 1\downarrow\rangle$ and $|F = 1\uparrow, 1\downarrow\rangle$ states are shown, whereas cooling for SrF using only the $|F = 1\uparrow, 1\downarrow\rangle$ is shown. (a) Saturation parameter = 1.5 shows that the cooling forces for both mechanisms in CaF and SrF are comparable. Note that this is a semilog plot and the range of δ_R over which cooling works is quite small. (b) Saturation parameter = 3 shows that as the power is increased the cooling gets more robust, but has roughly the same effectiveness for all mechanisms. (c) Saturation parameter = 8 shows that at large powers, for CaF, the $|F = 1\uparrow, 1\downarrow\rangle$ is more effective than $|F = 2, 1\downarrow\rangle$, and one can ultimately reach $5\,\mu$K with this method. **Right panel** (d)-(g) Time evolution of temperature for different saturation parameters for a fixed detuning $\delta_R/2\pi = 166$ kHz and cooling transition. (d) CaF cooling on $|F = 1\uparrow, 1\downarrow\rangle$ and (e) CaF cooling on $|F = 2, 1\downarrow\rangle$ shows that both methods are effective to reach $\sim 5-10\,\mu$K within 1 ms of cooling. (f) SrF cooling on $|F = 1\uparrow, 1\downarrow\rangle$ shows a clear power dependence, with fastest cooling at high powers, but ultimately, one can reach $\sim 5-10\,\mu$K within 1.5 ms of cooling. (g) SrF coolong on $|F = 2, 1\downarrow\rangle$ shows that cooling does not work on this transition at all, and in fact, for longer times, the random photon scatter overcomes the small cooling force and leads to heating. This finally explained our experimental results showing no cooling with $|F = 2, 1\downarrow\rangle$. Plot obtained from T. Langin.

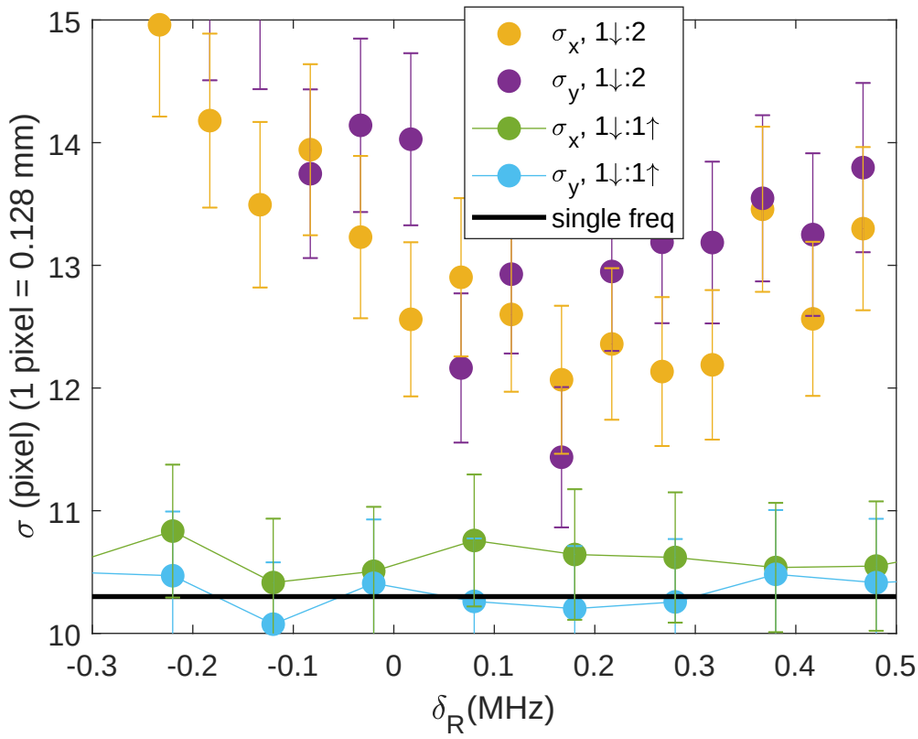

Figure 4.11: Comparison of the different cooling techniques. The width of the molecule cloud is plotted as a function of δ_R for Λ-cooling using (i) $|F = 2, 1\downarrow\rangle$ (yellow and purple), (ii) $|F = 1\uparrow, 1\downarrow\rangle$ (green and blue) and (iii) single frequency cooling (black line). The experimental conditions correspond to the best cooling we achieved in each case. As verified by the simulations, cooling using $|F = 2, 1\downarrow\rangle$ does not work for SrF and the widths are quite high, even at optimum δ_R. However, Λ-cooling using the $|F = 1\uparrow, 1\downarrow\rangle$ works quite well, and is very comparable to single-frequency cooling. Ultimately, both techniques on their own were able to cool down to $\sim 10\,\mu$K, good enough for our purposes.

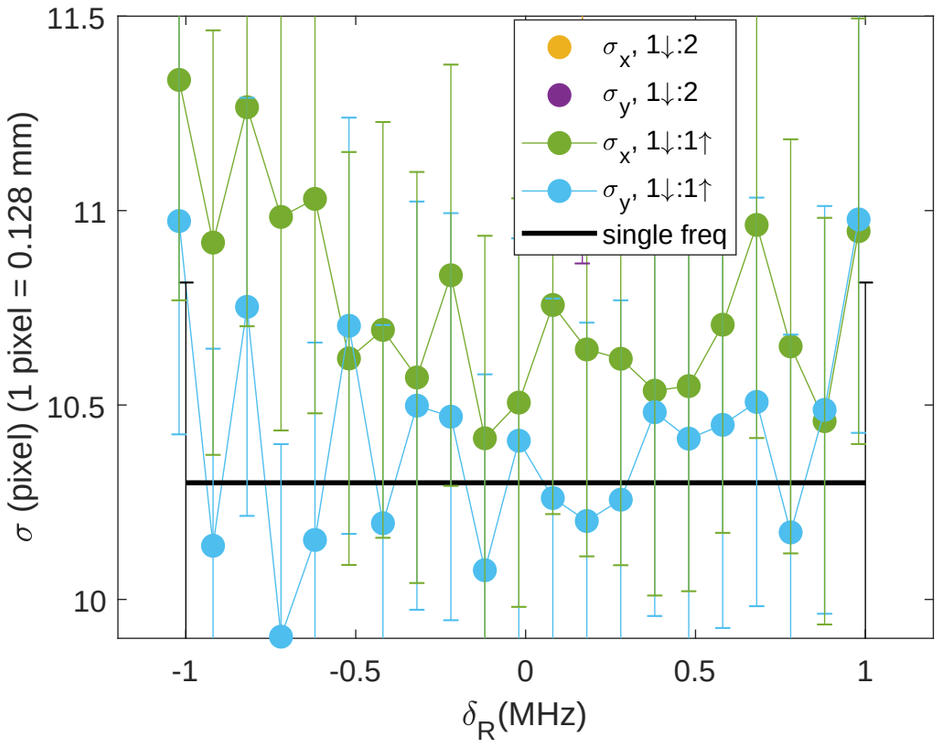

Figure 4.12: Comparison of single frequency and best Λ-cooling. The width of the molecule cloud is plotted as a function of δ_R for Λ-cooling using (i) $|F = 1\uparrow, 1\downarrow\rangle$ (green and blue) and (ii) single frequency cooling (black line). The experimental conditions correspond to the best cooling we achieved in each case. Λ-cooling using the $|F = 1\uparrow, 1\downarrow\rangle$ works quite well, and is very comparable to single-frequency cooling. Ultimately, both techniques on their own were able to cool down to $\sim 10\,\mu$K, good enough for our purposes.

4.6 Conclusions

In this chapter, we demonstrated successful deep laser cooling of SrF molecules to $T \approx 10\,\mu$K temperatures, mainly limited by our temperature measurement methods and off-resonant photon scattering. In particular, the demonstration of Λ-cooling is important because this can enable robust cooling even in the presence of AC Stark shifts from an ODT, as was demonstrated for CaF molecules [53]. Thus, this technique can be used to stochastically load into an ODT simply by cooling while overlapping the ODT beam with the cloud. If the ODT trap depth is sufficiently high, then successful loading can be obtained. This will be explored in the next chapter.

Chapter 5

Optical Dipole Trapping

5.1 Introduction

Until now, we were only trapping molecules in the MOT, which relies on active scattering of photons to keep trapping. This imposes limits on how dense the sample can be due to photon rescattering [113]. Furthermore, it also imposes a limit on how cold the molecules can get, due to photon recoils. Thus, there comes a point in any cooling experiment, where there is a need for a so-called conservative trap, i.e. a trap that does not rely on active scattering. These conservative traps have proven to be remarkably efficient at increasing the phase-space density (PSD) of atomic gases through either sympathetic [29] or evaporative cooling [30] techniques. There are several commonly used conservative traps in atomic cooling which have been equally successful: the magnetic quadrupole trap (MQT) [55, 114], the time orbiting potential (TOP) trap [115], the Ioffe-Pritchard trap [116], and the optical dipole trap (ODT) [84, 117]. Each of these has their own set of advantages which we detail below.

The MQT has several advantages: it can produce large trap depths of several 10s to 100s of mK with large capture volumes, and was used to capture molecules directly out of a buffer gas beam [48, 118, 119]. It has the disadvantage that at the center of the trap, the atoms can undergo Majorana spin flips and become untrapped [120, 121]. This can be overcome with different techniques, such as using the Ioffe-Pritchard and TOP traps which are designed to produce a non-zero magnetic field everywhere in the trap, thereby avoiding the Majorana losses. A further limitation of the MQT is that it cannot trap molecules in the absolute ground state, a consequence of Earnshaw's theorem that states that there can only be a field minimum in free space for static or quasi-static traps. Despite these limitations, owing to the large capture volume and trap depths, our group designed and implemented a MQT of SrF molecules [55]. At the temperatures achieved in that paper of $T \approx 50\,\mu$K, and for the field gradient applied in the trap, Majorana spin flip losses were expected to be below a rate of $R_{\max} \sim 0.05\ \mathrm{s}^{-1}$ [93], not a limiting concern by any means. Ultimately, our group was able to successfully capture and compress molecules in the MQT, leading to the first conservative trap of SrF molecules. The best densities obtained in that paper ($\sim 10^6\ \mathrm{cm}^{-3}$) were low enough that collisions between molecules were not yet observed. Thus there was an imminent need to improve the density in some way and get higher densities.

Around this time, the Doyle group demonstrated successfully demonstrated optical dipole trapping of CaF molecules using Λ-enhanced gray molasses to load the trap [53]. In their paper, they were able to transfer $\sim 5\%$ of the molecules from their MOT to the ODT, and achieve $\sim 5\,\mu$K

temperatures in the trap, a big jump in density as compared to the MQT ($\sim 10^9$ cm^{-3}). This was quite encouraging, and since we had also gotten Λ-cooling working (see Chapter 4), we were optimistic that we could also achieve these high densities and observe collisions between molecules. ODTs also have several advantages: ODTs typically have approximately equal trapping strengths for all internal states of the molecule, making it possible to trap in the rovibrational ground state. With the typical large detuning of the trap light from resonance, the scattering rate in an ODT is very low and thus long lifetimes and low temperatures can be achieved. Commercial high power 1064nm lasers are available, which makes it possible to achieve large trap depths even for a far detuned laser. Finally, a great variety of different trapping geometries can be realized including the standard harmonic trap, highly anisotropic traps or even multi-well potentials. Thus, we decided to switch from the MQT and pursue optical dipole trapping of SrF. The rest of this chapter goes over the ODT background, and details our experimental efforts towards realizing an ODT.

5.2 ODT principle

We briefly detail the principle of an ODT following our treatment from § 2.7. We recap that the interaction of a molecule with a detuned laser field leads to an energy shift of the internal levels given by:

$$\Delta E_{\text{AC Stark,i}} = \frac{e^2 I_0}{2c\hbar\epsilon_0} \sum_{f \neq i} \left(\frac{\langle i|\hat{r}\cdot\hat{p}^*|f\rangle \, \langle f|\hat{r}\cdot\hat{p}|i\rangle}{\omega_{if} - \omega} + \frac{\langle i|\hat{r}\cdot\hat{p}|f\rangle \, \langle f|\hat{r}\cdot\hat{p}^*|i\rangle}{\omega_{if} + \omega} \right) \tag{5.1}$$

where I_0 is the intensity of the laser beam, ω is the laser frequency and $\langle i|\hat{r}\cdot\hat{p}|f\rangle$ are the dipole matrix elements coupling the state $|i\rangle$ to the state $|f\rangle$ for a laser with polarization $\hat{p}$. For a ground state, when the laser is red-detuned of all the transitions, the Stark shift is negative, i.e. the energy of the ground state is lowered. Thus, if the laser has a spatially varying intensity (as is the case for a Gaussian beam), then the energy shift of the ground state is also spatially varying, as shown in Figure 5.1. The highest energy shift of the molecule is called the trap depth, and for usual ODTs, this is in the rage of $\sim 100 - 1000\,\mu$K. This creates a trap for the molecules, and in particular, this is a conservative trap because there is no damping force, there is only a restoring force. Thus, to load molecules in the trap, one needs to actively dissipate energy while loading, which is where the molasses cooling techniques are useful. As a molecule traverses the trap, if it scatters a cooling photon, while it is still in the trap, it loses energy and cannot escape. This also leads to a dramatic increase in density - typical ODT beam spot sizes are in the $30 - 40\,\mu$m range, thus producing very tight traps.

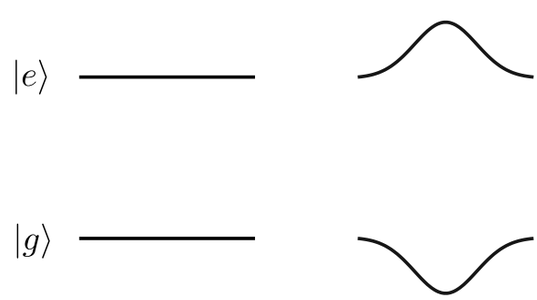

Figure 5.1: ODT Stark shift for a two level system in a laser beam with spatially varying intensity. The Gaussian nature of the laser beam produces an approximate harmonic trap at the trap minimum.

In the following, we will derive some useful quantities for our experiment such as the trap depth, the photon scattering rate etc. that will guide our experiments.

5.2.1 ODT trap depth and polarizability

The trap depth (U_T) is defined as the maximum AC Stark shift experienced by the molecules, i.e. the value of the AC Stark shift at the peak intensity. We shall be working with Gaussian laser beams throughout this thesis, though certainly other trapping geometries are actively used [122–124]. For a Gaussian beam propagating along the $\hat{z}$ axis, the spatial intensity profile is given by:

$$I(r,z) = I_0 \left(\frac{w_0}{w(z)}\right) \exp\left(\frac{-2r^2}{w(z)^2}\right) \tag{5.2}$$

where

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \tag{5.3}$$

where $z_R = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh range, and w_0 is called the beam waist or the radius of the beam at the focal spot. The peak intensity of the Gaussian beam is at $z = r = 0$ and is related to the power in the laser beam by:

$$I_0 = \frac{2P_0}{\pi w_0^2} \tag{5.4}$$

The peak Stark shift, and thus the trap depth is then

$$\Delta E_{\text{AC Stark,i}} = U_T = \frac{2P}{\pi w_0^2} \frac{e^2}{2c\hbar\epsilon_0} \sum_{f \neq i} \left(\frac{\langle i|\hat{r}\cdot\hat{p}^*|f\rangle \langle f|\hat{r}\cdot\hat{p}|i\rangle}{\omega_{if} - \omega} + \frac{\langle i|\hat{r}\cdot\hat{p}|f\rangle \langle f|\hat{r}\cdot\hat{p}^*|i\rangle}{\omega_{if} + \omega} \right) \tag{5.5}$$

We can immediately notice some properties of the trap depth based on this, namely that the trap depth is proportional to the ratio P/Δ. This is the reason why experiments prefer using 1064 nm light for ODTs, namely the high power available at this wavelength from conventional laser technology, allowing the generation of deep traps.

$\|i\rangle$	$\|f\rangle$	d_{fi}(Debye)	$E_f - E_i$ (cm^{-1})
$\|X^2\Sigma\rangle$	$\|A^2\Pi_{1/2}\rangle$	**6.22** [125]	**15076** [91]
$\|X^2\Sigma\rangle$	$\|A^2\Pi_{3/2}\rangle$	**6.24** [125]	**15357** [91]
$\|X^2\Sigma\rangle$	$\|B^2\Sigma_{1/2}\rangle$	**4.93** [126]	**17267** [127]
$\|X^2\Sigma\rangle$	$\|C^2\Pi_{1/2}\rangle$ & $\|C^2\Pi_{3/2}\rangle$	1.53 [128]	**27385** [91]
$\|X^2\Sigma\rangle$	$\|D^2\Sigma_{1/2}\rangle$	0.82 [128]	**27774** [91]
$\|X^2\Sigma\rangle$	$\|F^2\Sigma_{1/2}\rangle$	1.02 [128]	**32824** [129]
$\|X^2\Sigma\rangle$	$\|G^2\Pi_{1/2}\rangle$ & $\|G^2\Pi_{3/2}\rangle$	1.27 [128]	**34809** [129]
$\|A^2\Pi_{1/2}\rangle$	$\|A'^2\Delta_{3/2}\rangle$	6.89 [79]	4036 [79]
$\|A^2\Pi_{1/2}\rangle$	$\|B^2\Sigma_{1/2}\rangle$	0.53 [79]	**2195** [127]
$\|A^2\Pi_{1/2}\rangle$	$\|C^2\Pi_{1/2}\rangle$ & $\|C^2\Pi_{3/2}\rangle$	2.8 [128]	**12313** [91]
$\|A^2\Pi_{1/2}\rangle$	$\|D^2\Sigma_{1/2}\rangle$	5.34 [128]	**12702** [91]

Table 5.1: Transition dipole moments and energies used in AC Stark shift calculations. The sources for the values are various experimental (bold) results and theoretical calculations.

Now, to calculate the trap depth for our states of interest, we recall that the dipole matrix element between two states expressed in the Hund's case a basis is:

$$\langle \Lambda', S, \Sigma', \Omega', J', F', M_F' | r_p^1 | \Lambda, S, \Sigma, \Omega, J, F, M_F \rangle = (-1)^{F'-M_F'} \begin{pmatrix} F' & 1 & F \\ -M_F' & p & M_F \end{pmatrix}$$
$$(-1)^{J'+I+F'+1} \begin{Bmatrix} J' & F' & I \\ F & J & 1 \end{Bmatrix} \sqrt{2F+1}\sqrt{2F'+1}$$
$$(-1)^{J'-\Omega'} \sqrt{2J+1}\sqrt{2J'+1} \begin{pmatrix} J' & 1 & J \\ -\Omega' & \Omega'-\Omega & \Omega \end{pmatrix}$$
$$\langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle \, \delta_{\Sigma,\Sigma'}$$
$$= C_{\eta'\eta} \langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle \, \delta_{\Sigma,\Sigma'} \tag{5.6}$$

where the $\langle \Lambda' | r^1_{\Omega'-\Omega} | \Lambda \rangle$ are the transition dipole moments. We follow the general method of converting case b states into case a using eq. 2.26 and then calculating all the matrix elements as above. We list the transition dipole moments as well as the energy differences for the states used in the trap depth calculation in Table 5.1. Then, the trap depth is obtained by summing the contributions from all the rovibrational and hyperfine states within each electronic state, following eq. 5.5. We ignore the relative hyperfine energy splitting within each electronic level as they are of much lower magnitude than the electronic energy differences. We do include the J-mixing coefficients for each $^2\Sigma$ state for the full calculation.

5.2.1.1 Trap depths and polarizabilities for $\left|X^2\Sigma, N=1\right\rangle$ states

For our calculations, we choose $\lambda = 1064$ nm, the wavelength of our ODT laser. Our measurements indicate that roughly 50 W of the light makes it through the chamber, and the beam is focused to $w_0 \approx 40\,\mu$m. So we find that $I_0 = 2$ MW/cm^2. The trap depths for the various hyperfine sublevels in the $\left|X^2\Sigma, N=1\right\rangle$ are given in Table 5.2, calculated for $\hat{p} = \sigma^+$ light. We also list the corresponding energy shifts of each sublevel as well as the relative shift in each sublevel. Note that the values listed in Table 5.2 are lower by a factor of 2 than the correct values, however we have kept these values here to be consistent with our paper [111]. This error is because we missed a factor of 2 in the intensity of a Gaussian beam, where the intensity was taken to be $P_0/\pi w_0^2$ as opposed to $2P_0/\pi w_0^2$. In the rest of this chapter, this error is propagated forward - **all the temperatures, trap depths and AC Stark shifts are higher by a factor of 2**.

The AC Stark Hamiltonian can also be written as follows [130]:

$$\Delta E_{AC} = -\frac{1}{2}\alpha E^2 \tag{5.7}$$

where $\alpha = \alpha_S + \alpha_V + \alpha_T$ are called the scalar, vector, and tensor polarizabilities, and E is the electric field produced by the laser light. The polarizabilities are, in general, dependent on the wavelength of the light, and are indicative of how large the Stark shift is for a given electric field. This can be used to determine so called "magic" trapping conditions where, for a particular wavelength of the trap light, the polarizabilities for a pair of ground and excited states is the same, meaning that the transition is insensitive to fluctuations in the light intensity, making it very robust [131–133]. This is very important for clock experiments where any differential light shifts can pollute the measurement; these experiments go to great lengths to acquire the right

State $\|F, m_F\rangle$	U_T (μK)	$E_{AC}/(2\pi)$ (MHz)	$\Delta E_{AC}/(2\pi)$ (kHz)	$\frac{\alpha_V}{\alpha_S}$	$\frac{\alpha_T}{\alpha_S}$
$\|1\downarrow, -1\rangle$	553	-11.52	310		
$\|1\downarrow, 0\rangle$	599	-12.49	-665	-0.0019	0.0563
$\|1\downarrow, 1\rangle$	550	-11.47	355		
$\|0, 0\rangle$	568	-11.83	0	0	0
$\|1\uparrow, -1\rangle$	576	-12.00	-174		
$\|1\uparrow, 0\rangle$	558	-11.64	178	-0.0072	-0.0151
$\|1\uparrow, 1\rangle$	568	-11.83	-4		
$\|2, -2\rangle$	546	-11.38	444		
$\|2, -1\rangle$	580	-12.09	-264		
$\|2, 0\rangle$	591	-12.31	-487	-0.0036	0.0824
$\|2, 1\rangle$	578	-12.05	-222		
$\|2, 2\rangle$	542	-11.29	530		

Table 5.2: Calculated trap depths, energy shifts and polarizabilities for each state in the $\left|X^2\Sigma, N = 1\right\rangle$ manifold. First column indicates the $|F, m_F\rangle$ states in the $\left|X^2\Sigma, N = 1\right\rangle$ level. Second column is the trap depth for each level for σ^+ polarized ODT light with $I \sim 2.0$ MW/cm^2. Third column is the absolute energy shift of each level. Fourth column is the relative energy shift for each level, i.e. $E - E_s$ where E_s is the scalar Stark shift. Fourth and fifth columns are the extrapolated vector (α_V) and tensor (α_T) polarizabilities normalized by the scalar polarizability (α_S). These are only F dependent, and so are specified once for each hyperfine manifold. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

wavelength light for the trap. For our case, however, the trap wavelength is fixed to 1064 nm, and we will only extrapolate these polarizabilities based on our trap depth calculations. This is important for us because we would like to know the Stark shifts as a function of the ellipticity of the ODT beam. Here, we call the ellipticity of the ODT beam as γ_{ODT} with $\gamma_{ODT} = \pm 45$ corresponding to $\sigma^{\pm}$ and $\gamma_{ODT} = 0$ corresponding to $\hat{z}$ polarization.

For a $|F = 1\rangle$ state, the energy shift is given by:

$$\begin{aligned}\frac{H_{AC}}{E_s} = &\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \frac{\alpha_V}{\alpha_S} \begin{pmatrix} -\sin(2\gamma_{ODT}) & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sin(2\gamma_{ODT}) \end{pmatrix} \\ &- \frac{\alpha_T}{\alpha_S} \begin{pmatrix} \frac{1}{2} & 0 & \frac{3}{2}\cos(2\gamma_{ODT}) \\ 0 & -1 & 0 \\ \frac{3}{2}\cos(2\gamma_{ODT}) & 0 & \frac{1}{2} \end{pmatrix}\end{aligned} \tag{5.8}$$

where E_s is the scalar Stark shift (equal to the shift for $|0, 0\rangle$ and also the average shift over all $|F, m_F\rangle$ states in Table 5.2). Note that for $\hat{p} = \pm 1$ (circularly polarized light), this Hamiltonian is diagonal and so we can directly extract the ratios α_V/α_S and α_T/α_S from the calculated energy shifts from Table 5.2, and are also displayed in that table. The trap depth for the $|F = 1\rangle$ states as well as their Stark shift is shown in Figure 5.2 as a function of the ODT ellipticity.

For the $|F = 2\rangle$ state, the energy shift is given by:

$$\frac{H_{AC}}{E_s} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} - \frac{\alpha_V}{\alpha_S} \begin{pmatrix} -\sin(2\gamma_{ODT}) & 0 & 0 & 0 & 0 \\ 0 & -\frac{\sin(2\gamma_{ODT})}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{\sin(2\gamma_{ODT})}{2} & 0 \\ 0 & 0 & 0 & 0 & \sin(2\gamma_{ODT}) \end{pmatrix}$$
$$- \frac{\alpha_T}{\alpha_S} \begin{pmatrix} \frac{1}{2} & 0 & \sqrt{\frac{3}{8}}\cos(2\gamma_{ODT}) & 0 & 0 \\ 0 & -\frac{1}{4} & 0 & \frac{3}{4}\cos(2\gamma_{ODT}) & 0 \\ \sqrt{\frac{3}{8}}\cos(2\gamma_{ODT}) & 0 & -\frac{1}{2} & 0 & \sqrt{\frac{3}{8}}\cos(2\gamma_{ODT}) \\ 0 & \frac{3}{4}\cos(2\gamma_{ODT}) & 0 & -\frac{1}{4} & 0 \\ 0 & 0 & \sqrt{\frac{3}{8}}\cos(2\gamma_{ODT}) & 0 & \frac{1}{2} \end{pmatrix} \tag{5.9}$$

Again, this Hamiltonian is diagonal when $\hat{p} = \pm 1$, and the polarizabilities can be extracted directly from the calculated trap depths, as in Table 5.2. The trap depth for the $|F = 2\rangle$ states, as well as their Stark shift is shown in Figure 5.2 as a function of ODT ellipticity.

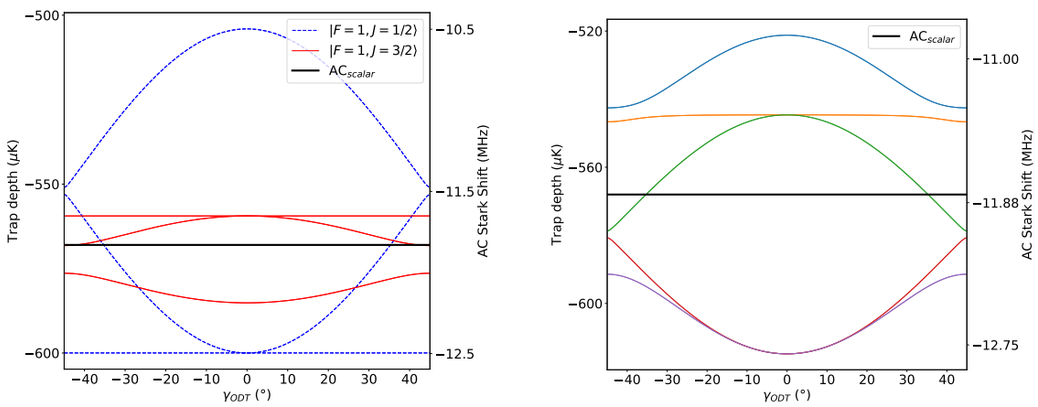

(a) AC Stark shift for $|F = 1\rangle$

(b) AC Stark shift for $|F = 2\rangle$

Figure 5.2: Trap depths and AC Stark shifts for the $|F = 1\rangle$ and $|F = 2\rangle$ levels. (a) Stark shift for the $|F = 1\rangle$ levels and (b) for the $|F = 2\rangle$ levels as a function of ellipticity. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

With the full power of the ODT beam, we see that we can produce roughly $U_T \approx 580 - 600\,\mu$K deep traps. The trap frequencies can be determined assuming the trap is harmonic near the minimum [84], and for our trap, are $\omega_{x,y,z} = 2\pi \times (1.7 \times 10^3, 1.7 \times 10^3, 9)$ s^{-1}. For optimal loading, the ratio of the trap depth to temperature should be $U_T/T \lesssim 10$ [134]. With our free space temperature $T \approx 10\,\mu$K, we should be able to load the molecules into this trap easily.

5.2.1.2 Trap depth for the $A^2\Pi_{1/2}$ states

Since the Λ-cooling couples the $\left|X^2\Sigma, N = 1, F = 1\uparrow, 1\downarrow\right\rangle$ to the $\left|A^2\Pi_{1/2}, F = 0, 1\right\rangle$ states, we are also interested in calculating the Stark shift for these manifolds. The states that couple to the $A^2\Pi$ state are displayed in Table 5.1 with their transition dipole moments and energy differences. We again perform the same calculation for the same laser conditions and the results are summarized in Table 5.3. Again, note that these values are lower by a factor of 2 to reflect the results in our paper [111]. We have extrapolated the polarizabilites as we did in the previous section. This manifold does not have any tensor shift, and is only subject to a vector shift, as shown in Figure 5.3.

State $\|F, m_F\rangle$	U_T (μK)	$E_{AC}/(2\pi)$ (MHz)	$\Delta E_{AC}/(2\pi)$ (kHz)	$\frac{\alpha_V}{\alpha_S}$	$\frac{\alpha_T}{\alpha_S}$
$\|1,-1\rangle$	-294	-6.17	-6560		
$\|1,0\rangle$	19	0.39	0	16.7106	0.0
$\|1,1\rangle$	334	6.96	6570		
$\|0,0\rangle$	19	0.39	0	0	0

Table 5.3: Calculated trap depths, energy shifts and polarizabilities for each state in the $\left|A^2\Pi_{1/2}\right\rangle$ manifold. First column indicates the $|F, m_F\rangle$ states. Second column is the trap depth for each level for σ^+ polarized ODT light with $I \sim 2.0$ MW/cm^2. Third column is the absolute energy shift of each level. Fourth column is the relative energy shift for each level, i.e. $E - E_s$ where E_s is the scalar Stark shift. Fourth and fifth columns are the extrapolated vector (α_V) and tensor (α_T) polarizabilities normalized by the scalar polarizability (α_S). These are only F dependent, and so are specified once for each hyperfine manifold. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

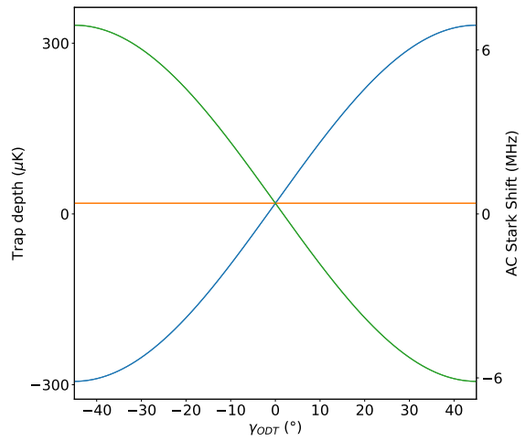

Figure 5.3: Trap depth and Stark shift for the $A^{\Pi}_{1/2}$ states. This manifold does not have a tensor shift. The blue and green traces correspond to $|F = 1, m_F = \pm 1\rangle$ respectively and the orange trace corresponds to both $|F = 1, m_F = 0\rangle$ and $|F = 0, m_F = 0\rangle$. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

5.2.2 Scattering rate

Even though the ODT beam is far detuned from the transition resonance, there is still a non-zero photon scattering rate, and it is important to estimate this scattering rate to understand if heating due to photon scatters is a problem. We can estimate the average scattering rate (Γ_{sc}) by following Ref. [84] to find:

$$\Gamma_{sc} = \frac{\Gamma}{\hbar\Delta}\left(U_T + \frac{3}{2}k_B T\right) \tag{5.10}$$

where $\Gamma = 2\pi \times 6.6$ MHz is the natural linewidth of the $X \to A$ transition, $U_T \approx -600\,\mu$K is the trap depth, $\Delta = -2\pi \times 160$ THz is the detuning, and $T \approx 10\,\mu$K is the temperature of the molecules in the trap. The molecules experience an average trapping potential $U_T + 3/2k_B T$. Then, we find:

$$\Gamma_{sc} \approx 2\pi \times 0.5 \text{ s}^{-1} \tag{5.11}$$

The heating rate (P) can then be estimated as:

$$P = 2E_{rec}\Gamma_{sc} \tag{5.12}$$

where E_{rec} is the photon recoil energy, and Γ_{sc} is the photon scattering rate. Using the above calculation, we find:

$$P = 7 \text{ nK/s} \tag{5.13}$$

which would not be a big concern in the experiment.

5.3 Experimental setup

The complete apparatus is shown in Figure 5.4 and here we go over the key details and some of our measurements.

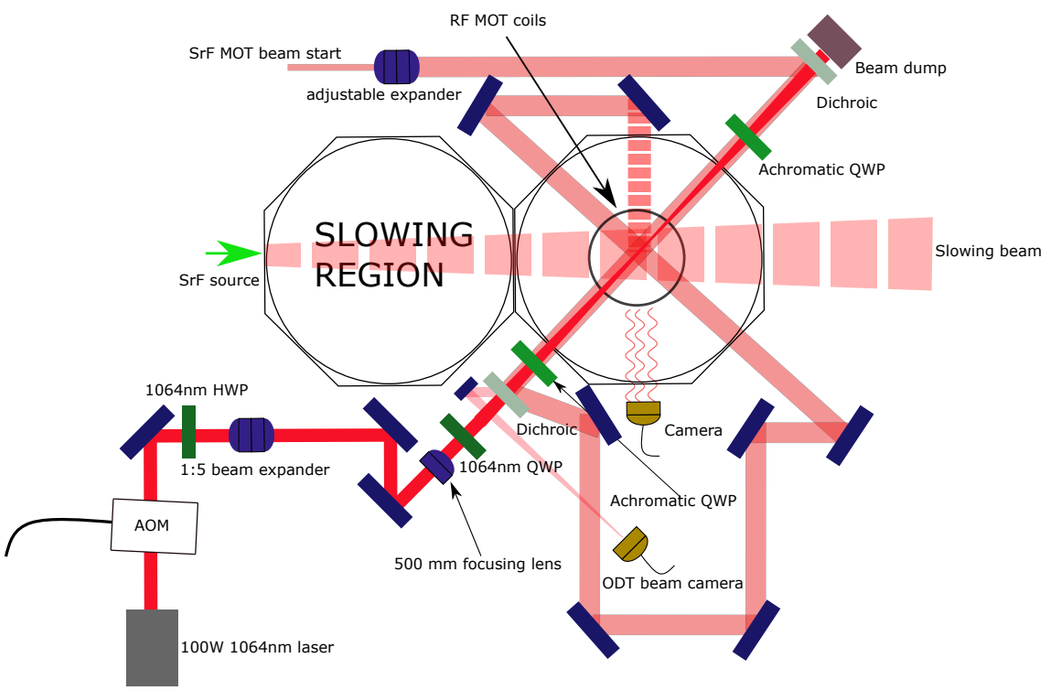

Figure 5.4: ODT beam setup. The ODT beam passes through an AOM, a HWP, a 1:5 beam expander, a focusing lens, a QWP, the MOT dichroic, and the MOT QWP before entering the chamber. On the other end of the chamber, the ODT beam is deposited in a beam dump. The 6 passes of the MOT beam are also shown.

5.3.1 ODT beam

Our ODT beam is produced from a 1064 nm, 100 W fiber laser. The beam coming out of the laser is around 3 mm $1/e^2$ diameter at full power[1]. The ODT beam is then passed through a high power isolator (Thorlabs IO-10-1064-VHP) before it is sent into a high power acousto-optic modulator (AOM). The AOM is resonant at ~ 80 MHz, and is used for fast control of the ODT intensity by applying a RF tone to it. The first order diffracted beam from the AOM is then passed through a half-wave plate (HWP) before it is expanded using a 1:5 beam expander (Edmund Optics 5X, 1064nm Vega Nd:YAG Laser Line Beam Expander). The beam is then passed through a 500 mm focusing lens, and a quarter-wave plate (QWP). The first HWP along with this QWP allow for arbitrary control over the ODT beam polarization. Finally, the beam is combined with the first pass of the MOT beam to go into the chamber. For this, the MOT mirrors along this path are replaced with dichroic mirrors that reflect below 900 nm and transmit above it so that the MOT beam is unaffected. The two MOT QWPs along this path are also replaced with achromatic QWPs (Edmund Optics 25.4mm Dia, 650 - 1100nm, $\lambda/4$ Achromatic Waveplate) so that the ODT beam does not burn it. On exiting the chamber, the ODT beam passes through the second achromatic QWP, the second dichroic mirror and gets dumped in a beam dump (Thorlabs LB2).

[1] We have seen that the waist is higher at low power and as the power is increased, the beam shrinks in size and the convergence of the beam changes. As a result, all the alignment is performed in the "high power beam, low power AOM" mode where the ODT beam is kept at full power, but only a fraction of it is passed through the AOM.

With a 500 mm focusing lens, and a 15 mm $1/e^2$ beam diameter before it, we would expect a focus spot of $w_0 \approx 20 - 40\,\mu\text{m}^2$, very close to what we achieve in the lab, see Figure 5.5. The beam profiling was carried out using a Thorlabs BC106N-VIS beam profiler placed a variable distance after the focusing lens. Figure 5.5 also shows the trap depth as a function of $\hat{z}$, the axis of propagation of the ODT beam. If the temperature of the cloud T is much smaller than the trap depth, the cloud size axially will be small compared to the Rayleigh range, in this case we can approximate the trapping potential along $\hat{z}$ as:

$$U_T(z) \simeq U_T \left(-1 + \frac{z^2}{z_R^2}\right) \tag{5.14}$$

which yields a $z_R \approx 5.3$ mm for this profile and $U_T \approx 597\,\mu$K.

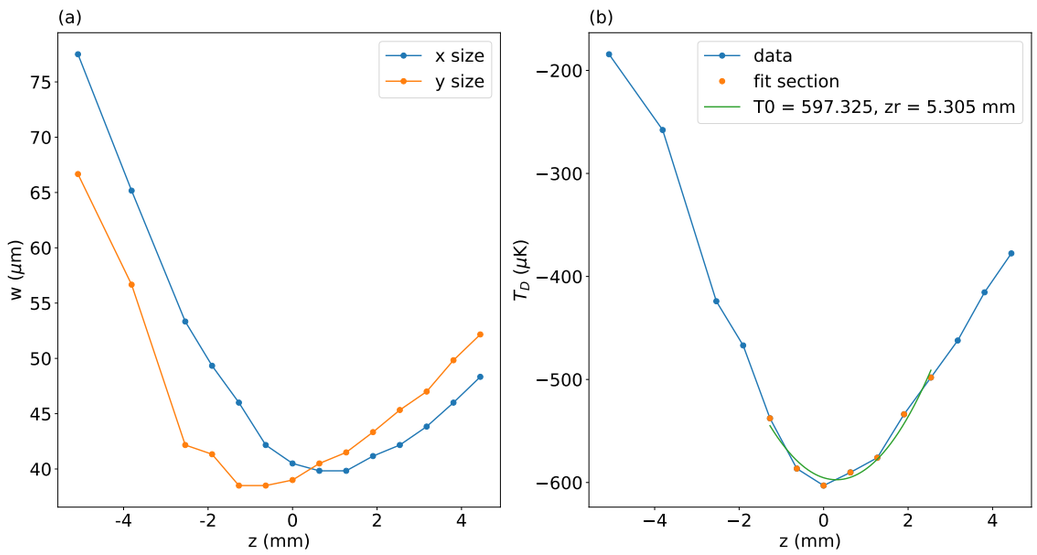

Figure 5.5: ODT beam Profile at Yale. (a) Beam profile data taken from a beam profiler. There is a slight astigmatism in the beam due to misalignment in the beam expander before the focusing lens. (b) The beam profile converted to trap depth and fitted to a quadratic at the minimum to obtain a Rayleigh length z_R that is used for temperature determination. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

There is some astigmatism in the beam due to misalignment through the 1:5 beam expander. The beam expander is made of very short focal length lenses which makes it critical to align the beam with exactly normal incidence and exactly centered on the lens - in practice this is quite tough and Figure 5.5 shows the best we could achieve for this at Yale.

Another concern with the ODT beam was about the mode quality, especially if it was single mode or not. Multi-mode beams can drive Raman transitions between different rovibrational states, which could lead to rapid loss. We set up a Fabry-Perot cavity to monitor the laser mode by picking off some leaked light from the isolator. Figure 5.6 shows that the ODT beam is indeed single mode, and stays single mode throughout. We do see that the ODT frequency drifts after turn on, until it is fully warmed up (around $\sim$ 1 hour) when it becomes stable. We just turn on the beam at the beginning of the day so that by the time we run the ODT, it is in a stable mode.

[2]The spot size for a lens is given by the formula $w_0 = \frac{1.27M^2\lambda f}{d}$ where M^2 is the laser beam parameter, λ is the wavelength, f is the lens focal length, and d is the beam diameter before the lens. Our beam quality is not exactly 1, but is around 1.5-2 leading to higher spot sizes.

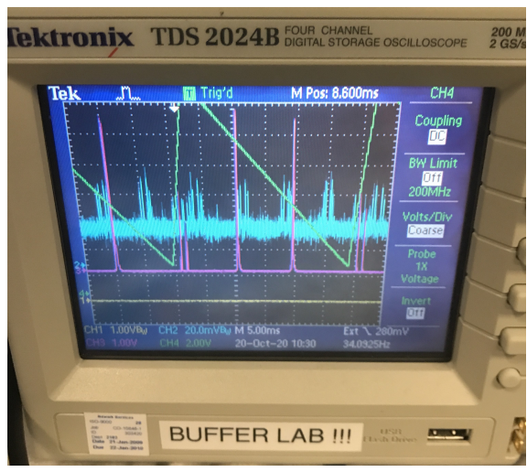

Figure 5.6: ODT beam mode quality on a cavity. The ODT beam (pink trace) is single mode, as seen on the cavity ramp (green trace), with V_{10}^{M} slave laser being monitored (blue trace) on another channel.

5.3.2 AOM control

The AOM is controlled by applying a RF tone to it. It has a broad resonance of around 15 MHz with the center at 80 MHz. At the resonant frequency, we were able to obtain a best diffraction efficiency of around 80%. Our initial plan was to use this AOM to implement a painted ODT to increase the trapping fraction, following [135]. For this, we needed to sweep the ODT beam spot location by dynamically changing the RF frequency to the AOM, and ultimately have it converge to a central frequency to compress the captured cloud. For this, we built a system with a microcontroller (Teensy 4.0) that was driving a direct digital synthesizer (DDS, AD9910). The teensy could be triggered externally, and we could configure arbitrary frequency waveforms on it, and this would then be used to control the DDS. An example sweep implemented with this is shown in Figure 5.7 where the ODT beam is swept along the horizontal axis to create a flat-top profile, while the vertical axis still maintains its tight focus. The diffraction efficiency for most of the flat top is around 80%, which is optimized at each frequency interval by tuning the RF powers and uploading that to the Teensy.

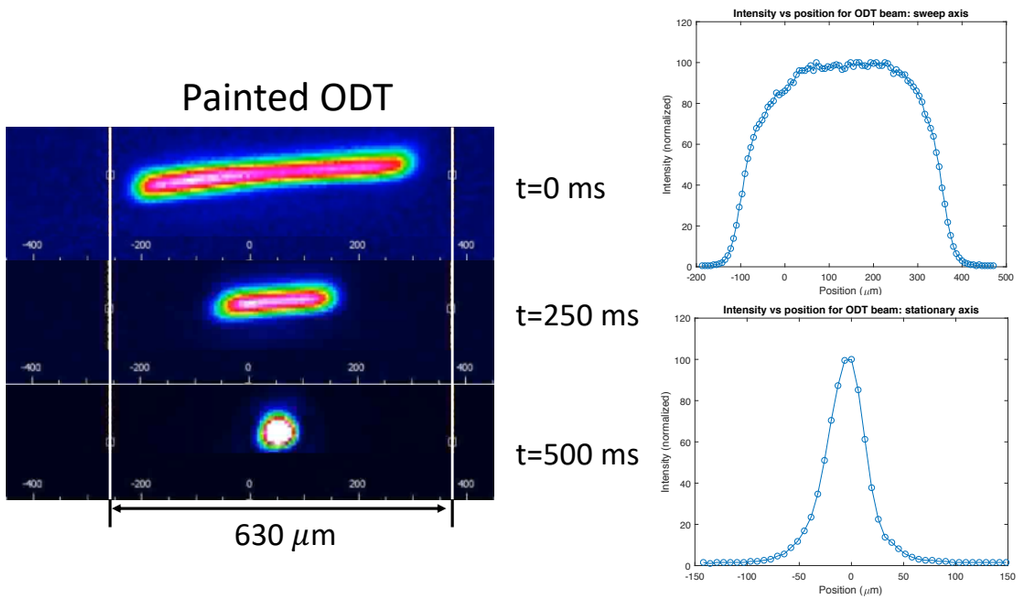

Figure 5.7: ODT beam sweep using the Teensy + DDS. (left) ODT beam sweep imaged on a camera. (right) The beam sweep intensities along the swept (top) and unswept (bottom) axes at $t = 0$ ms. The time averaged potential looks like a flat top along the swept axis, but maintains the gaussian profile along the unswept axis which is along gravity.

Ultimately, this idea was not implemented because the ODT loading is saturated only at higher trap depths. So sacrificing the time-averaged intensity to sweep the beam was not feasible. We replaced the Teensy driver with the standard driver that came with the AOM, and we were again able to achieve around 80% diffraction efficiency.

5.3.3 Polarization control

For the initial pass of the ODT, we were not too careful about the polarization, since we did not think it mattered too much, and our main focus was on getting roughly linear polarization since that is the easiest to analyze. To get linear polarization, we had only added a single QWP after the focusing lens, to counteract the MOT QWP rotation. However, we soon realized that the dichroic changes the polarization significantly, and furthermore, we also wanted the ability to control the polarization arbitrarily, i.e. not only make linear or circular, but also elliptical polarization with different signs. We added a HWP + QWP along the ODT beam path before it goes through the dichroic. To analyze the resultant polarizations, we resort to the Jones matrix formalism, where the effect of each optical element that can rotate the polarization is given by the Jones matrix. We treat the dichroic + MOT QWP as a single polarization element, and we measure the Jones matrix of this combination. The formalism in this section was developed by T. Langin. A generic Jones matrix can be written as:

$$J_{gen} = \begin{bmatrix} a & b \\ -e^{i\nu}b^* & e^{i\nu}a^* \end{bmatrix} = \begin{bmatrix} \cos[\psi]e^{i\alpha} & \sin[\psi]e^{i\beta} \\ -\sin[\psi]e^{i(\nu-\beta)} & \cos[\psi]e^{i(\nu-\alpha)} \end{bmatrix} \tag{5.15}$$

where $a = \cos[\psi]e^{i\alpha}$ and $b = \sin[\psi]e^{i\beta}$. There are 4 free parameters. We can look at the effect of this on linearly polarized light $[1, 0]^T$ to see:

$$\epsilon_{Lin}J_{gen} = \begin{bmatrix} e^{i\alpha}\cos[\psi] \\ -e^{i(\nu-\beta)}\sin[\psi] \end{bmatrix} = \begin{bmatrix} e^{i(\alpha+\beta-\nu-\pi)}\cos[\psi] \\ \sin[\psi] \end{bmatrix} \tag{5.16}$$

where an extra π is added which does not affect the results but makes the algebra easier. The parameters ψ and $\mu = \alpha + \beta - \nu - \pi$ can be determined by looking at the transmitted amplitude through a PBS cube for a few different configurations. Now, we can replace $\nu = \alpha + \beta - \mu - \pi$ in the Jones matrix, factor out $e^{i\beta}$ and define $\delta = \alpha - \beta$ to find:

$$J_{gen} = \begin{bmatrix} \cos[\psi]e^{i\delta} & \sin[\psi] \\ \sin[\psi]e^{i\delta}e^{-i\mu} & -\cos[\psi]e^{-i\mu} \end{bmatrix} \tag{5.17}$$

We have already measured ψ and μ, but measuring δ is a bit tricky. We can measure this by looking at the horizontal component after a PBS with circular and 45° linearly polarized light incident on the dichroic + MOT QWP. We have for circular polarization:

$$|\epsilon_{circ,x}|^2 = |J_{gen}(\psi, \delta, \mu) \begin{bmatrix} 1/\sqrt{2} \\ i/\sqrt{2} \end{bmatrix}|^2 = \frac{1}{2} + \cos[\psi]\sin[\psi]\sin[\psi] \tag{5.18}$$

and for 45° linearly polarized light:

$$|\epsilon_{lin,x}|^2 = |J_{gen}(\psi,\delta,\mu)\begin{bmatrix}1/\sqrt{2}\\1/\sqrt{2}\end{bmatrix}|^2 = \frac{1}{2} + \cos[\psi]\sin[\psi]\cos[\psi] \tag{5.19}$$

This is shown in Figure 5.8 where the transmission is plotted vs δ. By measuring the transmission for two different polarizations as stated above, a unique δ is found.

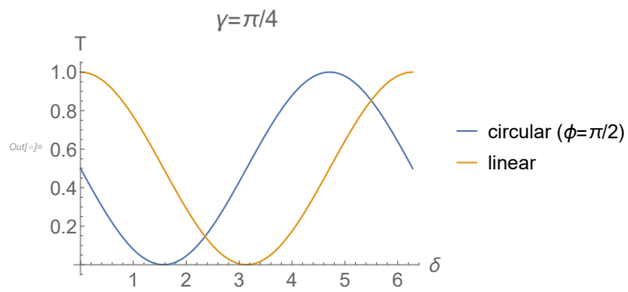

Figure 5.8: Transmission vs δ for Jones matrix determination. Measuring the transmission for two polarizations uniquely determines δ.

We used this technique to determine that our dichroic + QWP combination has the Jones matrix defined by $\psi = 0.82, \mu = 1.85, \delta = 5.87$.

With this measurement, we can now determine what we need to produce any arbitrary polarization. As shown in Figure 5.4, the ODT beam is horizontally polarized after the fiber, and it then passes through a HWP, a QWP and the dichroic + QWP combination. The total Jones vector on passing through these is simply:

$$\begin{bmatrix}E_x\\E_y\end{bmatrix} = J_{gen}(\psi,\mu,\delta)R[\omega_{QWP}]M_{QWP}R[-\omega_{QWP}]R[\omega_{HWP}]M_{HWP}R[-\omega_{HWP}]\begin{bmatrix}1\\0\end{bmatrix} \tag{5.20}$$

where

$$R[\omega] = \begin{bmatrix}\cos[\omega] & -\sin[\omega]\\ \sin[\omega] & \cos[\omega]\end{bmatrix} \tag{5.21}$$

handles the rotation of the waveplates, and

$$M_{QWP} = \begin{bmatrix}e^{-i\pi/4} & 0\\ 0 & e^{i\pi/4}\end{bmatrix} \quad M_{HWP} = \begin{bmatrix}e^{-i\pi/2} & 0\\ 0 & e^{i\pi/2}\end{bmatrix} \tag{5.22}$$

are the Jones matrices for the waveplates when the fast axis is aligned with the horizontal axis. The last piece of the puzzle is that the electric field evolves in time according to $\vec{E} = E_x\hat{x}e^{-i\omega t} + E_y\hat{y}e^{-i\omega t}$. We can then look at parametric plots of the magnitude of the field $(x_{mag}(\omega_{HWP},\omega_{QWP},t), y_{mag}(\omega_{HWP},\omega_{QWP},t))$ to determine the ellipticity as:

$$\gamma(\omega_{HWP},\omega_{QWP}) = \arctan\left[\frac{\min(\sqrt{x_{mag}^2+y_{mag}^2})}{\max\sqrt{x_{mag}^2+y_{mag}^2}}\right] \tag{5.23}$$

where the min and max are taken over a full rotation of $t = [0, 2\pi)$. Similarly, the orientation of the ellipse may be defined as:

$$\chi(\omega_{HWP},\omega_{QWP}) = \arctan\left[\frac{y_{mag}(t_{max})}{x_{mag}(t_{max})}\right] \tag{5.24}$$

where t_{max} is the time for which the total magnitude $\sqrt{x_{mag}^2 + y_{mag}^2}$ is maximized. χ thus defines the angle of the major axis of the ellipse. The results of this analysis are shown in Figure 5.9 where we show arbitrary polarization generation by varying the HWP + QWP angles to vary χ and γ. This allowed us to explore the full range of ODT polarizations, as was needed in the paper [111].

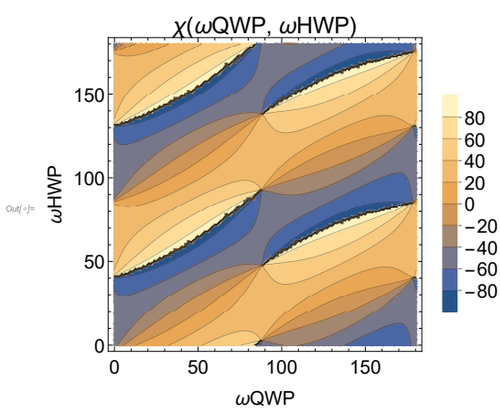

(a) χ vs $(\omega_{HWP}, \omega_{QWP})$

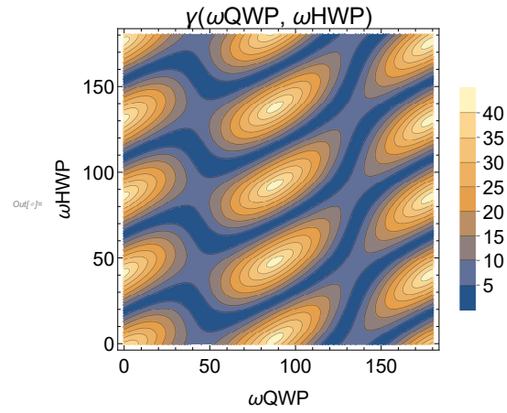

(b) γ vs $(\omega_{HWP}, \omega_{QWP})$

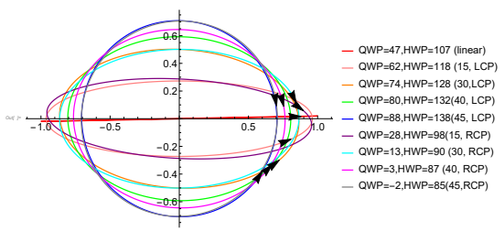

(c) Polarization ellipses for $\chi = 0°$

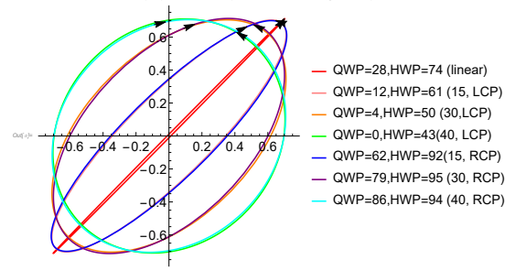

(d) Polarization ellipses for $\chi = 45°$

Figure 5.9: Arbitrary polarization generation using a HWP + QWP for the ODT. (a) Contour plot of χ and (b) γ vs $(\omega_{HWP}, \omega_{QWP})$. We can generate arbitrary polarizations by changing the HWP and QWP angles. (c) and (d) show same polarization ellipses for fixed $\chi = 0°$ and $\chi = 45°$, where the γ for each ellipse and its sense of rotation is indicated.

5.3.4 Experiment sequence

As in the previous chapter, all our experiments start with the molecule beam produced in the CBGB that is then slowed and trapped in a RF MOT ($\sigma \sim 1$ mm, $T \sim 1$ mK, $N \sim 3500$). Subsequently, the polarization switching and the RF gradient is turned off, V_{00}^M light is shuttered, and V_{00}^d frequency is jumped to the Λ-cooling configuration. The Λ-light is kept on until the ODT loading is complete, after which the light is shuttered, and untrapped molecules are allowed to fall out of the imaging area. The fall duration is varied between 50 ms, if imaging is being done with the Λ-cooling light in situ, and 150 ms, if the molecules are to be recaptured in the MOT and imaged. The Λ-cooling and MOT recapture imaging durations are nominally 100 ms, unless specified otherwise.

We also use an in-vacuum shutter in the slowing region for all these experiments. The slowing shutter is normally closed, and is only opened for the slowing duration of 35 ms. This was found to greatly increase the ODT lifetime by blocking ballistic helium.

5.4 Results

Here we go over our initial attempts at an ODT, and how we were able to use the ODT polarization to maximize trapping efficiency. The results of this section are also published in Ref. [111].

5.4.1 Initial attempts

Once we had Λ-cooling working and the ODT beam line was assembled, we were ready to try loading an ODT. Our initial guess for parameters such as loading time and the detuning were taken from the Doyle group paper [53] where they loaded the ODT for around 150 ms and with $\delta_R \approx 0.1$ kHz. We tried a variety of initial attempts where we varied the load times and the Λ-cooling times but did not really see any ODT signal. We tried aligning the ODT beam by making use of the MOT beam and centering it very well but it did not yield anything. At this point, we also checked the laser mode, to make sure the laser was indeed single mode, see Figure 5.6. The problem was that we were trying to image the molecules using the Λ-cooling light in situ, which makes it difficult to image very small numbers of molecules if the ODT is not optimized. At this point, we realized that we can actually image tiny signals if we can recapture the molecules in the MOT, and collect fluorescence from the MOT for a long time - this is only limited by the MOT lifetime which is ~ 300 ms. To check this method, we first experimented with Λ-cooling immediately after the MOT, where we cooled for ~ 50 ms and then dropped the molecules by extinguishing the light. We found that if we drop for $t > 150$ ms, the molecules cannot be recaptured as they have fallen too far. This would provide us a clear signal if the ODT was indeed capturing any molecules - if, after loading, we extinguish the light and recapture in the MOT after 150 ms, the only signal we should see is from molecules captured in the ODT. This was right around the time Covid hit the world and we got our first ODT signal a day before Yale closed the campus for the pandemic, see Figure 5.10. With this initial signal in hand, we can now try to optimize the other parameters - but that would have to wait until after Yale gave us the all clear.

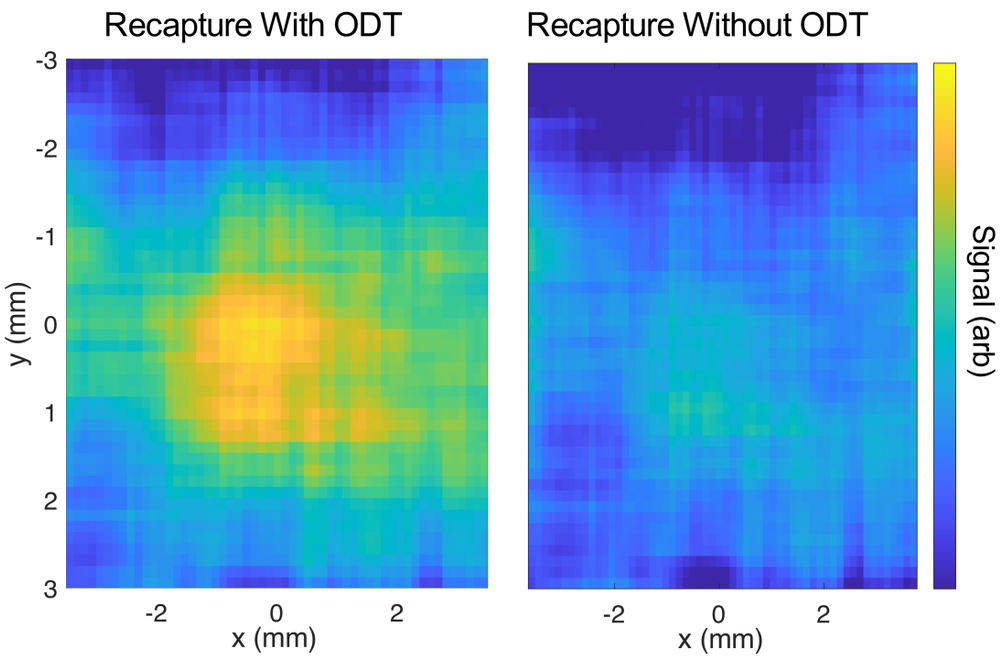

Figure 5.10: First ODT signal using MOT recapture method. Images show fluorescence from SrF molecules with (left) and without (right) the ODT beam, then released and recaptured in the MOT.

5.4.2 ODT loading time

After aligning the ODT to maximize signal, we first looked at the dependence of the ODT number on the loading time, see Figure 5.11. For this, we load the ODT for a variable time, and then

recapture the molecules in the RF MOT and image. We find that the loading is not very sensitive to the timing, and is optimized at around 150 ms, very similar to what the Doyle group saw [53]. Once the ODT + Λ-cooling light is turned on, the molecules slowly start diffusing into the ODT. The large size of the RF MOT ($\sigma \sim$ 1mm) makes the loading slow, but also makes the ODT insensitive to fine alignment of the ODT beam.

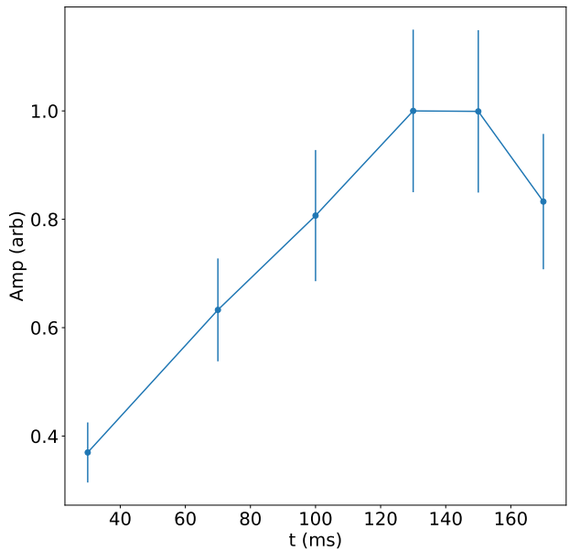

Figure 5.11: ODT loading using Λ-cooling. The loading is not very sensitive to timing, and is optimized around $t \sim 150$ ms.

For all the subsequent results, we fix the load time to 150 ms, as we observed no significant changes even after we revisited this after optimizing the other cooling parameters. Under fully optimal conditions, we can load around $\sim$ 4% of the molecules from the MOT into the ODT, meaning around $N \sim 160$ molecules.

5.4.2.1 Technical note on pockels cell

While performing ODT optimizations, we noticed that the optimal shim currents were fluctuating day-to-day. We then saw that this was due to polarization drift caused by the pockels cell that was being used for the polarization switching for the RF MOT. In particular, on stopping the polarization switching for the Λ-cooling, the polarization coming out of the pockels cell switched over the course of around 40 ms. We characterized this by monitoring the reflected and transmitted components after a PBS cube simultaneously. As shown in Figure 5.12, when the polarization switching is halted, most of the light is p-polarized, but over 40 ms, the polarization changes to s-polarized. While it is not clear what causes this drift, we account for this by setting the shim currents to minimze the width after 40 ms of cooling, instead of the usual 4 ms of cooling. This seems to solve the issue.

5.4.3 Temperature measurement

Our initial attempts at measuring the temperature by TOF proved to be very noisy due to the low molecule number and low signal from short exposures (we only want to expose for $\approx$ 1 ms to avoid artificial heating). We then decided to use the in situ ODT profile to determine the temperature. Here, we can image for far longer (around 100-150 ms) without heating the molecules, thus getting high signal to noise ratios. Ultimately, we also revisited TOF imaging with some tricks that showed that it was at least consistent with the in situ technique.

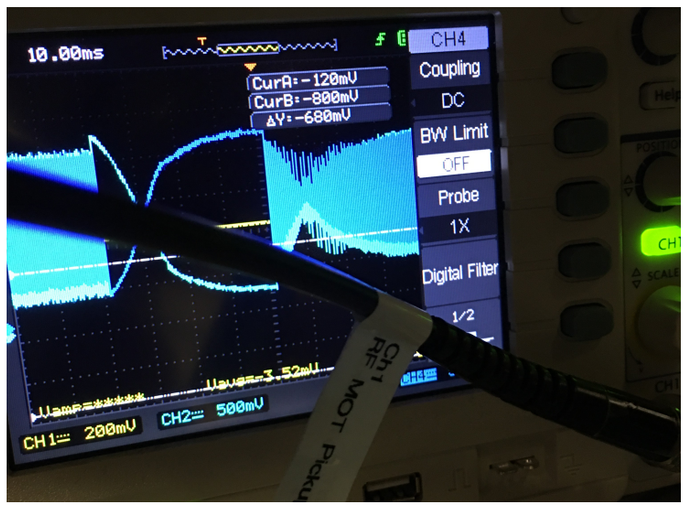

Figure 5.12: Pockels cell polarization drift. Oscilloscope traces vs time after the pockels cell switching is halted. Dark blue trace shows s-polarization while light blue shows p-polarization. Over around 40 ms, the polarization switches from p to s.

5.4.3.1 Temperature using trap profile

For this technique, we collect the fluorescence while the molecules are under the influence of the Λ-cooling light. With our current imaging system, we can only resolve the ODT profile along the long axial asis as shown in Figure 5.13. At the bottom of the trap, we can approximate the trap to be roughly harmonic, i.e. the trap potential is

$$U(r) = \frac{1}{2}m(\omega_x x^2 + \omega_y y^2 + \omega_z z^2) \tag{5.25}$$

Under such a potential, the density distribution is a Gaussian in all directions:

$$n(r) = n_0 \exp\left(-\frac{x^2}{2\sigma_x^2}\right) \exp\left(-\frac{y^2}{2\sigma_y^2}\right) \exp\left(-\frac{z^2}{2\sigma_z^2}\right) \tag{5.26}$$

with the cloud widths related to the temperature as $\sigma_i = \omega_i^{-1}\sqrt{k_B T/m}$, where ω_i are the trap frequencies. Along the axial direction, the trap frequency is $\omega_z = \sqrt{2U_T/mz_R^2}$. Thus, we see that the temperature can be determined from the axial width by the relation:

$$T = 2U_T \frac{\sigma_z^2}{z_R^2} \tag{5.27}$$

where the trap depth is expressed in K units.

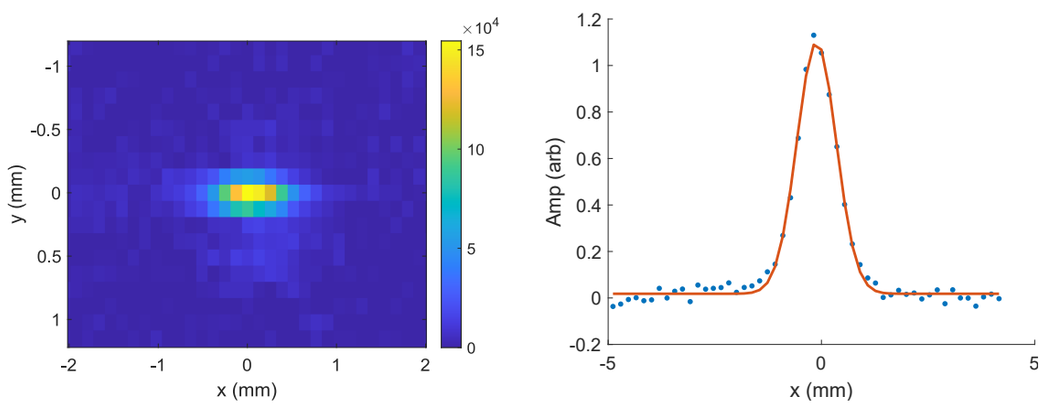

Figure 5.13: Best ODT image. (left) Heatmap of the ODT fluorescence from in situ Λ-imaging. (right) The integrated signal along the y axis is fit to a Gaussian and the ODT cloud size is extracted from the fit.

Our imaging system is at a 45° angle to the cloud and thus the actual width is $\sqrt{2}$x higher than the measured width. Thus,

$$T = 4U_T \frac{\sigma_{z,exp}^2}{z_R^2} \tag{5.28}$$

By integrating the signal along the short axis in the image, we can fit the resulting long axis density profile to a Gaussian, and extract the width, as shown in Figure 5.13.

We tried different ways to determine z_R. The first is to fit the bottom of the trap to a quadratic profile using the Rayleigh approximation, as shown in Figure 5.5. However, our trap has significant astigmatism, which causes the behavior to deviate significantly from an ideal trap. To that end, we also directly used the raw trap depth profile, and interpolated the trap profile using a 1D spline. Then, the density profile along the long axis is

$$n(z) \propto \exp\left(-\frac{U_T(z)}{T}\right) \tag{5.29}$$

This spline fit and the corresponding cloud density profiles for different trap temperatures are shown in Figure 5.14.

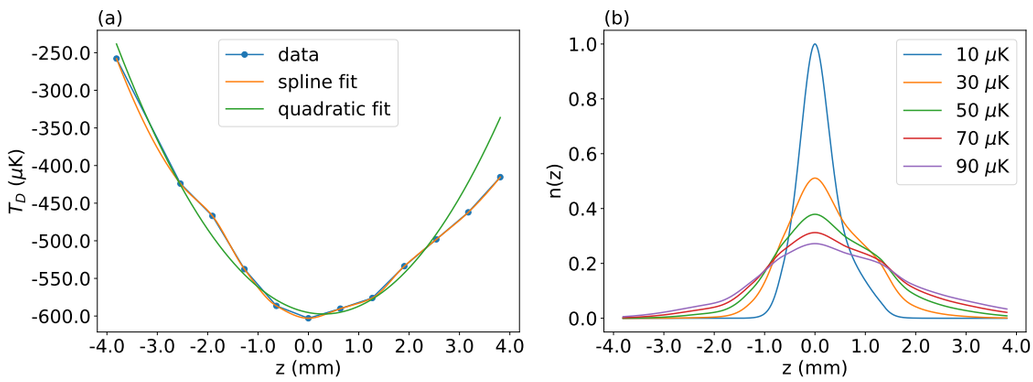

Figure 5.14: Trap and cloud profile using a 1D spline. (a) Different fits to the trap profile are shown. (b) Simulated cloud profiles determined by using the spline fit for different temperatures. For lower temperatures, the cloud still looks mostly Gaussian, while at higher temperatures, the shape deviates. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

For lower temperatures, the simulated cloud profiles are still roughly Gaussian, but there is significant deviation at higher temperatures. However, we may still fit a Gaussian to these simulated profiles, and then compare the widths to those of the in situ images to determine the temperature. We find that the temperatures as determined from the simple Rayleigh approximation are around 15% lower than those determined from the spline fit, not too big of a concern in this case. For a cloud of measured width $\sigma = 0.45$ mm, we find that the Rayleigh approximation gives $T = 16\,\mu$K whereas the spline fit gives $T = 19.5\,\mu$K.

5.4.3.2 TOF temperature

The analysis presented in this section was carried out by T. Langin. The initial problem we faced with TOF imaging was the poor signal to noise ratio due to the short imaging duration and the low molecule number. This made the images so noisy, that they could not be fit without smoothing the image. Fortunately, since we are fitting a gaussian to the image, we can always use a gaussian filter to smooth the image. The convolution of a gaussian with a gaussian is another gaussian,

with the resulting width

$$\sigma_{conv} = \sqrt{\sigma_{real}^2 + \sigma_{filt}^2} \tag{5.30}$$

Thus, in principle, the gaussian blurring should allow us to extract the real widths while reducing the noise. This is shown in Figure 5.15 where, on applying filters with progressively larger widths, the image becomes clearer, with a good filter width around 0.5 mm.

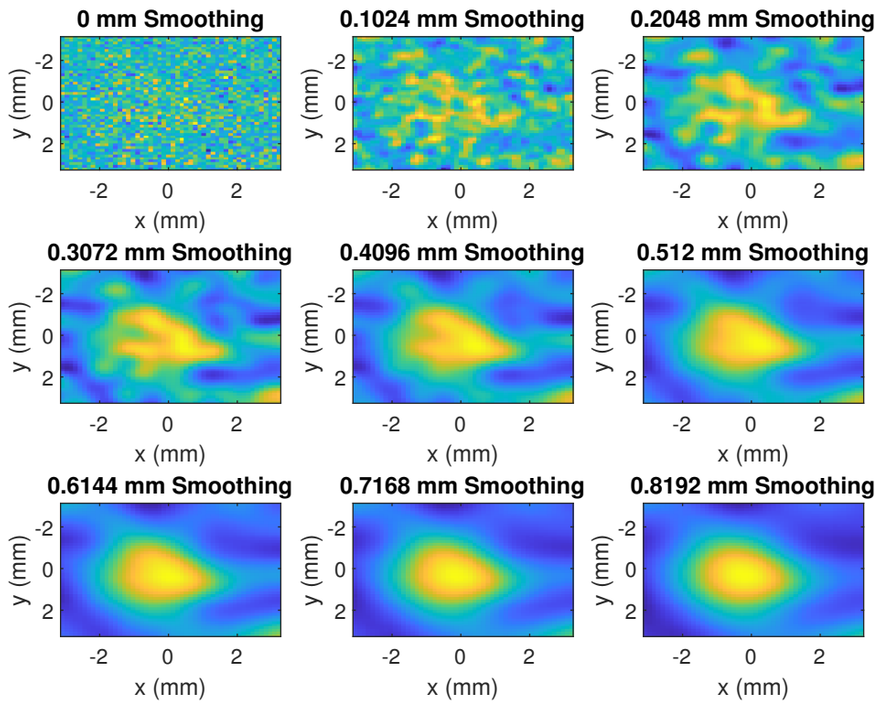

Figure 5.15: Example of smoothing the ODT TOF images. With no smoothing it is impossible to fit anything while on applying a around 0.5 mm filter allows us to fit the image.

The need for filtering also become less important the longer the cloud is allowed to fall for. We can use a maximum of 21 ms TOF before the cloud falls out of the imaging region, so we try this technique for two TOFs of 15 ms and 21 ms. To estimate the vertical temperature, we note that the initial vertical width should be only a few μms, so the temperature for a free expansion of t is:

$$T_{vert} = \frac{m}{k_B t^2}\sigma_{vert}^2 \tag{5.31}$$

For the horizontal axis, the initial width is non-zero, and is also measured at a 45° angle, so we get instead:

$$T_{hor} = \frac{m}{k_B t^2}\left(\sigma_{exp}^2 - \frac{1}{2}\sigma_{hor,0}^2\right) \tag{5.32}$$

We saw however, that there was still a lot of artificial blurring due to fluctuations in the cloud position itself. The second trick we used here was to take a lot of images for each TOF time, and create chunks of images that average only the minimum number of images that can fit a Gaussian. This is shown in Figure 5.16 where, for 2 images per chunk, we cannot fit reliably, but with 10 images per chunk we are able to fit easily.

Then, by fitting each chunk to a Gaussian and extracting the center positions, we can determine the fluctuations in the cloud centers, and try to account for them. Using this, we find that the standard deviation in the centers for the 15 ms TOFs are $(\sigma_{x,center}, \sigma_{y,center}) = (0.24, 0.27)$ mm

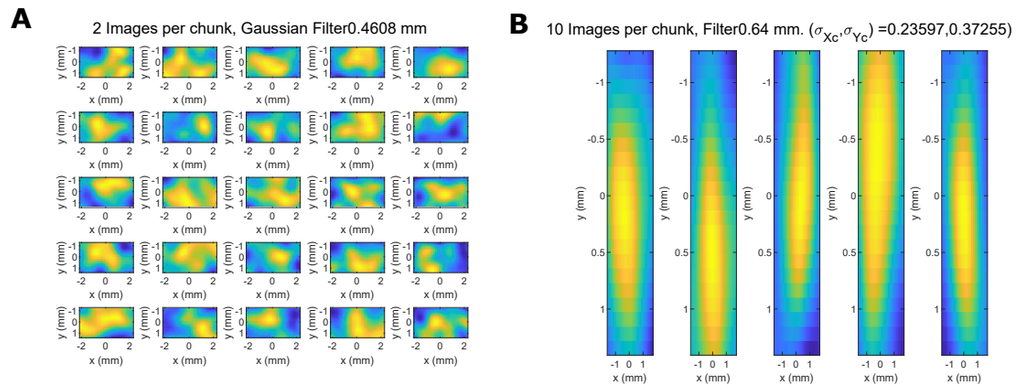

Figure 5.16: Example of chunking the ODT TOF images. (a) With only 2 images per chunk, we are not able to fit reliably, but (b) with 10 images per chunk, we can fit easily.

and for the 21 ms TOFs are $(\sigma_{x,center}, \sigma_{y,center}) = (0.21, 0.68)$ mm.

We also look at the blurring due to the imaging itself. We are imaging the cloud using the MOT light, with a scattering rate $R_s \approx 3 \times 10^6$ s^{-1}. This imparts an acceleration of $\dot{v} = \sqrt{R_s}\frac{\hbar k}{2m}\frac{1}{\sqrt{t}}$. Furthermore, the red detuning of the MOT light imparts sub-Doppler heating (which T. Langin simulated and found that $a = bv = 500v$ is the acceleration). Thus, we have that

$$\dot{v} = \frac{c}{2}\frac{1}{\sqrt{t}} + bv \rightarrow v(t) = \frac{c\sqrt{\pi}}{2\sqrt{b}}e^{bt}\text{erf}(\sqrt{bt}) \tag{5.33}$$

where $c = \sqrt{R_s}\hbar k/m$, assuming that the initial velocity is 0 (close enough for the ODT temperatures). Due to this heating, we can estimate the increase in cloud size by propagating the velocity over 1 ms, and find that $\sigma_{heat} = 0.14$ mm (the same, for 2 ms imaging, is $\sigma_{heat} = 0.57$ mm).

Finally, we estimate the blur from gravity as $\sigma_g = 0.5g \times (1\text{ ms})^2$ (how much the molecules have fallen in the 1 ms imaging time, which leads to a small smearing of the image) only in the vertical direction. This is 0.11 mm for 21 ms and 0.76 mm for 15 ms.

We can then subtract all of these from the fitted width in the same manner as the filter width, i.e. for example, $\sigma_{vert} = \sqrt{\sigma_{fit}^2 - \sigma_{y,center}^2 - \sigma_{heat}^2 - \sigma_g^2 - \sigma_{filt}^2}$.

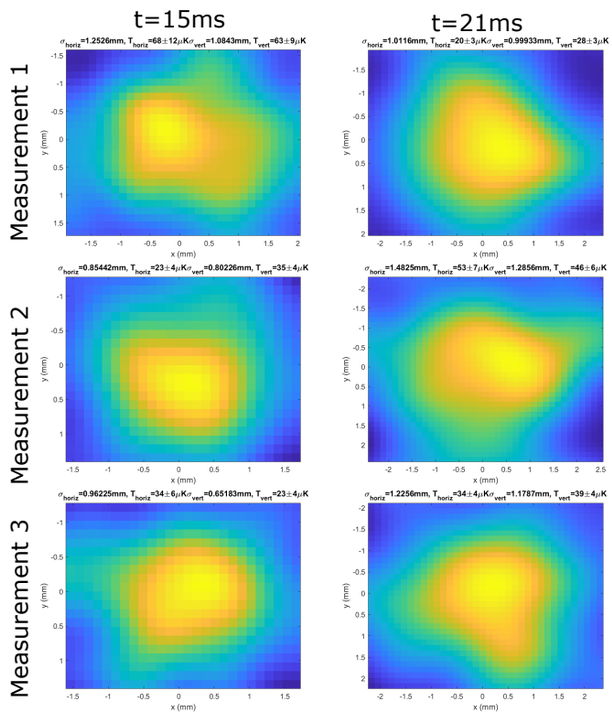

Figure 5.17: Example of TOF images using the full analysis. 15 ms and 21 ms TOF data is first blurred, then fit to a gaussian to extract the width.

The final analysis is shown in Figure 5.17 where different chunks are fitted using the above technique. With this method, we find that $T = 31 \pm 3.1\,\mu$K for the above images, for a day when the ODT temperature was around $T \approx 24\,\mu$K as determined by the first method. Since these are found to be consistent consistent within 25%, hence we consider that to be a typical systematic uncertainty in the temperature, and we resorted to using the first method since that is quicker and has less noise and assumptions associated with it.

5.4.4 Polarization dependence

As happens with a lot of new results, this one was also discovered somewhat accidentally. We had started the ODT optimization with linear polarization. However, we noticed one day that the ODT alignment had drifted, and in particular the polarization was not horizontal anymore, but slightly elliptical - and this was leading to better trapping! We were consistently getting at least a factor of 2 higher number and lower temperature at whatever this elliptical polarization was, see Figure 5.18. This was initially puzzling, until we dug deeper into the cooling mechanism we were using for the ODT.

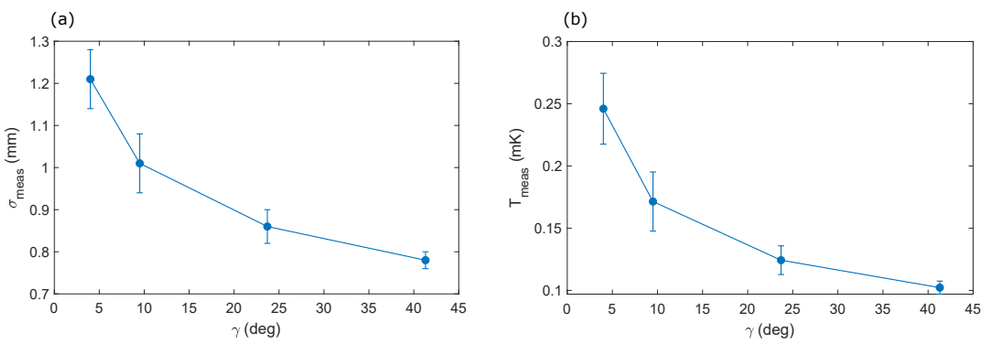

Figure 5.18: First hints of ellipticity dependence of ODT polarization. (a) Width and (b) Temperature vs ellipticity. Note that these initial temperatures are very large because we had estimated the trap depth incorrectly.

We are using Λ-cooling, which relies on the coupling of the $|F = 1\uparrow, 1\downarrow\rangle$ sublevels to make coherent dark states. Under the influence of the trap light, these states undergo differential AC Stark shifts which can destabilize the dark states, leading to worse cooling. However, there exist certain ODT polarizations, as shown in Figure 5.19, where the differential Stark shifts for these states are exactly identical. This leads one to believe that, for these ODT polarizations, Λ-cooling is extra effective, since the states coupled by the cooling light experience no shift inside or outside the trap, i.e. they can form coherent dark states inside and outside the trap. We also do not apply any external field to define a quantization axis - the shim coils are used to roughly cancel earth's field during Λ-cooling. Thus, the expectation was that the cooling should only depend on the ellipticity, and not the sense of rotation. Indeed, this was proposed as a means to compensate light shifts leading to loss of coherence in molecules [136].

Instead, we found that T can strongly depend on the rotation direction. As in § 5.3.3, we can define the ellipticity

$$\gamma_{ODT} = \frac{1}{2}\tan^{-1}\frac{S_3}{\sqrt{S_1^2 + S_2^2}} \tag{5.34}$$

where $S_{1,2,3}$ are the dimensionless Stokes parameters of the trap light. Here, we incorporate the sign to indicate the direction of rotation, with $\gamma_{ODT} > 0$ indicating clockwise when viewing along the light propagation axis. $\tan(\gamma_{ODT})$ is the ratio of the minor to major axis of the ellipse.

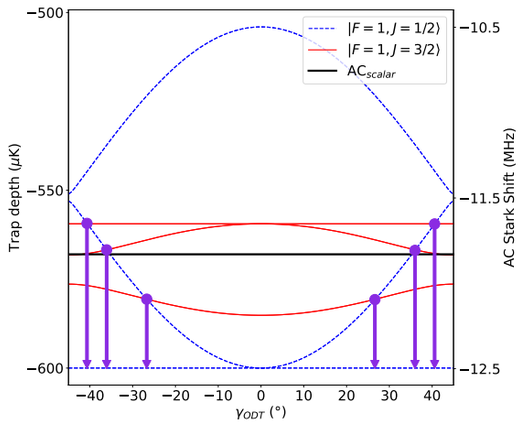

Figure 5.19: Expected optimal polarizations for in trap cooling in ODT. Curves of the same color are for the different superpositions of m_F sublevels corresponding to energy eigenstates. When the light shifts for two states coupled by the cooling light are equal, we would expect optimal cooling since the dark states would be maintained everywhere. This occurs at the six ellipticies as shown by the purple points. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

We found that the symmetry between positive and negative γ_{ODT} is broken, as seen in Figure 5.20. We tried a number of different configurations to figure out if this was indeed real - we tried loading the ODT at half power to reduce the Stark shifts, and we also tried to load at the maximum ODT power, but with lower cooling power to reduce scattering - however, in all cases, there was a distinctly prefered polarization.

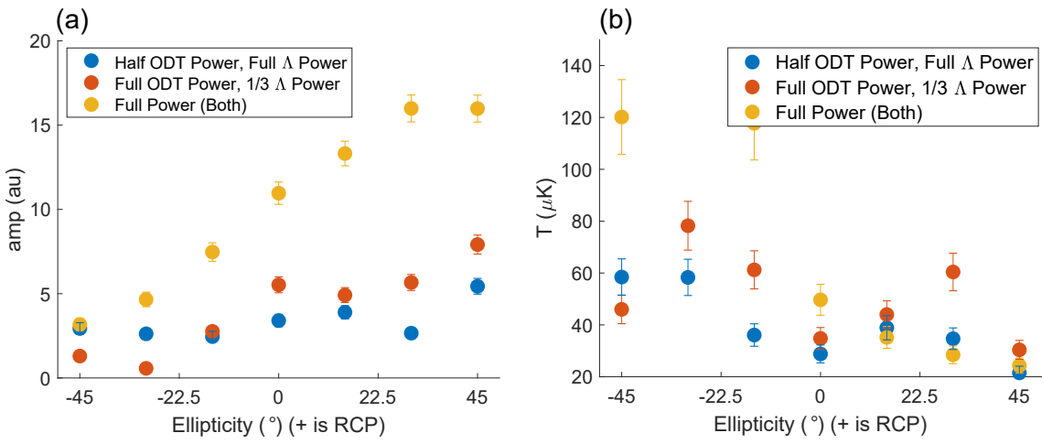

Figure 5.20: Evidence of broken symmetry between ellipticities. For all configurations of ODT power and/or Λ-cooling power, we found that there is always a preferred direction, in this case RCP light. (a) Number vs ellipticity and (b) Temperature vs ellipticity.

This was again quite puzzling, as there is, in principle, no preferred quantization axis but clearly a prefered polarization axis. Our guess was that something has to be creating a preferred axis, and we spent a lot of time trying to use the shim fields to affect the loading. We found that we could slightly improve the worse polarizations, but never really good enough to match the best polarization in Figure 5.20.

The only other degree of freedom that can break the symmetry is the Λ-cooling beam itself. In our setup, the ODT beam is co-propagating with the first/final pass of the MOT beams. As the MOT beam goes through all the optics, there is a substantial power loss, and thus there is a possibility of an intensity imbalance arising between the first and the final pass of the beam. To verify if this is the case, we measured the MOT beam profile after the first pass, and before the final pass as shown in Figure 5.21. Indeed, we found that the MOT beam sizes at the two points were more or less equal. However, there was around 26% power loss of the beam by the time the

beam came back for the final pass. This meant that, along the ODT path, the Λ-cooling beams had a roughly 26% intensity mismatch, and we quantify this by the quantity $R_{I\Lambda} = \frac{I_{weak}}{I_{strong}} = 0.74$ for this case.

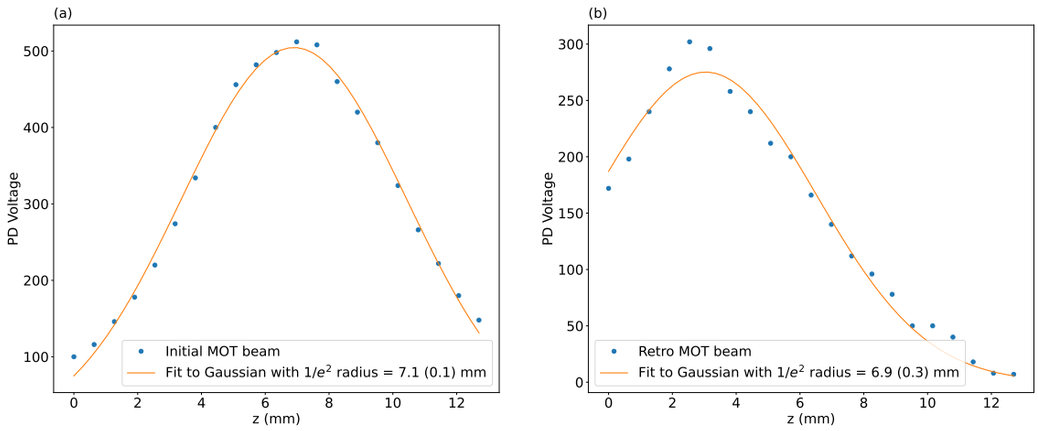

Figure 5.21: MOT beam profiles after the first pass and before the final pass. (a) MOT beam after the first pass with a $1/e^2$ radius of 7.1 mm and (b) MOT beam before the final pass with a $1/e^2$ radius of 6.9 mm.

We then saw that the trapping was optimized when the ODT polarization matched the polarization of the retro Λ-cooling beam, i.e. the weaker Λ-cooling beam. In Figure 5.20, the retro beam is σ^+ when it enters the chamber, and we see that the ODT polarization also wants to be σ^+. To verify if this intensity imbalance is indeed the cause of the broken symmetry, we had a simple knob to turn, we could just reverse the Λ-cooling polarizations, and then, if our hypothesis is true, the ODT polarization dependence should also reverse.

The effect of this broken symmetry is shown in Figure 5.22. Figure 5.22(a) shows the results for $R_{I\Lambda} = 0.74$ when the Λ beam co-propagating with the ODT is σ^+ polarized. We find that σ_{ax}, and thus T, is minimized when the ODT polarization matches the weaker, co-propagating beam ($\gamma_{ODT} = +45°$). This remains the case when the Λ polarizations are reversed Figure 5.22(b), where T is optimized when $\gamma_{ODT} = -45°$. If the intensity imbalance is reduced [$R_{I\Lambda} = 0.93$, Figure 5.22(c)], the dependence on γ_{ODT} is much less pronounced. We also do not see any dependence on the ODT polarization ellipse orientation angle $\psi = \frac{1}{2}\tan^{-1}(S_2/S_1)$.

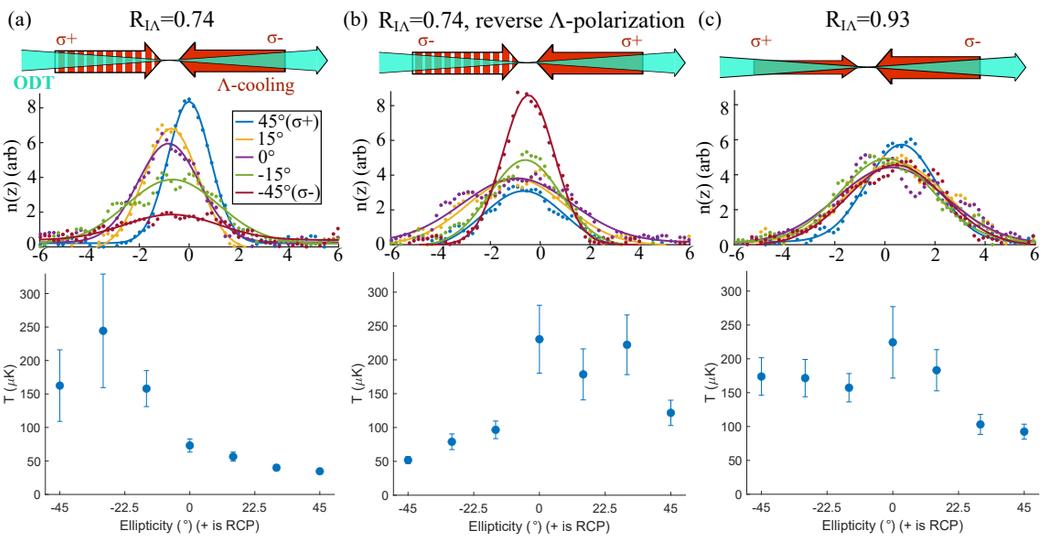

Figure 5.22: ODT profiles for different $R_{I\Lambda}$ and γ_{ODT}. (a) $R_{I\Lambda} = 0.74$ where the ODT is optimized for σ^+ polarization. (b) $R_{I\Lambda} = 0.74$ but with Λ polarizations reversed, and the ODT is optimized for σ^- polarization. (c) $R_{I\Lambda} = 0.93$ where there is no strong polarization dependence.

Ultimately, we find that T is globally minimized when the Λ beam intensities are deliberately

imbalanced — in particular, when $R_{I\Lambda} = 0.74$ and $\gamma_{ODT} = +45°$ and for the Λ polarizations in Figure 5.22(a), as shown in Figure 5.23. This configuration is used throughout the rest of this chapter.

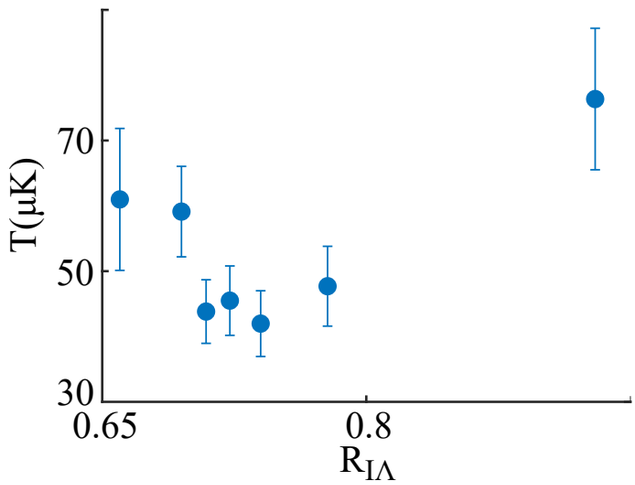

Figure 5.23: ODT temperature vs intensity imbalance. The temperature is roughly minimized for $R_{I\Lambda} = 0.74$, and for $\gamma_{ODT} = +45°$. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

Our first thought was that this could be arising because the Λ cooling beams are blue detuned from the transition, and could thus be inducing some AC Stark shifts that can counteract/aid the Stark shifts from the ODT light. To estimate whether the differential AC Stark shifts resulting from the Λ-cooling lasers could be large enough to compensate for those from the ODT, we solve for the AC Stark shifts resulting from a 1D pair of counter-propagating cross-circularly-polarized Λ beams with one photon detuning $\Delta/2\pi$ equivalent to the 'in-trap' value of 10 MHz, $\delta_R/2\pi = 1.2MHz$, $R_{1\uparrow,1\downarrow} = 2/3$, $R_{I\Lambda} = 0.74$, and the beam intensity (summed over both hyperfine addressing frequency components) of the initial (stronger) pass is 46 mW/cm^2, matching the experimental conditions. We found that $\langle 1\downarrow,-1|H_{AC}|1\downarrow,-1\rangle - \langle 1\downarrow,+1|H_{AC}|1\downarrow,+1\rangle = 2\pi \times 205$ kHz and $\langle 1\uparrow,-1|H_{AC}|1\uparrow,-1\rangle - \langle 1\uparrow,+1|H_{AC}|1\uparrow,+1\rangle = 2\pi \times 140$ kHz (m_F remains a good quantum number in the presence of only circularly polarized light). These differential shifts are comparable to the ones induced by the ODT when $\gamma_{ODT} = +45°$, and so could conceivably play a role in either mitigating or enhancing the effect of differential intra-hyperfine manifold shifts on the in-trap gray-molasses cooling.

We were however, not fully satisfied with this answer, and so we wanted to model this further. To understand this unanticipated behavior, T. Langin developed an OBE solver [67] that can incorporate intensity imbalanced, retroreflected beams. We explicitly add the AC Stark Hamiltonian from the ODT light (including vector and tensor shifts), while differential AC Stark shifts from the imbalanced Λ beams are included implicitly in the OBEs. This solver was benchmarked against results from comparable solvers [65, 67] and experimental observations, such as RF-MOT trap temperature [59] and capture velocity, Λ cooling [53], and single frequency cooling [54]. However, we were unable to reproduce the effects shown in Figure 5.22. The mechanism behind the observed interplay between ODT polarization and Λ-beam intensity imbalance thus remains an open question. We decided to take this win, and march on ahead.

5.4.5 Dependence on other cooling parameters

Next, we optimize the other cooling parameters to try and get the best loading and temperatures.

5.4.5.1 Temperature vs Raman detuning

We observe that T is optimized near the two-photon resonance ($\delta = 0$) with a broad minimum extending to $\delta > 0$, see Figure 5.24. This is also similar to what we observed in free space in Chapter 4.

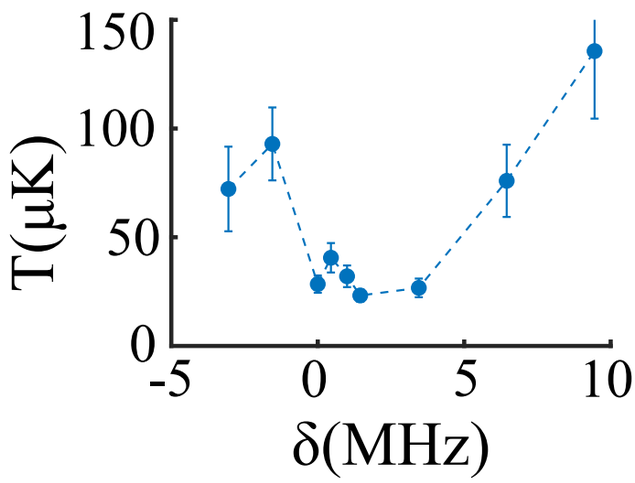

Figure 5.24: ODT temperature vs Raman detuning. Optimal cooling is observed around two-photon resonance, but has a broad minimum feature. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

Similar behavior has been observed in other experiments that use Λ-cooling [53, 98, 137]. The breadth of this feature is also comparable to the in-trap two-photon Rabi frequency between the coupled hyperfine manifolds ($\Omega_\Lambda \approx 2\pi \times 8$ MHz), as expected, see § 5.4.5.4. We use $\delta = 2\pi \times 1.2$ MHz throughout.

5.4.5.2 Temperature vs sideband ratio

Next, we optimize the temperature vs the intensity ratio between the two hyperfine levels, $R_{1\uparrow,1\downarrow}$, see Figure 5.25. We find that temperature is optimized for $R_{1\uparrow,1\downarrow} \approx 2/3$, a bit higher than the free space optimum. In particular, note that $R_{1\uparrow,1\downarrow} = 0$, i.e. single frequency cooling, is ineffective at cooling SrF in the ODT. Single frequency cooling was shown to cool to $\sim 10\,\mu$K in free space for CaF, and we also demonstrated similar temperatures in Chapter 4. However, it seems that the large Stark shits for the $|F = 2\rangle$ states lead to ineffective dark states, and thus poor cooling.

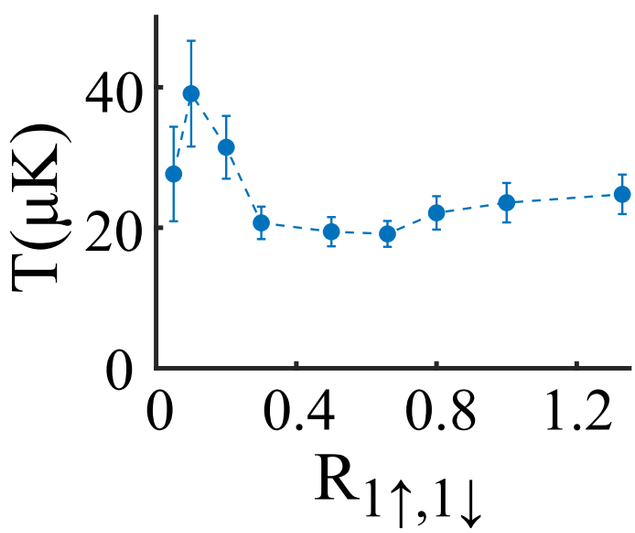

Figure 5.25: ODT temperature vs sideband ratio. Optimal cooling is observed around $R_{1\uparrow,1\downarrow} = 0.5$ but is not very sensitive at higher sideband ratios. Note that these values are a factor of 2 lower than the correct values, however, we have kept these to reflect our paper [111].

5.4.5.3 Temperature vs alignment

For optimal values of δ_R and $R_{1\uparrow,1\downarrow}$, we can regularly achieve $T \sim 20\,\mu$K. However, we also see that the ODT is very sensitive to the alignment of the Λ-cooling beam. In particular, we see that by adjusting the position of the retro mirror on our setup, we can dramatically make the ODT better or worse, an example of which is shown in Figure 5.26.

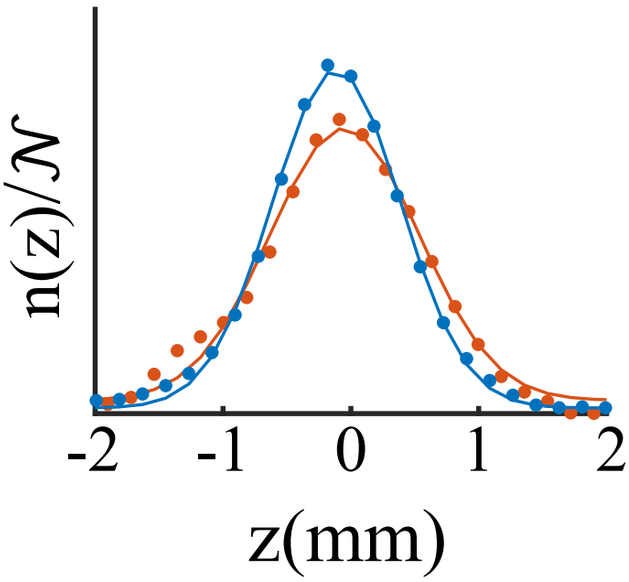

Figure 5.26: ODT temperature vs alignment. Optimal cooling, shown in blue, can be obtained by fine tuning of alignment of the Λ-cooling beams, but is very hard to maintain. Typical profiles are shown in red.

We then use a pseudo-deterministic procedure for optimizing trapping. The retro mirror is piezo actuated, and can achieve a minimum resolution of $0.7\,\mu$rad. We iteratively adjust the alignment by moving the retro mirror one tick and along one axis at a time, until a minimum is reached. It is not necessary that this will always be the global minimum, which is very hard to achieve. Figure 5.26 also shows the best temperature we could ever achieve, around $T \sim 14\,\mu$K. However, this can be maintained only for a few mintues after which the alignment tends to inevitably drift. This is also exacerbated by the fact that our MOT beam has to go over a path length of around 5-6 m and so even a small drift in any of the mirrors can make the alignment worse. The cooling beams pass through a lot of apertures as they propagate, and we do see some fringing on the beams by the time they have reached the retro mirror. The fine tuning of these fringes may be causing this dependence on alignment, but this is purely speculative.

The ODT beam alignment, not so surprisingly, does not seem to matter too much, as long as the beam is roughly centered on the MOT horizontally, and slightly lower vertically. The loading is also not super sensitive to the ODT focus position, which also makes sense because the Rayleigh length of 5 mm is much larger than the MOT cloud size of 1 mm.

5.4.5.4 ODT loading vs power

Next, we explored the ODT loading vs the ODT power. This is important to optimize the shape of the ODT beam, and to try something like a painted ODT for example. For this measurement, we use the MOT recapture method and image for 100 ms in the RF MOT. We find that the loaded number rises monotonically with trap depth, but appears to saturate for $U_T \gtrsim 500\,\mu$K, see Figure 5.27.

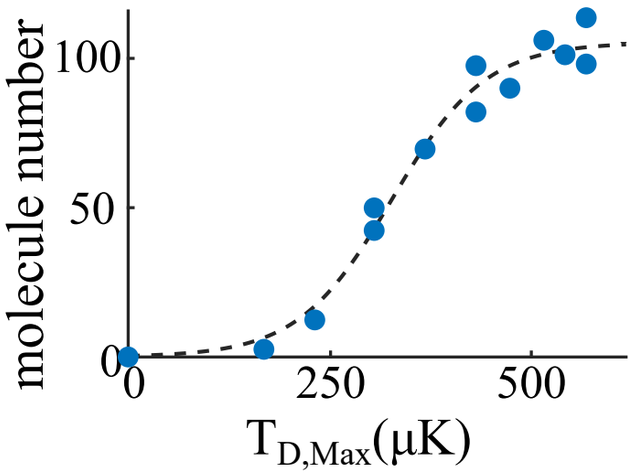

Figure 5.27: ODT loading vs trap depth. The data is taken at the optimal cooling parameters of $[\delta = 2\pi \times 1.2 \text{ MHz}, \Delta = 2\pi \times 22 \text{ MHz}, R_{1\uparrow,1\downarrow} = 2/3]$. The MOT number is typically $N \sim 3000$. The smooth curve is used to guide the eye.

This is strikingly different than what was observed in ODT loading of CaF molecules [53], where the efficiency peaked at $U_T \sim 130\,\mu$K. We suspect that this difference relates to the states chosen for Λ-cooling. We use the $|F = 1\uparrow, 1\downarrow\rangle$ states, which undergo J-mixing that modifies the transition strengths from the $\left|A^2\Pi_{1/2}, J' = 1/2, F' = 0, 1\right\rangle$ hyperfine states to these levels, see Table 4.2. In our case, $|F = 1\uparrow\rangle$ couples 56× more strongly to $|F' = 0\rangle$ than to $|F' = 1\rangle$. This has some important ramifications for optimal cooling, namely that the $|F = 1\uparrow, 1\downarrow\rangle$ are coupled through the $|A\Pi, F' = 0\rangle$ state, and not the $|A\Pi, F' = 1\rangle$ state. Consider the simple case when the two manifolds are coupled under the influence of σ^+ light. In this case, the branching ratios from the $|A\Pi, F'\rangle$ manifold to the $|X\Sigma, F\rangle$ manifold, $f_{F,F'}$ are shown in Figure 5.28. These branching ratios are proportional to the squares of the transition dipole matrix elements under σ^+ light.

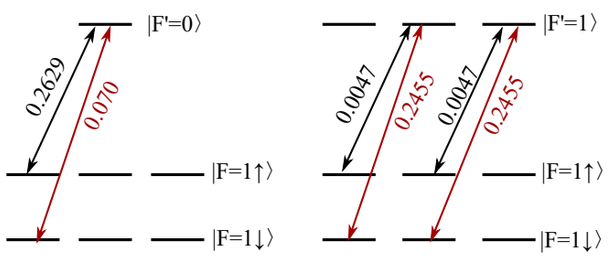

Figure 5.28: Branching ratios for the cooling transitions for σ^+ polarized light. Values mext to the arrows indicate $f_{F,F'}$. Left: $F' = 0$, Right: $F' = 1$

The two photon Rabi frequency, for the case when the amplitude through a single excited state F' dominates is given by:

$$\Omega_{F'} = \frac{\Gamma^2}{4\Delta}\sqrt{f_{1\uparrow,F'} f_{1\downarrow,F'}}\sqrt{\frac{I_{1\uparrow,F'} I_{1\downarrow,F'}}{I_{sat}^2}} \tag{5.35}$$

where $I_{F,F'}$ refers to the intensity of the light driving the transition between states $|X\Sigma, F\rangle \rightarrow |A\Pi, F'\rangle$, and $I_{sat} = \pi hc\Gamma/(3\lambda^3)$ is the saturation intensity ($\lambda = 663$ nm for this transition and $I_{sat} = 2.9$ mW/cm^2). Using the $f_{F,F'}$ values from Figure 5.28, the in-trap overall detuning $\Delta/2\pi = 10$ MHz, the total intensity ($I = 278$ mW/cm^2) and hyperfine ratio $R_{1\uparrow,1\downarrow} = 2/3$, we find

$$\frac{\Omega_0}{2\pi} = 8.5 \text{ MHz} \quad , \quad \frac{\Omega_1}{2\pi} = 2.1\text{MHz} \tag{5.36}$$

for each of the two excited manifolds. Thus, the main coupling is through the $|F' = 0\rangle$ state, and

we also see that $\Omega_0/2\pi$ is similar to the width of the Raman resonance width observed in the experiment.

The coupling through $|F' = 0\rangle$ enables robust cooling even at large trap depths as this state does not have any vector light shift as seen in Figure 5.3. In the case of CaF molecules, the use of $|F = 2, 1\downarrow\rangle$ necessitates coupling through the $|F' = 1\rangle$ state which can undergo large vector light shifts. Furthermore, the larger number of sublevels, each of which experience different light shifts may also limit the cooling efficiency.

The loading curve of Figure 5.27 does, however, mean that we cannot implement any sort of painted ODT without higher power. This is because, by sweeping the ODT, the average trap depth is reduced, which would in turn reduce the loading efficiency, thereby not really winning anything. If we can get twice the power in the beam, we can conceivably implement this.

5.4.6 Lifetime

The final quantity we measure is the lifetime of the molecules in the ODT cloud. This is important to understand collisional properties as well as to understand the limitations of the current vacuum system. Previously, in the MQT paper, our group had demonstrated a roughly 1 s lifetime [55]. To study collisions, it is necessary to have $1/\tau_{ODT} \sim \beta n$, where β is the collisional rate constant, n is the density and τ_{ODT} is the lifetime. We measure the lifetime by shuttering the Λ-cooling light for a variable time before reopening the shutter and imaging the remaining molecules. Single body loss can be modeled as:

$$\dot{N} = -\frac{1}{\tau_{ODT}} N \tag{5.37}$$

and the ODT lifetime is well described by $\tau = 910(200)$ ms, as seen in Figure 5.29. This is comparable to the lifetime measured in the MQT paper, and is hence believed to be mainly limited by collisions with background gas.

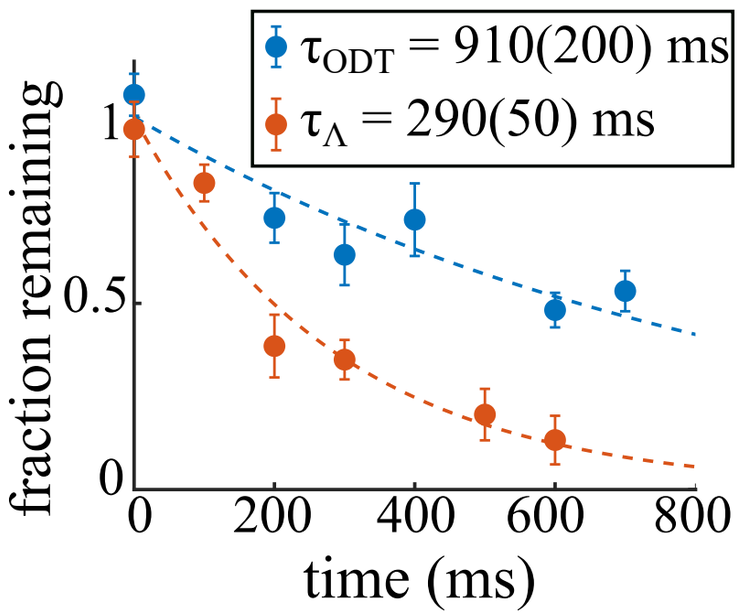

Figure 5.29: ODT lifetime with and without Λ-light. The lifetime in the presence of Λ-light is reduced, presumably due to light-assisted collisions.

It is known that the presence of the cooling light can induce additional collisions in the trap [53, 138–140]. This is because the atoms spend some time in the excited state, where they then interact with the dipole-dipole interaction and can gain significant energy before decaying back to

the ground state, thereby being ejected from the trap. To understand and verify if the Λ-cooling light has any adverse effects in our system, we measure the trap lifetime in the presence of the light. In our case, this also sets the time over which molecule loading is effective and over which in situ imaging can occur. To measure this lifetime (τ_Λ), we continuously apply the light for an additional $t_c = 650$ ms after the 150 ms loading and 50 ms release times, and image for 50 ms intervals during t_c.

As seen in Figure 5.29, we measure $\tau_\Lambda = 290(50)$ ms, similar to what was observed in an ODT of CaF [53]. In their cases, they had attributed the loss to spatial diffusion out of the trap induced by light scattering. They had a shallow trap ($\sim 130\,\mu$K), and so molecules could conceivably diffuse out. In our case, the trap depth is large, and Monte Carlo simulations indicate that spatial diffusion should contribute negligibly. This leads us to conclude that the main loss mechanism is probably light-assisted collisions, and indeed loss rate coefficients of $\beta \gtrsim 10^{-9}$ cm^3 s^{-1} have been observed in optical tweezer traps of molecules [141].

For applications where high fidelity detection is critical, such as studying single molecules in optical tweezers, it is important for the average number of photons emitted per molecule before loss to be large. The number of photons scattered is $\tau_\Lambda R_\Lambda$, where R_Λ is the scattering rate during Λ-cooling. We measure $R_\Lambda = 3.1 \times 10^5$ s^{-1} by comparing the fluorescence collected during Λ-imaging with that from the MOT recapture, where the scattering rate is known [59]. Thus, we find $\tau_\Lambda R_\Lambda = 9(2) \times 10^4$, ~ 3 times larger than demonstrated in an ODT of CaF [53] despite the similar τ_Λ. This is due to the larger R_Λ, which in turn is due to the smaller in-trap detuning - $\Delta_{trap}/\Gamma = 1.5$ in our case as compared to $\Delta_{trap}/\Gamma = 3.6$ in the CaF results [53].

5.5 Conclusions

ODT parameter	Optimized value
Molecule number N	160
Loading efficiency	4%
Trap temperature T	$14\,\mu$K
Peak density n	2×10^9 cm^{-3}
ODT Polarization	σ^+
ODT Power	50 W
ODT trap depth (U_T)	$600\,\mu$K
Intensity imbalance $R_{I\Lambda}$	0.74
Two-photon detuning δ	$2\pi \times 1.2$ MHz
One-photon detuning Δ	$2\pi \times 10$ MHz

Table 5.4: Optimized ODT parameters. Note that the temperature and trap depths are a factor of 2 lower than the correct values (density lower by $2^{3/2}$), however, we have kept these to reflect our paper [111].

In conclusion, by optimizing the combination of trap light polarization and intensity imbalance of Λ-cooling lasers, we have loaded $\sim 5\%$ of SrF molecules from a RF-MOT into a $\sim 600\,\mu$K deep ODT, at temperatures as low as $14(1)\,\mu$K. The large value of U_T/T implies strong compression, yielding density, and phase-space density higher than previously reported in bulk gases of directly cooled molecules, despite starting with 10 times fewer molecules. We find that several features of loading molecule ODTs using Λ cooling remain poorly understood, such as the observed interplay between the trap polarization and cooling light intensity imbalance and, more generally, the effect

of vector and tensor light shifts. Once these are better understood, higher trap compression may be achievable.

The next question we need to address is how to increase the molecule number in our traps, and also whether we can transfer more molecules into the ODT, so that we can observe collisions between the molecules. We can estimate the inelastic collision rate from a universal model where the short range dynamics are ignored. In this case, there is a free parameter, y, which describes the loss at short range, and $y \to 1$ corresponds to the universal limit with unity loss at short range [142–144]. The universal rate coefficient at zero temperature is given by

$$k_2^{\text{univ}}(T=0) = 8\hbar\frac{\bar{a}}{m} \tag{5.38}$$

[143]. where $\bar{a} = 0.477988 \times l_{vdW}$ and $l_{vdW} = (mC_6/\hbar^2)^{1/4}$.

For SrF in the ground rotational state, the C_6 coefficient is identical for all internal states and is given by [72]

$$C_6 = \frac{1}{(4\pi\epsilon_0)^2}\frac{d^4}{6B} \tag{5.39}$$

where d is the dipole moment and B is the rotational constant. For SrF, $d = 3.47$ D and $B = 2\pi\hbar \times 7.5$ GHz which gives $l_{vdW}^{SrF} \approx 561a_0$. This then gives us

$$k_2^{\text{univ}}(T=0, SrF) \approx 4.23 \times 10^{-10} \text{ cm}^3/\text{s} \tag{5.40}$$

At the current ODT density of $n \sim 2 \times 10^9$ cm^{-3}, we would expect a loss rate of 0.85 s^{-1} which corresponds to a lifetime of 1.2 s. If collisional loss is described by this universal loss rate coefficient, then, for the compression achieved here, a factor of 5 increase in N_{ODT} would lead to a collisional timescale $\tau_0 = (k_2 n)^{-1} \sim \tau_{ODT}$, allowing for the first studies of ultracold SrF molecule collisions.

To achieve this increase in density, we may try to increase the number captured in the MOT, which would directly translate to an increase in ODT number, or we may try to reduce the MOT size, which would increase the ODT loading efficiency. In the next two chatpers, we show how we can use both of these approaches to increase density and observe collisions.

Chapter 6

Sr + SF_6 source

6.1 Introduction

A very common and popular way of doing molecule experiments is to start with a cryogenic buffer gas beam source (CBGB) [145–160]. This method involves the ablation of a precursor target in a cryogenic cell (anywhere from 4-20 K). The ablation products are typically very hot (around 1000s of K) and these then collide with some buffer gas (Helium, Neon) which serves to both cool the products and also extract them from the cell.

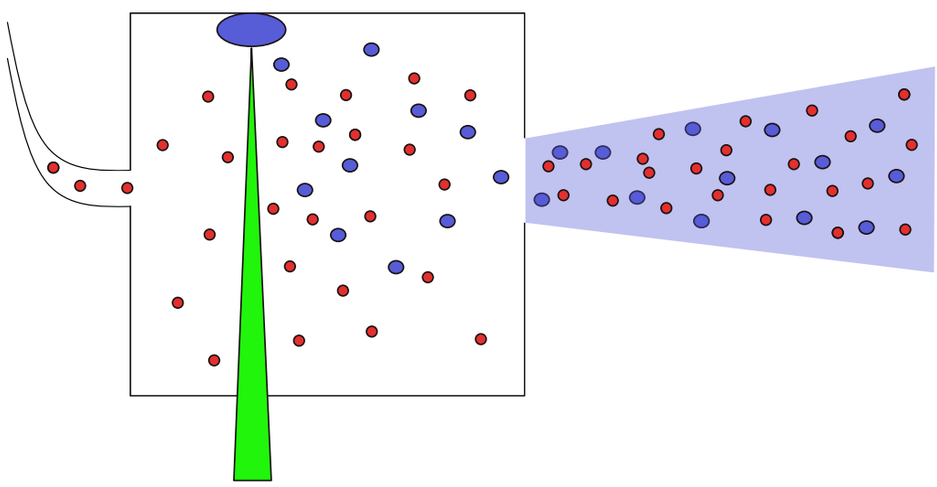

Figure 6.1: Diagram of a CBGB. The target (blue) is ablated, typically with a pulsed Nd:YAG laser (green), and the products collide with the buffer gas (red) to thermalize and become entrained in their flow, producing a molecular beam.

While the basic concepts of this system are understood, the actual dynamics are somewhat less clear, and so optimizing such a cryogenic cell is tricky and time consuming. However, there are several clearly articulated observations and suggestions in the literature [146, 161–164], which include:

1. The cell length should be long enough to ensure a sufficient number of collisions with the buffer gas occur before the molecules exit the cell. This ensures that the beam coming out is fully thermalized with the buffer gas. A typical feature of athermal beams are a lot of high-velocity molecules, that show up as long tails in the longitudinal velocity distribution [146].

2. Adding a diffuser in the back of the cell to diffuse the buffer gas as it enters the cell may be important. This ensures that the buffer gas is spread throughout the cell, and leads to more uniform extraction [165].

3. The target should not be protruding too much into the cell. If the target is too far in, it interferes with the buffer gas flow and leads to very poor extraction. Ideally, the target should be flush with the cell walls, or even a bit recessed.

4. Cell designs with different central bore sizes exist in the literature. While there is not a clear dependence of the molecule number on this, it is clear that the larger the bore size, the more buffer gas flow is needed to get maximum molecular output. The larger the bore, the larger the temporal extent of the molecule beam extracted from the cell [146].

5. Some modern buffer gas sources also make use of in-cell chemistry to produce the desired molecules (e.g. ablation of a pure Sr/Ca metal target and reacting it with SF_6 gas to produce SrF/CaF). In such a case, different designs exist that introduce the reactive gas at different spots, and again it is not clear whether any are better.

6. Adding a second stage to the cell leads to a clear reduction in the forward velocity distribution. It also leads to a lower molecule flux by a factor of $\sim$10 [165]. Thus for molecule cooling and trapping experiments, it is preferred not to have a second stage, since such experiments are hurting for molecule number. However, for precision measurement experiments, a second stage is used since it leads to longer interaction times.

With this in mind, we decided to try and optimize our beam source to produce a better and consistent beam of SrF molecules than we had typically achieved over the prior years of work with SrF in our lab.

6.2 Characterizing the old beam source

In this section, we will look at the characteristics of the old beam source and some of the problems associated with it.

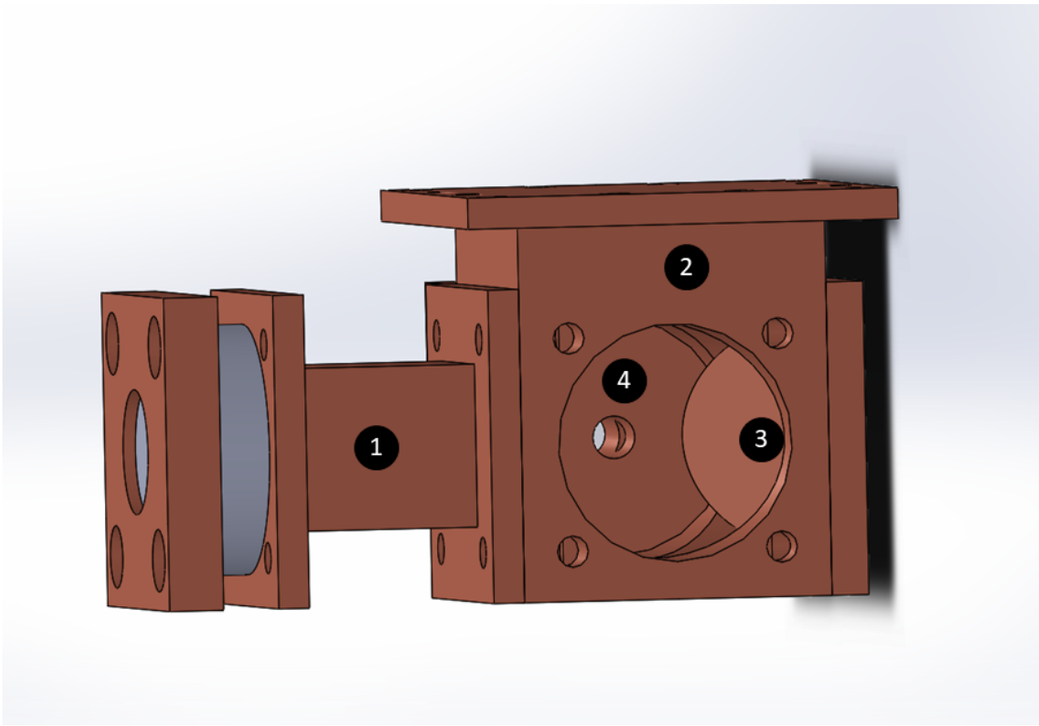

Figure 6.2: Old cell design. (1) Snorkel. (2) Cell. (3) Angled target piece. (4) Old Helium inlet which is just a single hole in the copper piece.

6.2.1 Ceramic targets

The original version of the SrF experiment used pressed and sintered SrF_2 targets. In some of the past theses, it is claimed that the old ceramic SrF_2 targets lasted as much as 10^6 shots per target or more, with around 2000 usable shots per ablation spot on the target. However, in recent times, the target quality has decreased by a lot, and ceramic targets are typically usable for only 2-3 weeks of continuous running before needing to be replaced. The newer ceramic targets lasted only around 100-200 shots per spot, and we needed to frequently tune the ablation spot. Furthermore, replacing a target was typically a two day cycle, thus leading to longer down times. The procedure for making ceramic targets is fairly complicated and often not reproducible, which makes it hard to reliably make good targets.

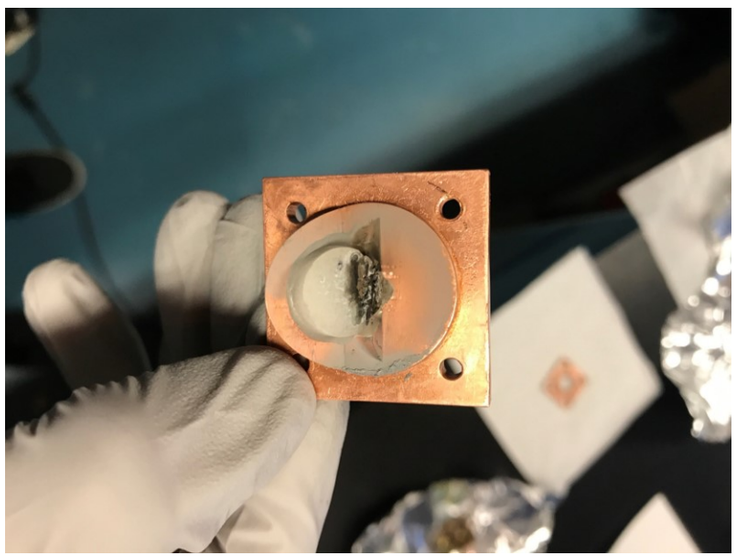

Figure 6.3: Usable portion of a typical ceramic target destroyed within 2-3 weeks of continuous use. Note the angled target piece used in the prior version of the source.

The Doyle and Tarbutt groups have shown that metal + reactive gas targets are much more

stable and last at least an order of magnitude time longer than ceramic targets [165]. Pure metal targets can be directly acquired without need for in-house fabrication. Thus, we concluded that it would benefit us as well to switch to a metal target + reactive gas scheme.

6.2.2 Angled target piece

As can be seen from Figure 6.3, the target was glued onto an angled piece and sits in the cell at an angle.This design was introduced very early in the development of ablation-loaded CBGB sources, prior to improved understanding of the importance of buffer gas flow patterns in the cell. No other CBGB source we have seen uses this kind of design. Since the angled part of the copper piece protrudes into the cell, it likely serves to disrupt some of the Helium flow in the cell, which could lead to poorer performance of the source. Because this design was so unique and did not give any obvious benefits, we decided to move to a more standard geometry where the ablation target is basically flush with the interior walls of the cell.

6.2.3 Forward velocity profile of the molecular beam

This is the most important aspect of the source. We desire the source to produce as high a molecular flux as possible, with as low a forward velocity as possible, in order to slow and capture the most molecules in a MOT. The velocity profile was measured using the technique from [75]. The detection scheme is a hyperfine free measurement using a 2-photon process that leads to near background-less detection. The typical result for the old ceramic source is shown in Figure 6.4. The velocity distribution is peaked around 140 m/s and has a long tail extending to higher velocities (> 200 m/s) which indicates incomplete thermalization in the cell. Our qualitative understanding is that the faster molecules are extracted quicker than the slower ones, and do not get enough time to collide with the buffer gas. We concluded that this could potentially be fixed by using a longer cell.

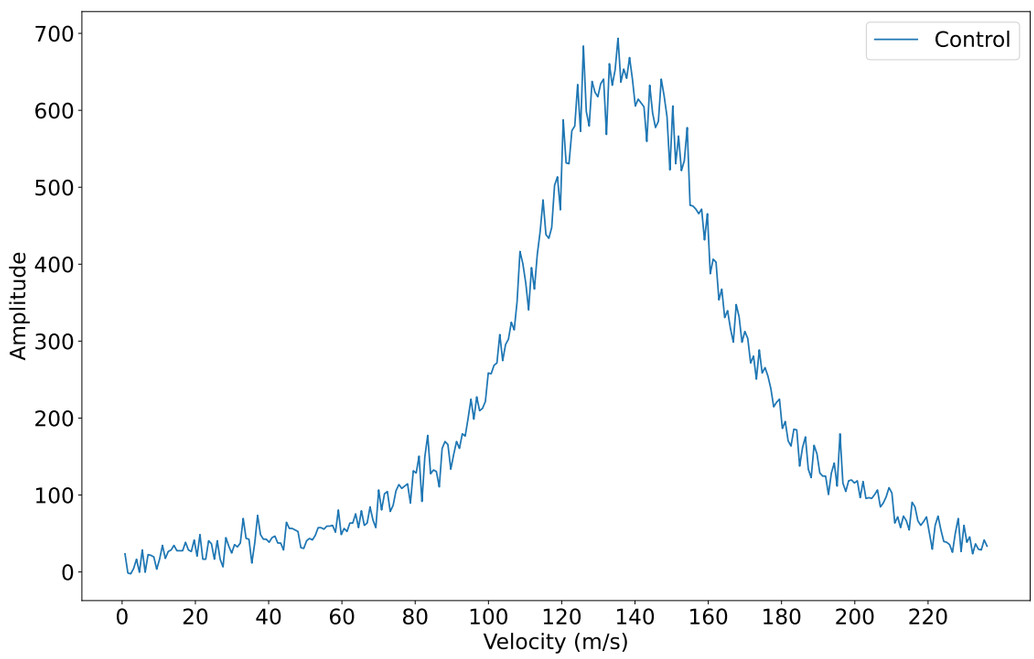

Figure 6.4: Velocity distribution of the old ceramic source. It is peaked around 140 m/s, and it has a long tail extending to higher velocities, indicating incomplete thermalization in the cell.

6.2.4 Helium inlet and cell length

The old design used a copper piece with a single hole in the center as the Helium inlet (see Figure 6.2) and did not have a diffuser. There is some evidence that such a design leads to inefficient extraction since the Helium does not spread evenly inside the cell. We concluded that this could potentially be fixed by adding a diffuser at the back.

6.2.5 Absorption outside the cell

To characterize performance of the old source, we directed an absorption laser beam (on the $|X\Sigma, N = 0, J = 1/2\rangle \rightarrow |A\Pi, J = 1/2^-\rangle$ transition) immediately outside the cell to monitor how many (if any) molecules are being produced in the beam. This is not the best measure, since the number of molecules produced in $N = 0$ may not always be in proportion to the number produced in $N = 1$. We found that the absorption depends a fair bit on the ablation spot as well as the Helium flow rate, but in general, higher absorption outside the cell has led to higher numbers in the MOT. Typically, for the best MOT signals with the old source, we saw around 2-3% absorption just outside the cell, but we also decent MOTs with as low as 1% absorption.

6.3 New designs for source improvement

6.3.1 Switching to a Sr + SF_6 source

This is the biggest change in the new design. Apart from the long target lifetimes reported for the earlier iterations of ceramic targets, all recent SrF_2 targets (as well as ceramic targets for other molecules such as CaF_2 and BaF_2) have very short lifetimes. Numerous groups have seen that switching to a metal + reactant source leads to a more stable and longer lived source. Thus it would be good for the experiment to make this switch as well.

There are a few design considerations for such a switch. First, SF_6 freezes at around 150 K in vacuum, whereas the cell is at 4 K. Thus, there needs to be excellent thermal insulation between the SF_6 line and anything else in the source. Second, the molecule beam properties depend sensitively on the temperature of the cell, with a hotter cell in general producing fewer and faster molecules. Thus, there also needs to be excellent thermal insulation at the point where the SF_6 tube enters the cell. Third, there should not be any additional holes in the cell from where Helium can leak out, since we found (and others reported [165]) that this greatly degrades the molecule beam properties. Lastly, Sr is highly reactive in air and hence there needs to be a robust procedure for attaching the target to the cell before it fully reacts.

6.3.1.1 SF_6 tube thermal isolation

To maintain the tube above 150 K througout the source, we need to make sure it does not come in contact with any of the thermal shields. The temperature will be mainly limited by radiative cooling, and so keeping the length as short as possible is vital. We also wanted to use copper since it is easy to work with, can be heated remotely due to its good thermal conductivity, and one can get small diameter tubes quite easily. To keep the thermal conductivity only moderately high, we chose 122 copper and a 1/16" OD tube (McMaster 7190K71).

The SF_6 line goes through a MKS gas flow controller, and through a KF50 gas feedthrough with Swagelok connections on either end. Inside the chamber, the 1/16" Cu tube is connected with a couple of Swagelok adapters to the feedthrough. The tube is then bent through a large hole in the 30 K shield (see Figure 6.13), and passed through large holes already present on the sides of the 4 K shield. This way, the tube is not in direct contact with anything colder than its desired temperature. Finally, the tube is bent under the cell and plugged into a Vespel piece that acts as a relatively gas-tight thermal insulator between the tube and the cell (see Figure 6.6). Vespel is a very poor thermal conductor, and is the best we could find to introduce the SF6 into the cell without heating the cell too much, and also leaving no gaps. This is important because we want to keep the Helium sealed in the cell, since any leaks will cause worse molecule extraction. The tube is inserted into the Vespel piece until it is just about flush with the inner edge of the cell. We also added a couple of Teflon washers under the screws on the Vespel piece to add further insulation.

A few Platinum RTD elements (Omega 1PT100KN1515) were epoxied using Arctic Silver at different places along the SF6 line to keep track of the tube temperature. With the current tube geometry, we see a pretty smooth gradient from room temperature outside the shields, down to about 157 K at the mouth of the Vespel piece. This is enough to keep the SF6 flowing, while making little difference in ultimate cell temperature (4K shield temperature without the tube was around 3.4 K, and with the tube was around 3.5 K). Hence, we concluded that further thermal insulation need not be done.

6.3.2 Increasing cell length:

As mentioned before, it has been reported by the Doyle group [165] that a longer source length allows for more collisions with the buffer gas and leads to a slower molecule beam. One has to be careful while trying to increase the length because the distance from the front of the cell to the collimator plate is critical, and should not be changed [76]. Our 4K shield is also pretty small, and one needs to take into account space constraints and logistics of target changing, but it can be done with our geometric constraints as shown in the cell design in Figure 6.5. The additional 1/4" length was expected to give more time for buffer gas collisions and potentially lead to produce an overall slower beam.

Figure 6.5: (Left) Longer cell with 0.25" extra length. (Right) Current cell design.

6.3.3 Adding a diffuser + back windows:

The Doyle group also reported [165] that adding a diffuser to the back of their cell to convert the directional buffer gas into diffuse buffer gas led to an increased number of collisions and thus slower molecules. We also have the added constraint of requiring a window in the back of the cell for the retro-reflected slowing beam repumpers and also the push beam in the future (see Chapter 9). We accomplish this as shown in Figure 6.6, which shows a new design incorporating the diffuser and the longer cell length.

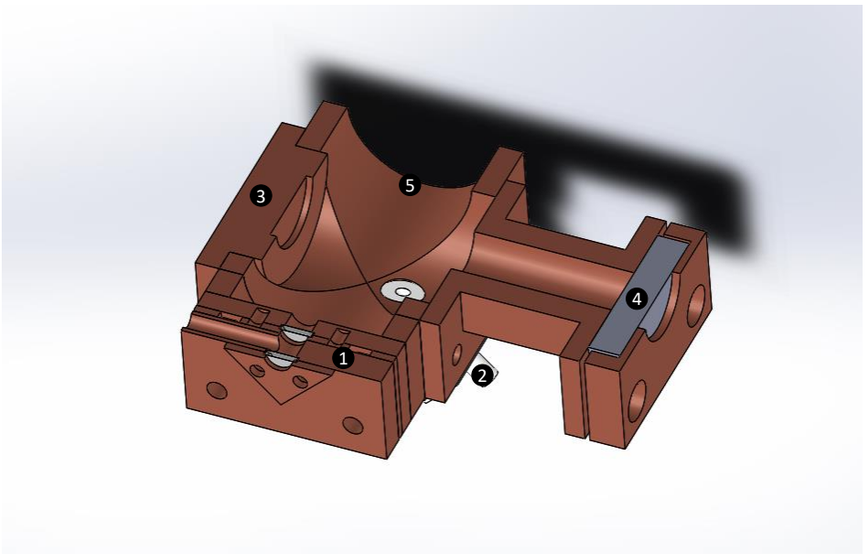

Figure 6.6: Cross section of new cell design with longer length and diffuser. (1) Helium inlet with diffuser. Helium enters from the side in between the two windows and is then diffused. (2) SF6 inlet. (3) Target. (4) Snorkel. (5) Cell exit. The addition of the diffuser and longer cell length should result in slower molecules.

6.3.4 Flat target piece

The old target piece had an angle and for some reason, we could get good absorption signal and a good MOT signal only from ablation spots closer to the front edge of the target (see Figure 6.3). I think this is because the front edge of the target is just at the right distance inside the cell to make good molecules, but the rest is not, and also the angle obstructs Helium from reaching the target spots very close to the upper edge of the angle. We designed a new flat piece which would place the entire target at this "good" distance inside the cell as shown below.

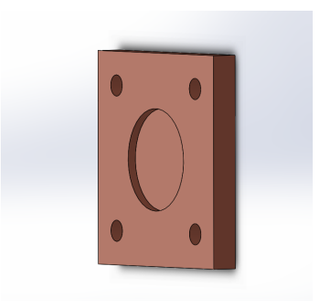

Figure 6.7: New flat target piece designed such the target (4 mm thickness) will end up at the correct distance inside the cell.

Table 6.1 shows the differences between the old and new designs.

	Old Design	New design
Cell length	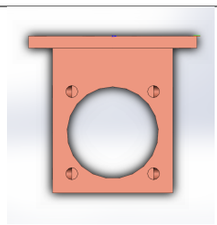	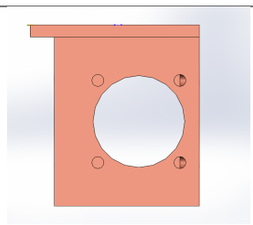
Target holder piece	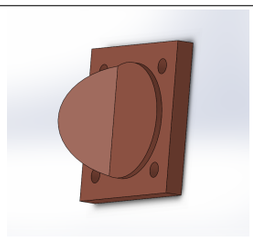	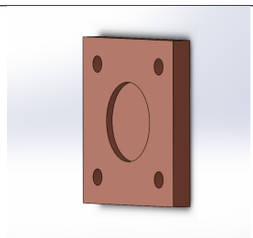
Cell back	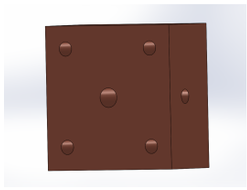	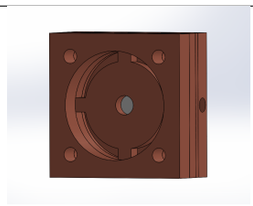
Ablation target	SrF_2	Sr + SF_6

Table 6.1: Old cell design vs new cell design

6.4 New source design characteristics

6.4.1 Testing with a ceramic target

We first decided to test the new source design with a ceramic target to make sure everything works as expected before putting in the Sr. With this flat target piece, we were able to get a similar level of absorption outside the cell as with the old design, indicating that at least the new source is not somehow worse.

Next, we set out to measure the velocity distribution of the new source. This measurement is shown in Figure 6.8. As expected, the new source has a peak at a slower velocity (because of the longer cell length), and does not have a long tail at higher velocities, which means thermalization with the buffer gas is effective. This data was taken at the same He flow rate as the old source (around 6.5 sccm).

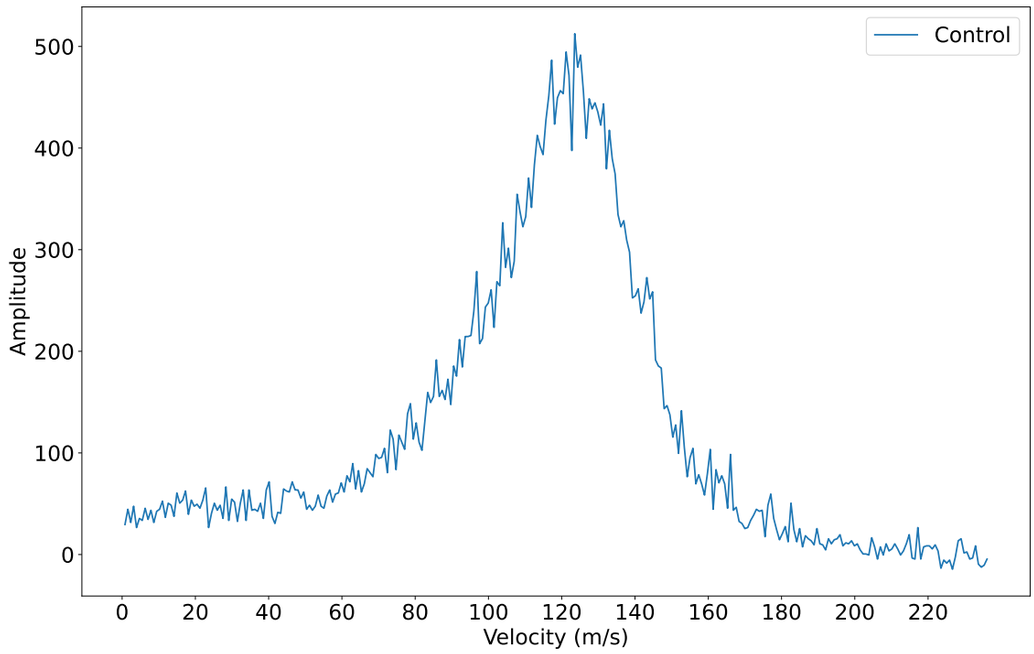

Figure 6.8: Velocity distribution of the new source with a ceramic target. It is peaked around 120 m/s, and it does not have a long tail extending to higher velocities, indicating complete thermalization in the cell.

These are good indicators that the new source was working well, so we decided to swap out the ceramic target for a Sr target with SF6 gas.

6.4.2 Sr + SF_6 target

Since this was a substantial change in the beam source, the first thing to check was if we are producing any molecules at all by looking at absorption outside the cell. After a bit of tweaking different parameters, we were able to find absorption outside the cell and then optimize this signal. After optimization, we observed upwards of 10% absorption outside the cell, which is around 5x larger than the old source (see Figure 6.9). This data was all taken at around 7 sccm He flow rate and around 0.04 sccm SF_6 flow rate as measured by the gas flow controllers.

The absorption signal takes some time (around 5 minutes) to saturate after ablation for the first time. This seems to primarily be because the SF_6 takes some time to diffuse in the cell and freeze on the walls and the target, after which point a new equilibrium is reached and the absorption saturates. Decreasing the SF_6 flow rate after this point does not do anything immediately, but after a while the signal starts decreasing, presumably because the frozen SF_6 starts getting used up. Increasing the SF_6 flow rate beyond 0.1 sccm however drastically reduces the absorption signal, presumably because the additional SF_6 gas heats up the molecules. Ultimately, the optimal ablation parameters will be determined by the MOT signal, but optimizing for absorption just outside the cell, we find the optimum He flow rate to be around 7-8 sccm and the SF_6 flow rate around 0.04-0.06 sccm.

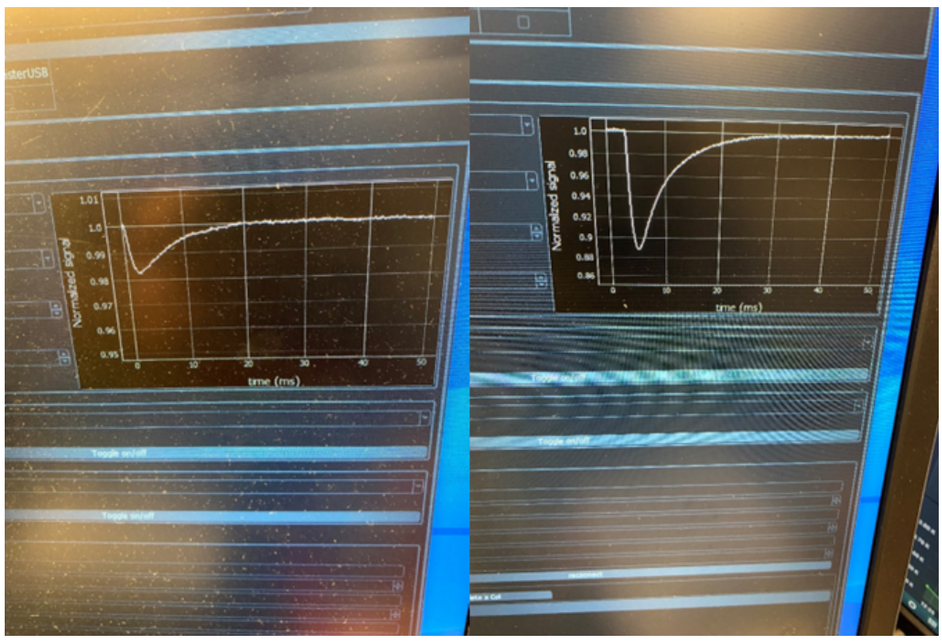

Figure 6.9: Comparison of absorption outside the cell between the old and new designs. (Left) Old source with ceramic target, typically gave around 2% absorption outside the cell. (Right) New source with Sr + SF_6 target, can give upwards of 10% absorption outside the cell.

6.4.3 Velocity distribution and optimization of the Sr + SF_6 source

Here, we will characterize the properties of the new source. In this section, all the data were taken using the slowing detection scheme outlined before, except without any slowing light. This gives us the bare velocity distribution from the source.

6.4.3.1 Velocity distribution versus He flow rate

Here, the ablation laser (YAG) power is fixed to around 8.5 mJ/pulse and the velocity distribution is measured as a function of the Helium flow rate. As expected, the total number of molecules extracted from the cell increases with the flow rate up to a certain point, beyond which more molecules are extracted but at faster velocities. This can be seen from Figure 6.10 where the overall curve shifts to higher velocities and more yield up to around 9 sccm, after which any extra molecules extracted are predominantly at very high velocities above 150-160 m/s. This higher flow rate compared to the older source can possibly be understood because the new cell is longer, and we also have a diffuser, which presumably reduces the effective flow rate in the cell. Thus there is a trade-off between number of molecules extracted and their velocity vs He flow rate, and at around 9 sccm He flow, we get the maximum number of molecules at around 140 m/s, which suits our slowing parameters very well.

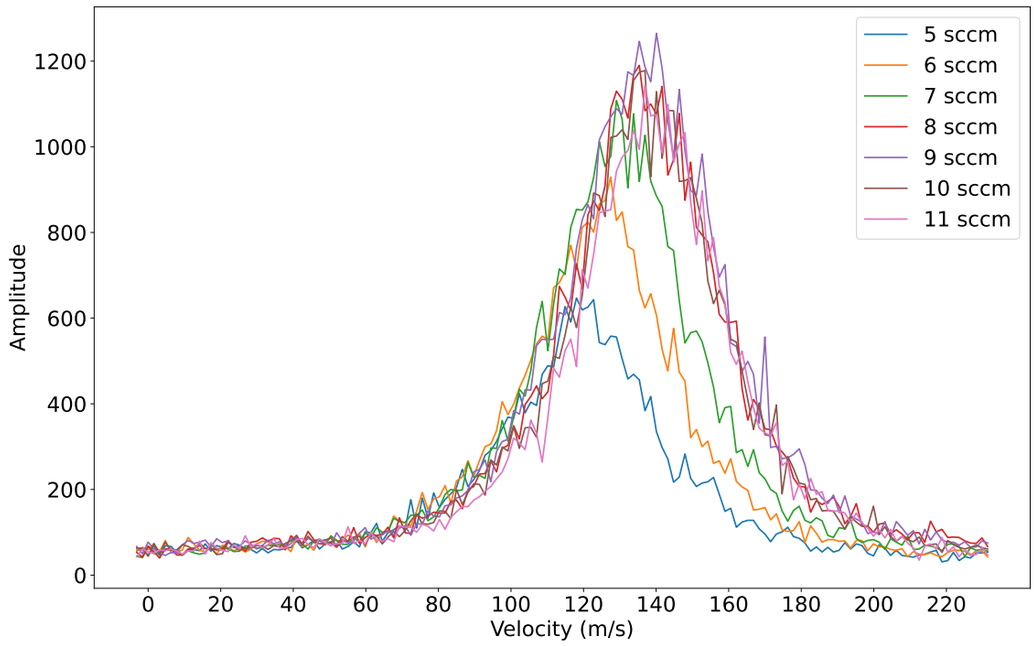

Figure 6.10: Velocity distribution vs Helium flow rate for the new source. The best extraction occurs around 9 sccm He flow.

6.4.3.2 Velocity distribution versus YAG power

Here, the He flow rate is fixed around 9 sccm, and the velocity distribution is measured as a function of the YAG power (see Figure 6.11). There is a sharp threshold effect around 5 mJ/pulse above which the SrF beam intensity increases abruptly. This may be due to an onset of reactions between ablated Sr and the SF_6 gas, or to a threshold in yield of ablated Sr.

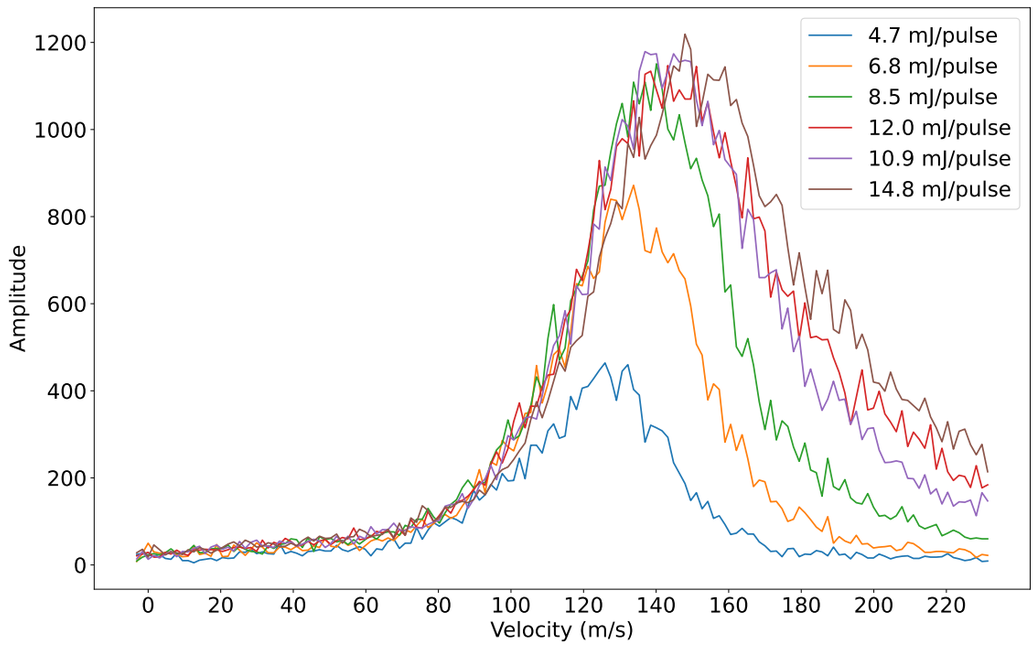

Figure 6.11: Velocity distribution vs YAG power. The best extraction occurs around 9 mJ/pulse YAG energy

Beyond this point, the extracted number increases with ablation energy upto around 9 mJ/pulse, where the velocity distribution is still narrow. Above this YAG energy, the additional molecules are again seen only in the very high velocity classes above 150-160 m/s. The excess YAG energy only serves to make everything faster and does not contribute to any useful molecules.

6.4.3.3 Optimal parameters for the new source

Figure 6.12 shows the velocity distribution from the source under optimal conditions, which are: He flow rate around 9 sccm, YAG power around 9 mJ/pulse, SF_6 flow rate around 0.05 sccm, SF_6 tube temperature around 145 K. With these parameters, we can consistently get 10% or more absorption outside the cell, and the velocity distribution is narrow and peaked around 140 m/s, thus indicating complete thermalization in the cell.

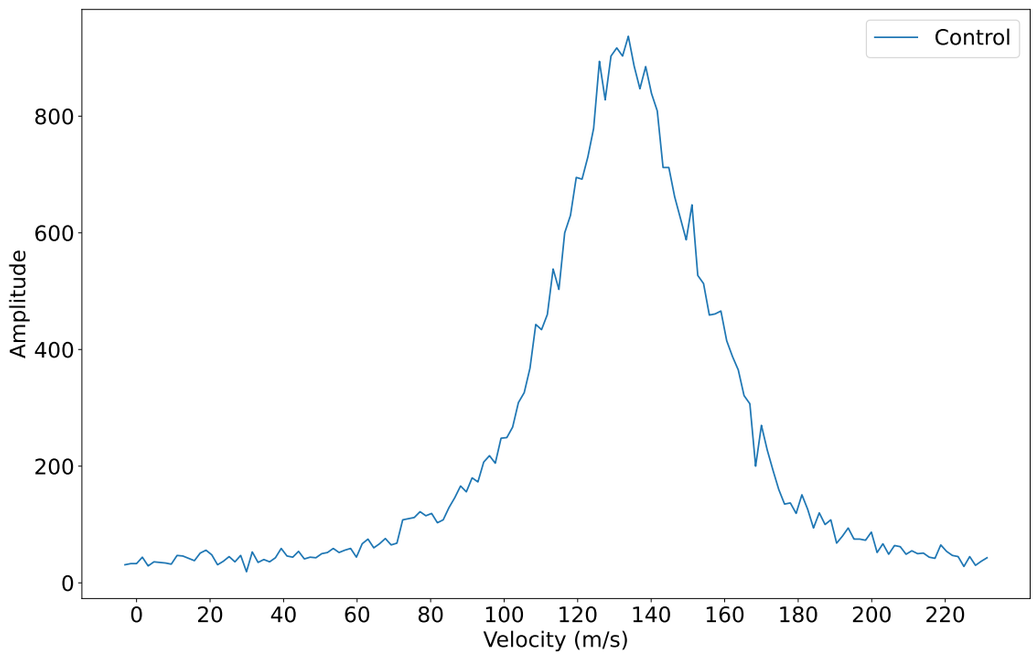

Figure 6.12: Velocity distribution of the new source with the Sr + SF_6 target at optimal operation. It is peaked around 140 m/s, but is much narrower and does not have the long tail extending to higher velocities, indicating complete thermalization inside the cell.

6.5 Using the new source to make a MOT

All of the above data was taken while trying to optimize the absorption signal seen outside the cell. However, the ultimate test will be when we try to slow these molecules and make a MOT. Using the same detection scheme as before, but with the slowing laser now turned on, we were able to characterize that the slowing was working as before, and perhaps we could make a MOT with this target. When we finally got the MOT working again, we were able to perform detailed scan of the ablation parameters.

For optimal MOT numbers, we see that the YAG power wants to be similar, but the Helium flow rate wants to be much lower, to around 3.5 sccm. This is very different from before, and is probably due to the diffuser spreading the Helium more evenly throughout the cell. We also see that at best operation, the absorption outside the cell is limited to around 4-5%. Trying to increase the absorption by increasing the YAG power, or by increasing the Helium flow rate only leads to worse MOTs. My guess is that this is due to the molecules having larger transverse velocities at these higher YAG powers/flow rates, which, in combination with the long slowing length, lead to more of them being lost than captured in the MOT. Perhaps, by shortening the slowing path we will be able to capture more of these molecules. Ultimately, it was found that the MOT numbers were about the same as with a ceramic target, maybe around 20-30% more.

The biggest win was in the lifetime of these targets. A typical target now lasts around 2

months of heavy use, as compared to the 2 weeks from before. After this point, MOT numbers start degrading, even though the absorption signal is unaffected. This is probably due to a lot of dust accumulating in the cell, which I think causes the SF6 to not freeze on the walls, but instead is free, which leads to a hotter beam. The targets themselves are barely dented during this. However, while replacing a target, we need to expose the chamber to air, which causes the Sr to react quickly, so a new target is installed each time. On cleaning the cell and installing a new target, we are again able to get back to usual MOT numbers within a day of running.

6.6 Some mistakes along the way

An initial attempt at isolating the SF6 tube from the 30 K shield was made by drilling a hole in the shield and adding a PEEK spacer through which the SF6 tube would pass. The argument was that the thermal conductivity of PEEK is much lower than copper, and so it should be able to keep the tube at the high temperature. This proved to be a mistake however, because there was significant thermal contact between the tube and the shield in spite of the PEEK spacer only touching the shield in a couple of places. We then realized that the best way to get thermal isolation is by having nothing between the tube and the shield, and so a larger opening was milled in the shield, with the tube just passing through it. No contact is better than any contact (see Figure 6.13). The first iteration also involved the use of a 1/8" OD, OFHC copper tube, while has much higher thermal conductivity. The larger diameter was found to heat the cell to above 4 K, making it unusable.

Another mistake was making the target piece protrude too far into the cell. The first attempt was to get the target at the same height as in the old design, however that proved to be wrong again as explained in §6.3.4. With this bad design, we were unable to see any absorption outside the cell, nor were we able to produce a MOT.

The final mistake was that the SF6 tube was protruding too far into the cell in the first attempt. This was seen to both increase the cell temperature, and we were unable to make any MOT, probably due to the molecules being too hot.

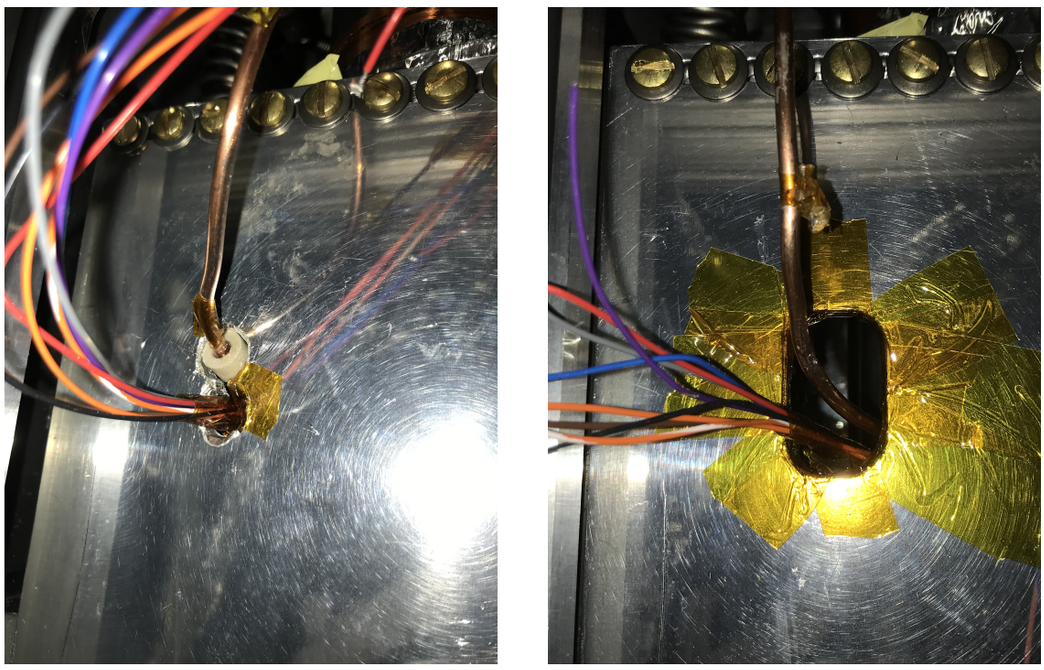

(a) Initial attempt at thermal isolation

(b) Final design

Figure 6.13: SF6 tube isolation attempts. (a) Initial attempt where a PEEK spacer was used between the SF6 tube and the 30 K shield. Even the small contact between the spacer and the shield caused the tube temperature to drop below the freezing point of SF6. (b) Final design with no spacers. The tube goes through vacuum, and is able to maintain the necessary temperature to flow SF6.

6.7 Conclusions

In this chapter, we have demonstrated a new CBGB source involving the chemical reaction of Sr and SF6 gas to produce SrF molecules. We demonstrated techinques to introduce the hot SF6 gas in a cold 4 K environment with negligible heating of the cell, and we also see behavior that is consistent with some of the conventional wisdom regarding CBGB sources: namely a longer cell length leads to slower and more thermal molecules, diffuser yields better extraction, and that there is no reason to have an angled target piece. We have shown that we can increase the yield out of the cell by a factor of 10 as compared to ablation of a ceramic SrF_2 target; however this did not result in an increase in the MOT number. We did see nearly a factor of 10 increase in target lifetime, resulting in the experiment having significantly more up-time. This has enabled the quick realization of the results in the following chapters.

Chapter 7

Red and Blue DC MOTs

7.1 Introduction

The history of molecular magneto-optical traps (MOTs) is quite interesting. The first MOT was demonstrated by our group in 2014 [58], wherein a static magnetic field gradient was used, along with red detuned lasers to provide the trapping. However, this MOT was weakly confining and had a short lifetime, leading to low densities. Mike Tarbutt then realized that this was because the first paper used the wrong laser polarizations, and he proposed an improved trapping scheme [112] with the correct polarizations. In the original DC MOT, it was assumed that the choice of polarization only depended on the sign of the g-facor of the ground state hyperfine levels - however, ignoring dual-frequency effects, the choice of polarizations for any MOT is actually guided by both the sign of the excited state g-facor *and* also whether the transition is $F > F'$ (e.g. $2 \rightarrow 1$) or $F' >= F$ (all other cases). Shortly thereafter, this was implemented in SrF [166], leading to the so-called DC* MOT. This was a better MOT in many aspects, but still had a relatively short lifetime, high temperatures ($T \gtrsim 10$ mK) and weak confining forces as compared to atomic MOTs. This was attributed to the presence of dark states in the ground state manifold.

During this period, Jun Ye's group showed that you can modulate the field gradient and laser polarizations at a radio-frequency (roughly equal to the typical scattering rate of a MOT $\lesssim \Gamma/4$) to make 1D and 2D MOTs of YO molecules [63]. Mike Tarbutt, in his paper, also showed that a 3D RF MOT should be better than the DC* MOT by virtue of getting rid of dark states. Our group then demonstrated the first 3D RF MOT of SrF molecules [59], with longer lifetime and stronger confining forces compared to the DC* MOT. During this time, Mike Tarbutt also realized that you can realize confining forces in a DC MOT that are comparable to those of an RF MOT by addressing one of the transitions with two lasers of opposite detuning and opposite circular polarizations [66]; his group used this technique to realize the first MOT of CaF molecules [56]. Soon thereafter, the Doyle group demonstrated a RF and DC MOT for CaF [62], providing a direct comparison between the two methods, and found that the RF MOT yields marginally higher densities and lower temperatures. Soon thereafter, the RF MOT was demonstrated for YO [64], and more recently, a DC MOT [57] where it was found that the DC MOT is much better for YO than the RF MOT due to the peculiar level structure of YO. Most recently, a polyatomic molecule CaOH was also trapped in an RF MOT [69], showing that both these techniques have now become standard workhorses for molecule trapping experiments.

However, both RF and dual-frequency DC molecular MOTs were found to be quite large ($\sigma \sim 1$

mm) and quite hot ($T \sim 1$ mK, much higher than the Doppler limit $\sim 150\,\mu$K) as compared to typical atomic MOTs. This is due to the sub-Doppler heating mechanism for red-detuning as discussed in Chapter 4. This was the primary motivation for finding ways to create molecular MOTs using primarily blue-detuned light which would provide both sub-Doppler cooling and spatial confinement. This so called blue MOT was first proposed in [65] and then demonstrated to work for Rb atoms in [167]. Very recently, it was also demonstrated to work for YO molecules [168] where they demonstrated cloud sizes $\sim 200\,\mu$m, a roughly 10^2 increase in density. This would certainly bode quite well for loading an ODT due to better mode matching with the small focus size.

In this chapter, we outline our red and blue MOT attempts. First, we discuss our approach to implementing a dual frequency DC MOT for SrF. Next, we show that a simple change in the laser frequencies allows us to directly switch to the blue MOT. We also discuss some peculiarities of the blue MOT and find an optimal trapping configuration that achieves $\sim 10^2$ gain in density. Finally, we show that we can also implement Λ-cooling after the blue MOT, heralding high density ODT loading.

7.2 Red DC MOT

In this section, we detail the principle, implementation and results of the Red DC MOT for SrF.

7.2.1 Red DC MOT Principle

The original scheme for CaF [66] uses the so-called dual-frequency mechanism on the $|F = 2\rangle$ state. The principle behind the dual frequency mechanism is summarized in Figure 7.1. Light of opposite detuning and circular polarization address the same electronic transition. For the example of $|F = 1\rangle \rightarrow |F' = 1\rangle$ transition in Figure 7.1, where the molecule is at positive x, both the $|m_F = \pm 1\rangle$ absorb preferentially from lasers opposing the displacement from the trap center (for $|m_F = 0\rangle$ the molecule is likely to absorb from all lasers), leading to a net restoring force on the molecule. This also eliminates dark state problems with the older DC MOTs, where the molecules in the dark states ($|m_F = -1\rangle$ for the example of Figure 7.1) would not scatter photons.

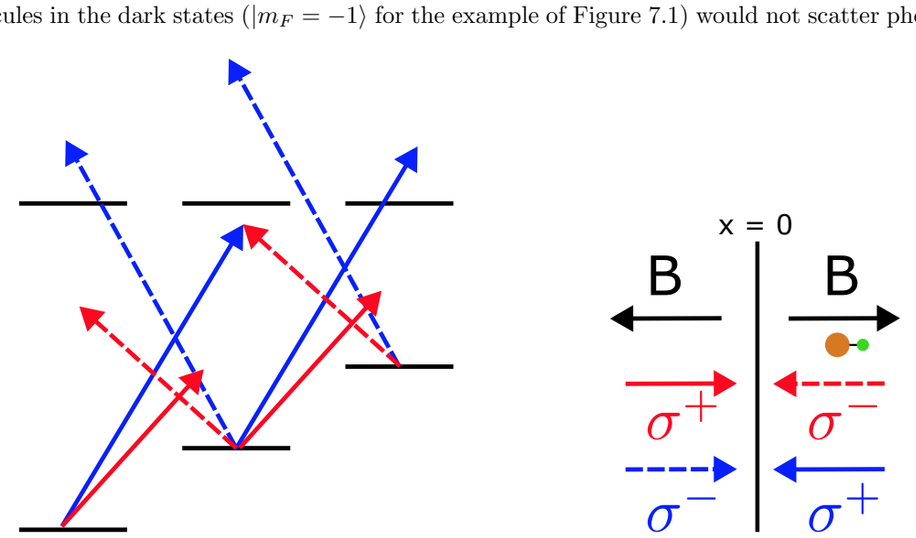

Figure 7.1: Dual frequency red DC MOT principle. Light of opposite detuning and circular polarization address the same electronic transition. For a molecule at positive x, the red detuned σ^- laser addresses the $|m_F = +1\rangle$ level, while at the same time, the blue detuned σ^+ laser addresses the $|m_F = -1\rangle$ level, which in previous DC MOTs, was dark. This provides an overall restoring force, leading to trapping.

In the case of the dual-frequency CaF MOT (see Figure 7.2), all the four hyperfine levels are addressed by red-detuned light (detuned by $\sim -\Gamma$), with the polarization for $|F = 1\uparrow, 0, 1\downarrow\rangle$ opposite that used for $|F = 2\rangle$. Because of the small hyperfine splitting between the $|F = 1\uparrow\rangle$ and $|F = 2\rangle$ states (only around 3Γ), the red-detuned light for $|F = 1\uparrow\rangle$ serves to act as the blue-detuned, oppositely polarized light for the $|F = 2\rangle$ state ($\sim +2\Gamma$).

This scheme can be adapted for SrF with a few modifications. Because of the larger hyperfine splitting in the $N = 1$ manifold ($|F = 1\uparrow\rangle$ and $|F = 2\rangle$ are separated by $\sim 6.3\Gamma$), the dual-frequency mechanism that you get for free in CaF is no longer easily obtainable in SrF. Thus, the dual-frequency mechanism needs to be generated with additional lasers. Furthermore, because of the larger hyperfine splitting, it is experimentally hard to generate the dual-frequency mechanism on the $|F = 2\rangle$ state with the current laser setup. Fortunately, optical bloch equation (OBE) simulations performed by our postdoc T. Langin indicate that the dual-frequency mechanism works just as well on the $|F = 1\uparrow\rangle$, with a similar restoring force as on the $|F = 2\rangle$ state. This is easier to generate, with the resulting level diagram shown in Figure 7.2.

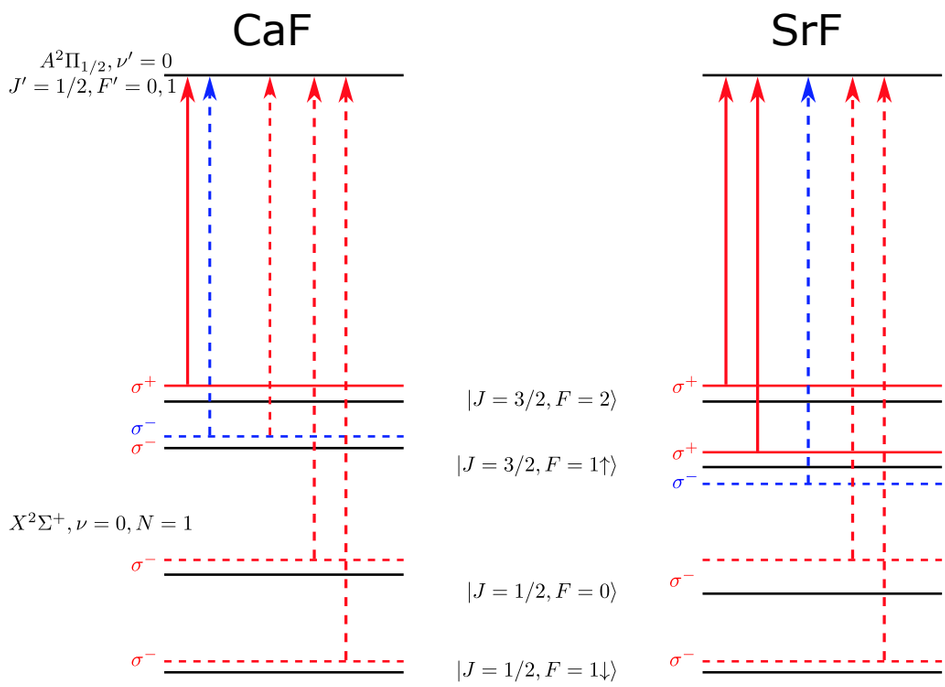

Figure 7.2: Red DC MOT level diagram for CaF and SrF. Solid (dashed) lines indicate $\sigma^+(\sigma^-)$ laser polarization, and color indicates red/blue detuning. For CaF, the dual frequency mechanism is shown on the $|F = 2\rangle$ state, where the red-detuned component for the $|F = 1\uparrow\rangle$ also provides the blue-detuning for the $|F = 2\rangle$ state. For SrF, the dual frequency mechanism is shown on the $|F = 1\uparrow\rangle$ state, where both red and blue detuned light of opposite polarizations act on it.

7.2.2 Implementing the Red DC MOT

The main challenge in implementing this is to get the correct polarizations and detunings for the hyperfine levels. The laser system is shown in Figure 7.3. A Precilaser sum frequency generation module generates up to 2 W of 663 nm light. This light is then split into two arms, V_{00}^M and V_{00}^d, both of which are fiber coupled to two double pass AOM setups, which are capable of around 80 MHz overall frequency adjustment. The V_{00}^d arm has an additional AOM which blue shifts the frequency by 138 MHz (171 MHz) for the DC MOT (RF MOT), such that it is closer to the $|F = 0\rangle$ ($|F = 1\downarrow\rangle$) level when V_{00}^M is tuned to primarily address $|F = 2\rangle$. Each arm also has a fiber EOM, which can generate sidebands up to around 500 MHz. Each arm is then amplified by

successive injection locked slave lasers and tapered amplifiers (TAs), each capable of producing 300 mW of power. An additional half wave plate is used after V_{00}^{d} TA to rotate the polarization 90° with respect to V_{00}^{M}, thus generating light of the opposite polarization, as needed for the DC MOT. These two TAs are then combined on a PBS cube, sent through a single pass AOM which increases the frequency of both by 70 MHz (this AOM is used to control the power of the MOT lasers), coupled to a fiber with the rest of the repump light and sent to the experiment.

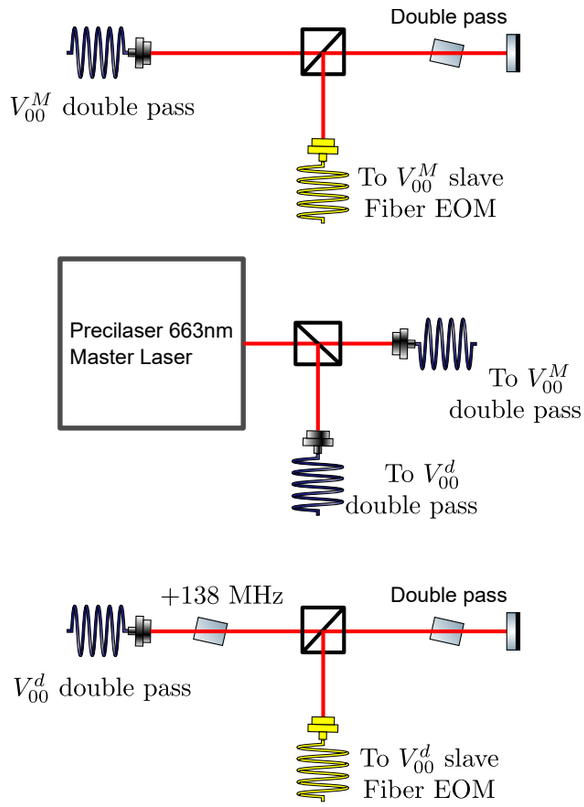

Figure 7.3: DC MOT laser configuration. The light from the precilaser master laser is split in two arms. The V_{00}^{M} arm has a double pass AOM and a fiber EOM to produce the correct frequencies for the MOT. The V_{00}^{d} has a single pass AOM which shifts it relative to V_{00}^{M}, along with a double pass AOM and a fiber EOM to produce the correct frequencies for the MOT. The outputs from both fiber EOMs are sent to respective slave lasers, which then seed tapered amplifiers, producing the final MOT light.

To realize the dual frequency red MOT, the precilaser frequency is set such that after all the frequency manipulations, the V_{00}^{M} carrier is red detuned with respect to the $|F = 2\rangle$ level. The fiber EOM then generates the sideband that provides the red detuning for $|F = 1\uparrow\rangle$, as shown in Figure 7.2. There is an additional sideband that is red of all the levels, but it is far off resonance and does not play any role. Due to the single pass AOM before V_{00}^{d}, the carrier for this laser is closest to the $|F = 0\rangle$ level, but is blue of this state. The double pass AOM is then used to bring this back to the red of $|F = 0\rangle$. The fiber EOM then simultaneously generates the red detuning for the $|F = 1\downarrow\rangle$, and also the blue detuning for the $|F = 1\uparrow\rangle$. Crucially, this blue detuned light is of the opposite polarization as the red detuned light, providing the necessary dual-frequency mechanism.

The magnetic field gradient is generated by modifying the in-vacuum RF MOT coils for DC operation. The impedance matching circuit and the RF amplifier were disconnected, and a DC power supply along with a MOSFET for current control were connected instead. There are additional shim coils on all 3 axes to center the MOT.

7.2.3 Red MOT characterization

After ablation, slowing is applied for 14.5 ms, during which time the MOT beams and gradient are kept on. The initial magnetic field gradient is set to 16 G/cm and the initial per beam laser intensity is set to $I \sim 100$ mW/cm^2 during the capture phase. The MOT is loaded for a total of 37 ms after the slowing pulse is applied, with molecules slowly accumulating in the MOT. These values were found to maximize the number of molecules captured in the MOT. This capture MOT is pretty hot and large ($T \sim 6$ mK and $\sigma \sim 3$ mm) due to the large scattering rates and sub-Doppler heating forces resulting from the high intensity needed to maximize the capture velocity.

Next, we apply a compression stage, similar to what was done for the RF MOT [59], by ramping up the gradient while reducing the laser intensity to reduce the scattering rate. The gradient is ramped up from 16 G/cm to 29 G/cm, while the laser intensity is reduced from $I \sim 100$ mW/cm^2 to $I \sim 10$ mW/cm^2, over 40 ms. We have found that this ramp can be done as fast as 20 ms, and as slow as 60 ms, without affecting the size or number, and so 40 ms is chosen as a good medium. Lowering the final power any further results in the loss of molecules from the trap. After compression, the molecules are held in this compressed MOT for 20 ms in order for molecules to equilibrate. At the end of this sequence, the molecule cloud is compressed to a final size of $\sigma \sim 1$ mm and a final temperature $T \sim 1$ mK with negligible loss of molecules. The final gradient can be further increased to compress the cloud even more, by flowing more current through the coils, but we choose to avoid doing this since the in-vacuum coils heat up substantially in this process (heat goes as I^2). The primary concern here is the Indium solder used on the coils, which melts at only 140°C.

The trapping can be optimized for number or temperature; we choose to optimize for number since the next steps in the experiment involve sub-Doppler cooling. To optimize trapping, we can scan a few different laser parameters. First is the overall detuning of the precilaser master laser, which sets the detuning of the laser addressing $|F = 2\rangle$. This typically wants to be around -0.8Γ from $|F = 2\rangle$. Next, we can scan the red detuning for the $|F = 1\uparrow\rangle$ by scanning the sideband frequency on V_{00}^M. In the old RF MOT results, it was thought that the red detunings for all levels would want to be the same, and so a single resonant EOM was used to generate the sidebands. However, now that we are using a fiber-EOM, it is easy to arbitrarily scan the frequency and amplitude of the sidebands. Indeed, upon scanning this frequency, we found that the optimal trapping occurred at a sideband frequency of 34.6 MHz, which corresponds to around -1.8Γ detuning. We also found that the optimum amplitude ratio is 4:1.

We fix the single pass AOM frequency on V_{00}^d to 138 MHz, and then we can scan the double pass AOM frequency to scan the detuning of the V_{00}^d carrier from $|F = 0\rangle$. This detuning wants to be around -2.7Γ. Then, similarly, we can scan the frequency and amplitude of the sidebands on V_{00}^d. This scan actually convolves two different detunings: the blue detuning that provides the dual-frequency mechanism for $|F = 1\uparrow\rangle$, and the red detuning for $|F = 1\downarrow\rangle$, see Figure 7.2. The optimal sideband frequency is found to be around 56.9 MHz which corresponds to $+0.8\Gamma$ detuning from $|F = 1\uparrow\rangle$ and -1.4Γ from $|F = 1\downarrow\rangle$. The optimal sideband ratio is found to be 1:2:1, with roughly equal powers addressing the two $F = 1$ levels, and higher power for $|F = 0\rangle$ because the laser is tuned to be further away from this state.

After everything is optimized, we can typically get around $N \approx 2.5 \times 10^4$ molecules in the red

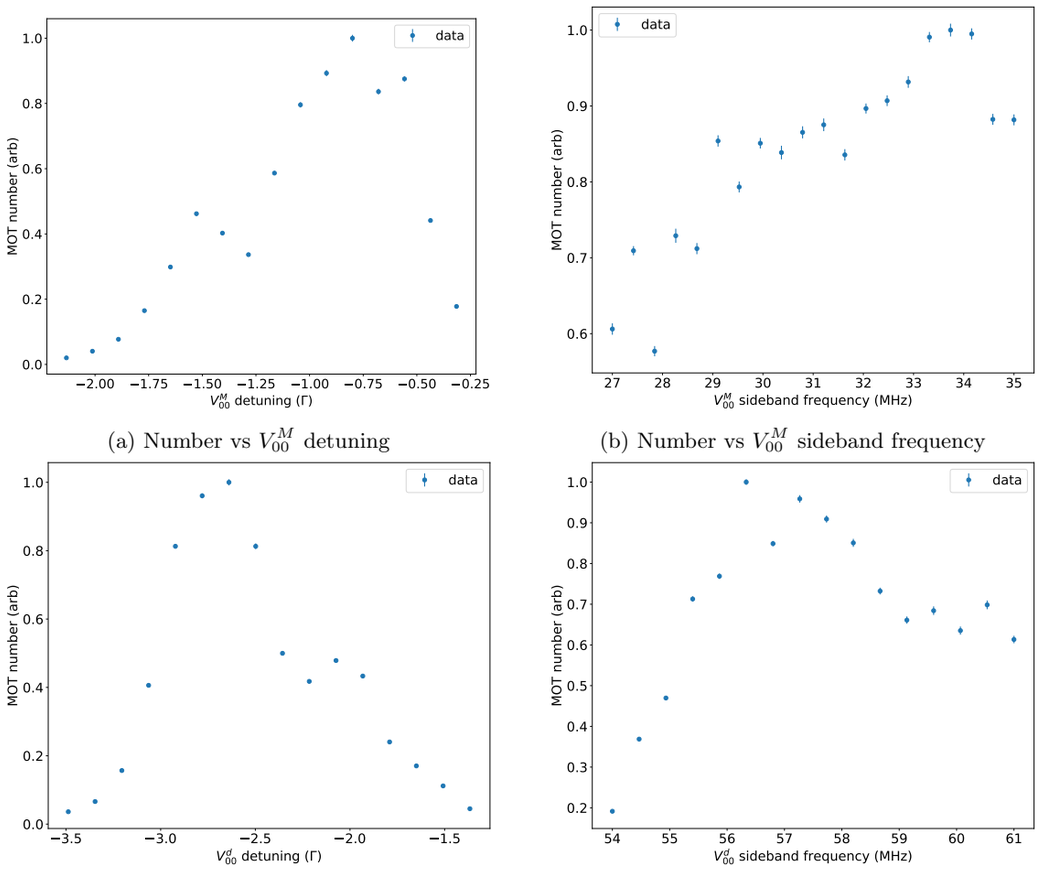

(a) Number vs V_{00}^{M} detuning

(b) Number vs V_{00}^{M} sideband frequency

(c) Number vs V_{00}^{d} detuning

(d) Number vs V_{00}^{d} sideband

Figure 7.4: Red MOT parameter scans.

MOT at a temperature $T \approx 1.5$ mK and a size $\sigma \approx 1$ mm as shown in Figure 7.5. This is around 30% worse than the RF MOT, which is also consistent with what has been observed in CaF [62]. The final laser parameters are shown in Table 7.1.

Laser	Red MOT
V_{00}^{M} carrier	-0.8Γ of $\|F=2\rangle$
V_{00}^{M} sideband	34.6 MHz (-1.8Γ of $\|F=1\uparrow\rangle$)
V_{00}^{d} carrier	-2.7Γ of $\|F=0\rangle$
V_{00}^{d} sideband	56.9 MHz ($-1.4\Gamma, +0.8\Gamma$ of $\|F=1\downarrow, 1\uparrow\rangle$)

Table 7.1: Frequencies/detunings for the different lasers used for the red MOT.

7.2.4 Molecule number measurement

To determine the MOT number, we typically take a 2 ms fluorescence image taken at the maximum laser intensity of $I \sim 170$ mW/cm^2. To convert the camera counts to a molecule number, we need to know the scattering rate of the molecules while the image is being taken. The image is taken at the same detuning as the red MOT for ease of operation.

One way to determine the scattering rate is as outlined in [76]. We can measure the MOT lifetime in the presence of the light for two cases: one with the repump V_{32}^{M} present, and one with the repump V_{32}^{M} shuttered. Here V_{32}^{M} is the laser addressing the $v = 3 \rightarrow v' = 2$ transition to repump molecules from the $v = 3$ vibrational state. If the lifetime in the absence of light (i.e. the vacuum lifetime and any other loss mechanisms) is longer than any of these individual lifetimes, then the additional decay in the absence of V_{32}^{M} is solely determined by the scattering rate and

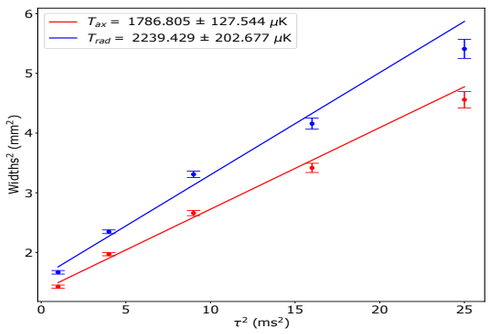

(a) Temperature measurement of the red MOT

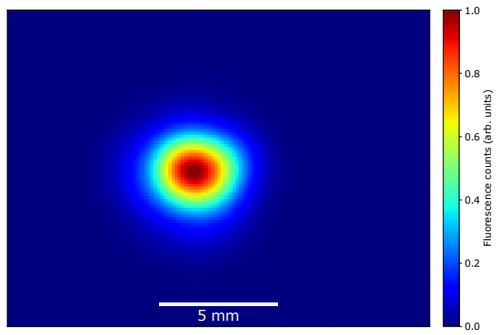

(b) False color image of the compressed red MOT

Figure 7.5: Temperature measurement and red MOT image. (a) Temperature measurement of the red MOT using the time of flight method (TOF). Lines show fits to the cloud expansion with extracted temperatures. (b) False color image of the compressed red MOT.

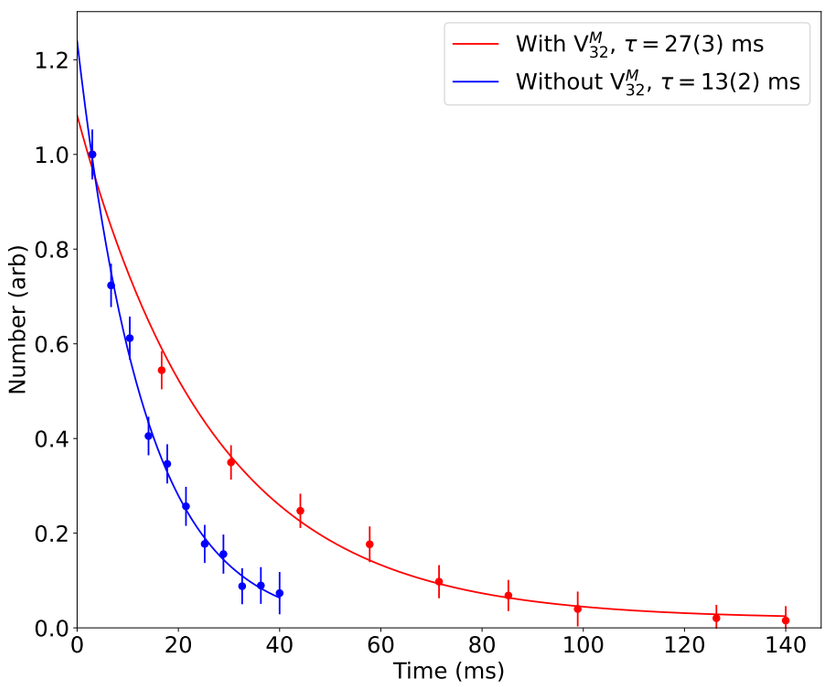

Figure 7.6: Example of a scattering rate measurement. MOT LIF signal as a function of time is measured with and without the V_{32}^M laser. Then the additional loss without this laser only occurs due to decay to the $v = 3$ vibrational level, and by knowing the branching ratio to that state (b_{03} per [76]), the scattering rate is determined.

branching ratio (b_{03}) to that state from the $A\Pi$ state. Thus by knowing this branching ratio, and by measuring the loss into $v = 3$, we can determine the scattering rate.

The scattering rate is strongly dependent on the laser power, and so this measurement must be performed at the same laser power as that used for imaging. Since we are using the maximum laser power, the MOT lifetime is quite low. Figure 7.6 shows an example of a scattering rate measurement, where the lifetime is measured with and without the repump laser. We repeat this measurement a few times to get a better estimate of the lifetime for each case, and we find that $\tau_{\text{with}} = 27(1)$ ms, while $\tau_{\text{without}} = 13.5(1.0)$ ms. Then, the additional decay into $v = 3$ is given by $\gamma_{v=3} = R_{sc} b_{03} = \frac{1}{13.5(1.0)} - \frac{1}{27(1)} = 37(6)$ s^{-1}, where the uncertainties are added in quadrature since they are uncorrelated.

Given a value of b_{03}, we can then determine the scattering rate. However, there is significant uncertainty in the literature about the value of b_{03}. There are two calculations that provide differing values: the one from [169] states $b_{03,Barry} = 9.6 \times 10^{-6}$ while the one from [79] provides $b_{03,NL-eEDM} = 1.1 \times 10^{-5}$. These give two different values of the scattering rate:

- $R_{sc,Barry} = 3.85(0.62) \times 10^6$ s^{-1}
- $R_{sc,NL-eEDM} = 3.36(0.54) \times 10^6$ s^{-1}

We choose to use an average scattering rate, and add another uncertainty corresponding to the difference from the mean. This gives us a final scattering rate of $R_{sc} = 3.6(6) \times 10^6$ s^{-1}, i.e. around a 17% uncertainty in each number measurement.

To convert from camera counts to molecule number, we use the relation [76]

$$N_{obs} = \frac{N_c}{G \cdot \eta_{qe} \cdot \eta_{geo} \cdot \eta_{fil} \cdot \eta_{tra} \cdot N_{per}} \tag{7.1}$$

where N_c are the integrated camera counts, G is the camera gain, η_{qe} is the quantum efficiency of the camera, η_{geo} is the geometrical collection efficiency, η_{fil} is the interference filter efficiency and η_{tra} is the transmission efficiency of the collection optics. All these have been characterized before [76], and we add an additional 10% error for the calibration of these efficiencies. $N_{per} = R_{sc} \cdot t_{exp}$ is the number of photons scattered per molecule. For a 2 ms exposure, we find that

$$N_{obs} = N_c \times 1.13 \times 10^{-2} \tag{7.2}$$

A typical fully optimized MOT after a fresh target change will have around 2 million camera counts in the image, giving a number of around $N \approx 2.5 \times 10^4$. This usually persists for a couple of weeks, after which the number starts slowly dropping. On a typical day, the MOT count can fluctuate from 1.2 - 1.6 million.

7.3 Blue DC MOT

Even though we take a hit of 30% in number by going from RF to DC MOTs, the reason we switched to the red DC MOT was to easily implement the blue DC MOT. This reduces the temperature and size of the trapped cloud relative to the red MOT, as outlined in this section.

7.3.1 Blue DC MOT principle

A typical atomic MOT is fairly dense and cold, since the same lasers that provide trapping are also providing strong cooling for free via the well-known polarization gradient cooling (also referred to as Sisyphus cooling) mechanism [28, 30, 65, 94, 110, 170]. This mechanism exists when transitions between a ground state with $F_g > 0$ and an excited state with $F_e > F_g$ (often referred to as type-I transitions, see Chapter 4) are driven. In contrast to this, a typical molecular system has a type-II transition ($F_g \geq F_e$) with he consequence that red-detuned light now provides Doppler cooling, but leads to sub-Doppler *heating* [65, 110]. This helps explain why most molecular MOTs are much worse than their atomic counterparts. The sub-Doppler cooling force you get for free in atoms is no longer present in molecules; instead you have a competing heating force that limits what temperature the MOT can achieve.

This behavior is clearly seen in the simulation curves shown in Figure 7.7. These curves are obtained by solving the optical bloch equations with the appropriate three-dimensional laser and magnetic fields incorporated [110], and were produced by T. Langin. From the curves, one can see that above a velocity of around 2 m/s and above a displacement of around 1.5 mm, there is

a damping and a restoring force acting on the molecules, corresponding to the typical Doppler cooling for red detuning. However, below these values, the sign of the force reverses, i.e. we now have sub-Doppler heating, which places a limit on how cold and well-confined red-detuning can get us. This is the main reason why molecular ODTs contain so few molecules: the small spot size of the ODT beam has a very small overlap volume with the large size of a typical MOT, leading to low transfer efficiencies ($\lesssim 5\%$) [53, 68, 111, 124, 171].

By the same logic though, if you flip the detuning to blue, you would now get Doppler heating, but sub-Doppler *cooling*. This mechanism is seen in the gray-molasses cooling technique widely used in atoms [98, 172], and also the Λ-enhanced gray molasses used in molecules [53–55, 57, 111], which has enabled molecular clouds to be cooled to as low as $5\,\mu$K, see Chapter 4. However, these mechanisms usually require zero magnetic fields (since they rely on stable Zeeman dark states), and thus provide only a damping force but no confining force. The molecular cloud size thus remains the same during Λ-cooling.

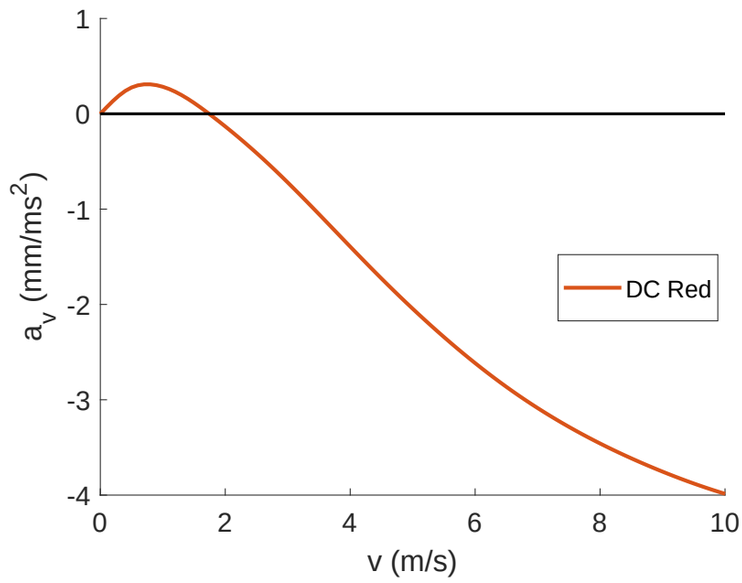

(a) Acceleration vs velocity

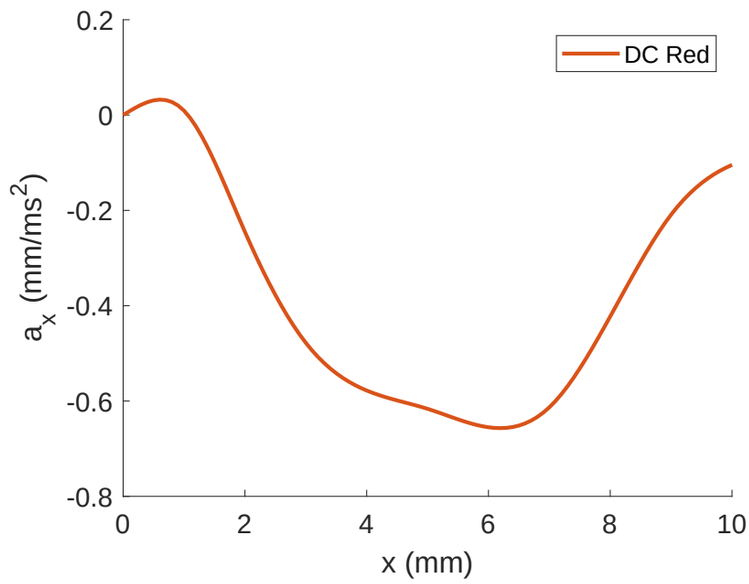

(b) Acceleration vs position

Figure 7.7: Simulated acceleration curves for the red DC MOT averaged over position and velocity ($a(v) = \int_{-5\text{mm}}^{5\text{mm}} a(x,v)dx$ and $a(x) = \int_{-3\text{m/s}}^{3\text{m/s}} a(x,v)dv$). (a) Acceleration vs velocity. (b) Acceleration vs position. These help explain why the red DC MOT is hot and large. For velocities above around 2 m/s and for positions above around 1.5 mm, Doppler cooling is observed, where there is a damping and a restoring force on the molecules. However, below these values, the sign of the force is reversed, and there is sub-Doppler heating instead, imposing a limit on the size and temperature. Curves obtained from T. Langin.

This naturally led to interest in 'blue-detuned' MOTs (blue MOTs), which can exhibit sub-Doppler cooling while simultaneously maintaining strong confining forces, for Type-II transitions. The possibility of such a blue MOT was first suggested in [65], and then demonstrated in Rb atoms [167]. The key realization is that, for type-II systems, the position dependent force remains the same if the detuning and polarization are reversed, but the velocity dependent force reverses sign if the detuning is flipped, regardless of polarization. This is crucial because the main problem in type-II MOTs is the sub-Doppler heating from the velocity dependent force. The natural way to try to achieve a blue MOT is to first capture in a red MOT, where you get Doppler cooling, and then switch to a blue MOT, for sub-Doppler cooling. In the Rb paper [167], they showed that such a blue MOT can give higher density and phase space density than in a typical type I red MOT.

Very recently, a blue MOT was also shown to work for YO molecules [168], the first implementation of a blue MOT for molecules. They observed around 50% molecule transfer from a red MOT to a blue MOT, with a final MOT size $\sigma \approx 220\,\mu$m and $T \approx 35\,\mu$K reached after 170 ms in the blue MOT. YO is a special case, because it has a magnetically insensitive ground state with $F \neq 0$ (the $|G = 0, F = 1\rangle$ state). This feature has been observed to increase the robustness of sub-Doppler cooling in magnetic fields up to 4 G [57], and is unlike molecules such as SrF or CaF where, due to non-zero ground state magnetic moments, magnetic fields destroy the delicate dark states needed for sub-Doppler cooling. Because of this, the blue MOT works for YO with just 2 lasers, creating both a Λ-enhanced gray molasses, and a MOT. Indeed, they see that, by tuning the two-photon detuning, they can easily switch between which of the two mechanisms dominates: at close to zero two-photon detuning, they can get colder temperatures but larger MOTs, because of the low excited state population during Λ-enhanced gray molasses, and for higher detunings they can get tighter confinement but larger temperatures, corresponding to the blue MOT mechanism.

For SrF, the story is quite different. There are no magnetically insensitive $F \neq 0$ states and, as such, Λ-enhanced gray molasses works best in zero field, ruling out this cooling mechanism. How can we implement a blue MOT in this case? One possibility is to, again, resort to the dual-frequency mechanism. It had previously been observed that, in static magnetic fields, simple gray molasses cooling works up to a few G [55, 56], and can cool the molecules down to a minimum of $\sim 50\,\mu$K, with the temperature rising to $\sim 200\,\mu$K at a few G. Thus, it seemed reasonable that a sub-Doppler MOT could work, given our red MOT cloud size of ~ 1 mm and our MOT gradient of 3 G/mm. To obtain a trapping force, we decided to try introducing a red-detuned component to implement a dual frequency mechanism, similar to the red MOT [168].

Simulations indicated that our approach should work, as can be observed in the OBE solver curves shown in Figure 7.8 for a dual-frequency mechanism on $|F = 1\uparrow\rangle$ and blue detuned light addressing the other states (see Figure 7.9). As shown previously, the red MOT exhibits sub-Doppler heating below 2 m/s and 1.5 mm. However, by flipping the detuning to blue, we can generate the sub-Doppler cooling force as mentioned before. The critical thing to note here is that, for the laser parameters chosen in each case (Figure 7.2 for the red MOT and Figure 7.9 for the blue MOT) the sub-Doppler cooling force is effective to higher velocities and positions than the sub-Doppler heating force from red detuning. This is important because it suggests that all the molecules trapped in the red MOT should be within the capture parameters of the blue MOT. Another important feature is the large acceleration for blue detuning - this suggests strong confining forces, which should lead to a tight and dense blue MOT.

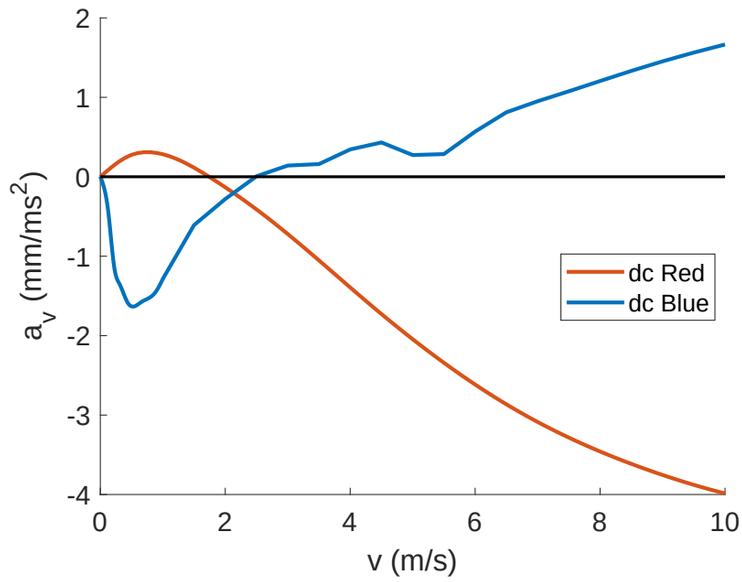

(a) Acceleration vs velocity

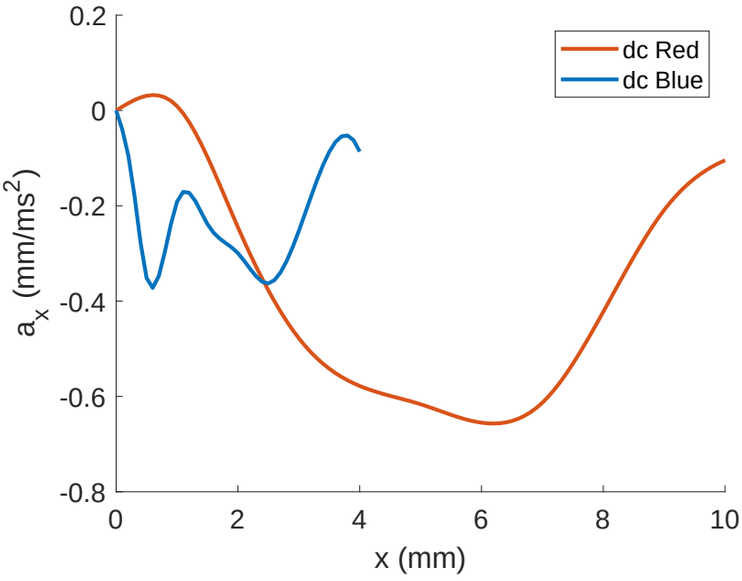

(b) Acceeleration vs position

Figure 7.8: Simulated acceleration curves for the blue DC MOT averaged over position and velocity ($a(v) = \int_{-5\text{mm}}^{5\text{mm}} a(x,v)dx$ and $a(x) = \int_{-3\text{m/s}}^{3\text{m/s}} a(x,v)dv$). (a) Acceleration vs velocity. (b) Acceleration vs position. The red MOT has the characteristic sub-Doppler heating features. However, by flipping the detuning to blue, we can generate a sub-Doppler cooling force. Crucially, the sub-Doppler cooling from blue-detuning is effective out to both higher velocity and larger position than where the sub-Doppler heating from red-detuning takes over. Thus, one can capture the molecules in a red MOT first, and then switch the detuning to effectively capture in the blue MOT. Curves obtained from T. Langin.

7.3.2 Implementing the blue DC MOT

With the requirement of a dual-frequency mechanism on $|F = 1\uparrow\rangle$ understood, we now need to be able to generate the required frequency components for both the red and blue MOT using our double pass AOMs and fiber EOMs (Figure 7.3). The red MOT frequencies are as detailed before (Figure 7.2), and the blue MOT frequencies are shown in Figure 7.9. To shift from the red MOT to the blue MOT, the V_{00}^{M} carrier is blue shifted by increasing the frequency of the double pass AOM, putting the carrier between the $|F = 2\rangle$ and $|F = 1\uparrow\rangle$ states. This laser acts as the blue detuned component for $|F = 2\rangle$ giving a sub-Doppler cooling force, while also acts as the red-detuned light for the $|F = 1\uparrow\rangle$ to provide the dual-frequency mechanism. The fiber EOM for this laser then generates the sideband targeting $|F = 0\rangle$. The detuning of this frequency component does not really matter as long as it is close to $|F = 0\rangle$, it mainly serves to repump molecules out of $|F = 0\rangle$. V_{00}^{d} carrier is also blue shifted such that the carrier now lies blue of $|F = 1\downarrow\rangle$, again providing a sub-Doppler cooling force. The fiber EOM for this laser then generates the sideband that provides the blue detuning for the $|F = 1\downarrow\rangle$, as necessary for the dual-frequency mechanism. We also note that the signs of the polarizations addressing $|F = 1\uparrow\rangle$ are the same for the red and blue MOT, so there is no need to reverse the field. By tuning each of these frequencies and the sideband ratios, we can optimize the blue MOT. The blue MOT field is generated with the same set of coils as used for the red MOT.

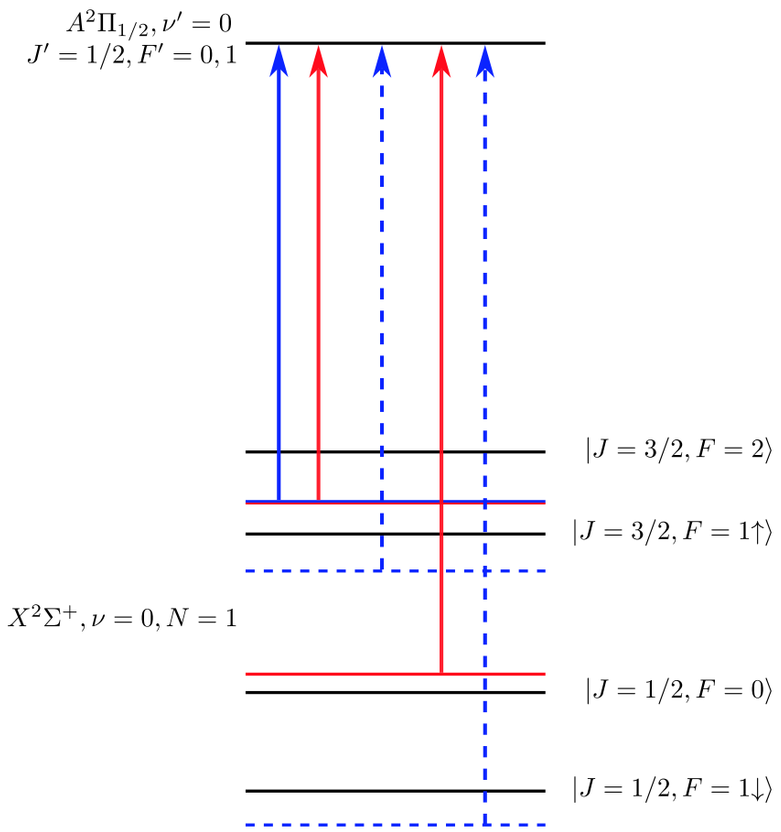

Figure 7.9: Blue DC MOT level diagram. Solid (dashed) lines indicate $\sigma^+(\sigma^-)$ laser polarization, and color indicates red/blue detuning. The dual frequency mechanism is on the $|F = 1\uparrow\rangle$ state, which is addressed by red and blue detuned light of opposite polarizations.

7.3.3 Blue MOT characterization

As stated before, after the red MOT compression, the temperature is $T \approx 1$ mK and the size is $\sigma \approx 1$ mm, both within the expected capture range of the blue MOT. Next, we instantaneously jump to the blue-MOT configuration. The laser frequencies are changed to those in Figure 7.9, and the intensity is increased to $I \sim 170$ mW/cm^2, corresponding to $I/I_{\text{sat}} \sim 60$, where I_{sat} is the saturation intensity. The magnetic field gradient is kept at $b = 29$ G/cm after the red MOT. While we were initially searching for the blue MOT, we used the optimal detunings found from the simulation, with the set of frequencies as listed in Table 7.2. We know that we want very low power in the V_{00}^M sideband addressing $|F = 0\rangle$, and we want roughly equal powers addressing the two $|F = 1\rangle$ states. The RF amplitudes driving the EOMs are set accordingly, with a roughly 1:6:1 and 1:1:1 sideband:carrier:sideband ratios for V_{00}^M and V_{00}^d respectively.

Laser	Red MOT	Blue MOT
V_{00}^M carrier	-0.8Γ of $\|F = 2\rangle$	$+2.8\Gamma, -3.48\Gamma$ of $\|F = 2, 1\uparrow\rangle$
V_{00}^M sideband	34.6 MHz (-1.8Γ of $\|F = 1\uparrow\rangle$)	92 MHz (-1.6Γ of $\|F = 0\rangle$)
V_{00}^d carrier	-2.7Γ of $\|F = 0\rangle$	$+1.72\Gamma$ of $\|F = 1\downarrow\rangle$
V_{00}^d sideband	56.9 MHz ($-1.4\Gamma, +0.8\Gamma$ of $\|F = 1\downarrow, 1\uparrow\rangle$)	132.5 MHz ($+2.1\Gamma$ of $\|F = 1\uparrow\rangle$)

Table 7.2: Inital list of frequencies/detunings for the different lasers for the blue MOT.

This first blue MOT was only a modest improvement over the red MOT, both in size ($\sigma_{ax,rad} \approx 400\,\mu$m, with ax being the axial and rad being the radial direction) and in temperature ($T_{ax,rad} \approx 400\,\mu$K). However, we did find that at least it captured 80% of the molecules from the red MOT. We tried a lot of optimizations here, and during this period we found that the blue MOT is very sensitive to the sideband ratio of the lasers. We saw that the sideband ratio from our fiber EOM was fluctuating by over 50% over several seconds, and, while this did not matter too much for the red MOT, it led to a noticeable effect on the blue MOT size and number. This problem was traced to the RF switches touching the aluminum casing of the RF boxes driving the fiber EOMs; upon putting plastic tabs under the switches, the drift was almost completely eliminated. Even after eliminating this problem, we were initially unable to improve the blue MOT during these early optimization attempts.

We then considered that this might be because of whatever intensity imbalance is present in the MOT beams, causing some weird polarization gradients that were detrimental. We had seen this effect when Λ-cooling in the ODT [111], and figured that this might play a role here too. In fact, at Yale, we had seen that by changing the alignment of the MOT retro mirror, we could "walk" the Λ-cooling beams into an optimal spot for cooling and loading into the ODT - presumably we just changed the polarization gradients at the location of the ODT by changing the alignment of the final mirror. We use a motorized mirror mount (Newport Picomotor 8816-6) to provide fine adjustment of the alignment of the MOT retro mirror. To monitor how much we moved this retro mirror, we can look at the very first iris after the MOT fiber; by closing this iris partially we are able to see the retro beam position, which is normally adjusted to overlap with the first pass. We tried misaligning the last mirror, and saw that by moving the retro 15 ticks to the "right" on the iris, we were able to reduce the size of the blue MOT down to $\sigma \approx 320\,\mu$m. We then saw that making the V_{00}^M carrier bluer (from $+2.8\Gamma$ to $+5.2\Gamma$ of $|F = 2\rangle$; correspondingly, we changed the

sideband such that the detuning from $|F = 0\rangle$ was now -0.5Γ) dramatically reduced the size, to as low as $\sigma_{ax,rad} \approx 147, 149\,\mu$m. We believe this is because this change results in an increase in the blue detuning with respect to $|F = 2\rangle$, potentially increasing the effectiveness of the sub-Doppler cooling force, while simultaneously bringing it closer to, but still red of, resonance with $|F = 1\uparrow\rangle$, which may increase the effective restoring force. This change also resulted in a colder blue MOT, with a final temperature $T_{ax,rad} \approx 200\,\mu$K.

The effectiveness of the blue MOT can be most clearly observed in the fast capture and compression time of the molecules from the red MOT to the blue MOT, shown in Figure 7.10. Within 30 ms of the switch, around 80% of the molecules from the red MOT have been completely compressed and cooled. The final size corresponds to a peak number density $n_0 \approx 4 \times 10^8$ cm^{-3}, a gain of $\sim 10^2$ with respect to the compressed red MOT. The final temperature can be further lowered to $T_{ax,rad} \approx 60\,\mu$K by reducing the laser intensity to $I \approx 34$ mW/cm^2 (20% of its initial value), though this results in an increase in cloud radius to $\sigma_{rad} \approx 230\,\mu$m and $\sigma_{ax} \approx 153\,\mu$m. The blue MOT reaches a maximum phase space density $\Phi_0 \approx 1.6 \times 10^{-9}$, a gain of $\sim 10^4$ with respect to the compressed red MOT. We always run the blue MOT at the highest intensity, since we care more about compressing the cloud than lowering the temperature.

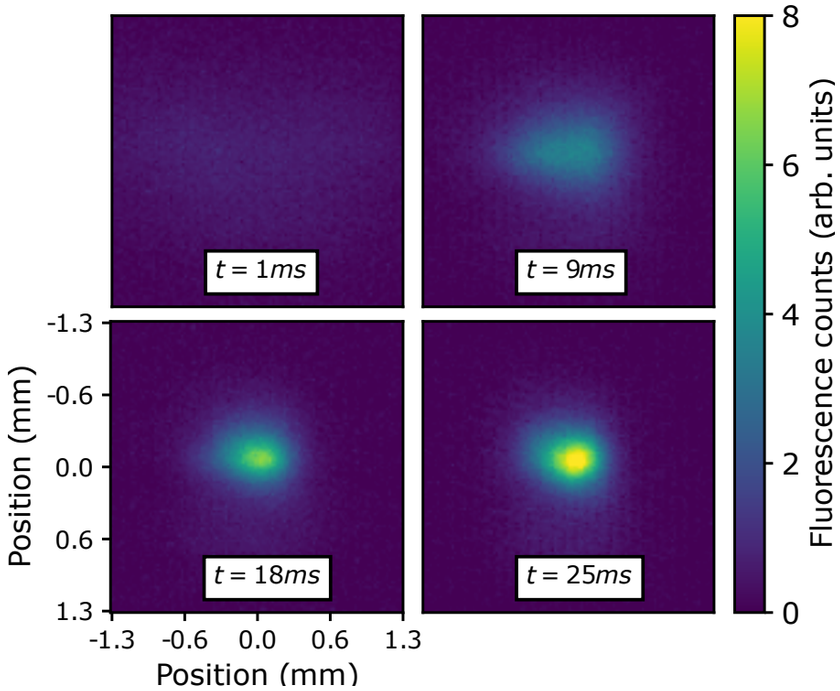

Figure 7.10: Capture in the Blue DC MOT. Fluorescence images showing capture into the blue-MOT as a function of time t after switching from the red-MOT (2 ms in situ exposure starting at time t). The loading is quick and efficient, with $\approx 80\%$ of molecules captured within 30 ms of the start of loading.

The trapping mechanism used here (dual-frequency, with trapping resulting from the magnetic moment of the ground state ($g_{F=1\uparrow} \sim 1$) is substantially different from the mechanism used in the YO blue MOT, where hyperfine levels were only addressed by blue light, and trapping results primarily from the magnetic moment of the excited state ($g_{YO,A\Pi} \sim -0.06$). As a result, the two blue MOTs exhibit starkly different behavior. We observe a stronger restoring force (~ 10 times faster compression) and smaller minimum trap volume (by a factor of 2) at the cost of a higher minimum blue-MOT temperature ($60\,\mu$K vs $38\,\mu$K) compared to YO. We also observe a higher spring constant (from the equipartition theorem, $\kappa = \frac{k_B T}{\sigma^2}$) of $\kappa \sim 10^{-19}$ N/m as opposed to YO with a spring constant $\kappa \sim 5 \times 10^{-21}$ N/m. We attribute the stronger confining forces

demonstrated here to the fact that $g_{F=1\uparrow} \gg g_{YO,A\Pi}$.

Additionally, we point out that the approach used for YO, which relies on the robustness of their gray-molasses cooling resulting from their use of a ground state with near zero magnetic moment, may not be generally applicable, since most other molecules that have been laser cooled (SrF, CaF, CaOH, SrOH, CaH, CaD), or are proposed as good candidates for laser cooling (MgF, BaF, YbF, RaF, CH, etc.) do not have this feature. The dual-frequency mechanism proposed here, however, should be generalizable to these cooling and trapping candidates, and has recently been demonstrated in CaF already [173].

The blue MOT is highly sensitive to the MOT beam alignment. As mentioned before, optimizing the blue MOT required misaligning the retro. The range of misalignment over which a good blue MOT is observed is quite small - only around 7-8 "ticks" of the last mirror. In any of the directions away from this optimum, the blue MOT deteoriates very quickly, with size increasing to over 300 μm and capture efficiency decreasing sharply. This is also similar to the ODT at Yale, which was highly sensitive to the retro position. In a similar vein, the blue MOT is also highly sensitive to the laser polarizations. As it turns out, even though a fiber is labeled polarization maintaining, it can still drift over time. In our case, the polarization out of the MOT fiber drifts even on touching the fiber, and also there is a general drift over time. To compensate for this, we periodically check the polarization coming out of the MOT fiber with a PBS, and minimize the transmission for V_{00}^{d} using a combination of a half and quarter wave plate placed after V_{00}^{M} and V_{00}^{d} are combined. In principle, one should only need a half wave plate, since the polarization out of the TAs are linear, and the two lasers are also combined on a PBS; however, the fiber is clearly introducing some ellipticity, and so a quarter wave plate is also used. Almost every time the blue MOT signal is worse, it can be fixed either by readjusting the retro or by fixing the polarization.

7.3.4 Further cooling after the blue MOT

The final temperature of 200 μK is too hot to try and load into a $\approx$ 1.2 mK deep trap. Thus, we use Λ-cooling (see Chapter 4) after the blue MOT for further cooling and loading into the ODT. Our initial attempts at Λ-cooling all failed to produce temperatures below $T \approx 100\,\mu$K. This was quite in contrast to the results at Yale, where we were easily able to achieve $T \approx 15\,\mu$K routinely. We scanned all the shims and as far as we could tell, we had optimized all the Λ-cooling parameters to no avail. Our initial guess was that there might be eddy currents induced in the MOT coils/surrounding chamber, which would cause the magnetic field to change over time, and worsen cooling. This was also the first time we were trying Λ-cooling at UChicago, and so we decided to switch to the RF MOT, try Λ-cooling since there would be no eddy currents, optimize Λ-cooling parameters and shims, and then redo for the blue MOT. Even after this switch, it was hard to get anywhere below $T \approx 50\,\mu$K for a long time. During this period, we found that there was cross-talk across some of the power supplies used for the shims, where the output of one would change depending on the output of another. We also tried various different MOT beam shapes in the hopes that was the problem, and also tried a few different MOT fibers in case the polarization was drifting too much. Ultimately, to our surprise, we found that simply gluing the dichroics used to combine the ODT and the MOT beams to their mirror mount helped us reduce the temperature to $\sim 30\,\mu$K, which is sufficient for loading into the ODT, albeit worse than our old best results

from Yale.

On switching back to the blue MOT after this, however, everything else had gotten a lot worse. The blue MOT size had risen to $\sigma \approx 350\,\mu$m, and the blue MOT temperature was now $T \approx 300\,\mu$K. We were able to get Λ-cooling to work after the blue MOT and reduce the temperature down to $T \approx 35\,\mu$K consistently but now the size of the blue MOT was much larger than the optimum, which would preclude good transfer to the ODT. We tried to redo the steps we had taken after switching to the RF MOT, but we were not able to reduce the size back to the original values, although we were able to reduce the temperature back down to 200 μK. We then decided to try and scan the blue MOT frequencies again, in case something changed. The results are shown in Figure 7.11 and are quite surprising. The biggest change was the optimal V_{00}^d EOM sideband frequency changed drastically from 132.5 MHz to 122.5 MHz, and V_{00}^d carrier shifted from being $+1.72\Gamma$ to $+0.5\Gamma$ of $|F = 1\downarrow\rangle$, both of which are big shifts. The shift in sideband compensated for the red shift of the carrier such that the blue detuning from $|F = 1\uparrow\rangle$ remained more or less the same. We also found that the optimum retro location had now shifted to the "left" by 30 ticks, instead of "right" by 15 ticks. Ultimately, this new blue MOT was able to reach a final size of $\sigma_{ax,rad} \approx 220, 240\,\mu$m and we were able to apply a Λ-cooling pulse after the blue MOT to reduce the temperature to $T \approx 35\,\mu$K without additional increase in size. The source of this change in size and blue MOT parameters is yet to be tracked. Possibilities include issues with the shim coils, issues with the MOT fiber, or even issues with the polarization out of the fiber.

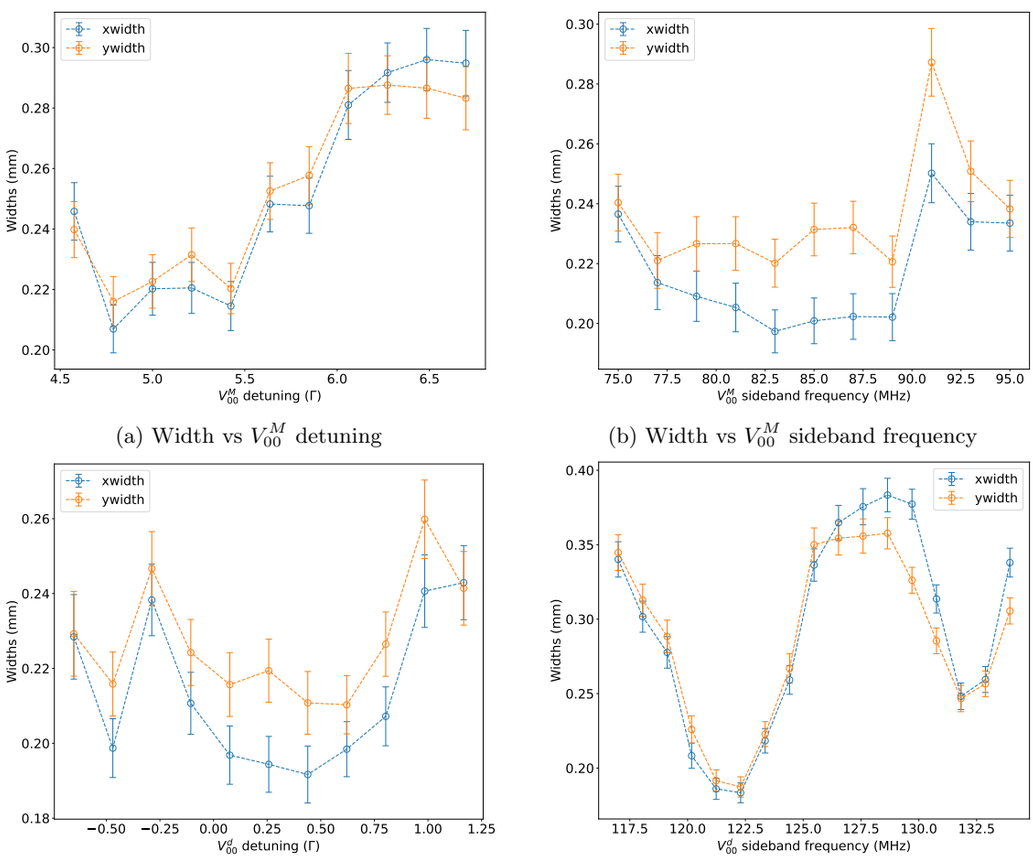

(a) Width vs V_{00}^M detuning

(b) Width vs V_{00}^M sideband frequency

(c) Width vs V_{00}^d detuning

(d) Width vs V_{00}^d sideband

Figure 7.11: Blue MOT parameter scans. There is an obvious double peak feature in (d) where we see that the worst confinement is, roughly speaking, achieved for the lambda condition ($delta_R = 0$). This is probably for dark-state coupling reasons, similar to in the YO blue MOT, where the lambda condition led to cold, but large, clouds.

7.4 Conclusions

In this chapter, we described how we successfully implemented a red, dual-frequency DC MOT of SrF molecules. By switching the detunings of some of the lasers, we were able to demonstrate a blue, dual-frequency DC MOT with a size around $\sigma_{ax,rad} \approx 220, 240\,\mu$m and $T_{ax,rad} \approx 200\,\mu$K. This represents a factor of 10^2 gain in density and 10^4 gain in phase space density as compared to the red MOT. The tight volume of this blue MOT heralds large loading efficiencies in the ODT as will be demonstrated in the next chapter.

Laser	Red MOT	Blue MOT
V_{00}^M carrier	-0.8Γ of $\|F=2\rangle$	$+5.2\Gamma, -1.1\Gamma$ of $\|F=2,1\uparrow\rangle$
V_{00}^M sideband	34.6 MHz (-1.8Γ of $\|F=1\uparrow\rangle$)	85 MHz (-0.5Γ of $\|F=0\rangle$)
V_{00}^d carrier	-2.7Γ of $\|F=0\rangle$	$+0.5\Gamma$ of $\|F=1\downarrow\rangle$
V_{00}^d sideband	56.9 MHz ($-1.4\Gamma, +0.8\Gamma$ of $\|F=1\downarrow, 1\uparrow\rangle$)	122.5 MHz ($+1.8\Gamma$ of $\|F=1\uparrow\rangle$)

Table 7.3: Final list of frequencies/detunings for the different lasers for the blue MOT.

Chapter 8

Observation of collisions in the ODT

8.1 Introduction

The quest for a BEC of polar molecules has been a long winded affair, with the first high phase space density gas of a bialkali molecule produced in 2008 [40]. However, these experiments were unable to cool the molecules to a BEC using conventional techniques. This is because collisional cooling requires a sufficiently high ratio of elastic to inelastic collisions in order to cool efficiently [30, 43]. However, to date, all the experiments with trapped ultracold molecules, be they directly cooled or bialkali molecules, typically see rapid loss due to inelastic molecule-molecule collisions. Loss mechanisms include chemical reactions, and, with molecules that cannot react, "sticky collisions" where long-lived collision complexes are formed, which are then kicked out of the trap by absorbing a trap light photon or by colliding with a third body [72, 74, 143, 174–182]. These inelastic loss rates are quite high, on the order of the universal loss rate $\beta_{in} \sim 10^{-10}$ cm^3 s^{-1}, making any subsequent evaporation impossible. Recently, several theory proposals have been published to use the dipolar nature of these molecules to generate a repulsive interaction between them, such that the molecules can never reach short range, and are shielded from these inelastic losses [80, 82]. These proposals have been proven successful with recent experiments with assembled bi-alkali molecules that have demonstrated evaporative cooling by suppressing the inelastic collision rate using microwave fields [43, 81, 183–185] or static electric fields [83], while enhancing the elastic collision rate. This opens a path towards collisional cooling of molecules, if the density is sufficient to observe collisions.

For directly laser cooled molecules, inelastic collisions have been reported between pairs of CaF molecules in tweezers [72], where subsequent microwave shielding was demonstrated [183], and between molecules and atoms in a magnetic trap [73, 74]. Thus far, however, bulk gases of directly laser cooled molecules have been too dilute for either elastic or inelastic molecule-molecule collisions to be observed. There are two primary reasons for this. First, standard red-detuned molecular MOTs (red-MOTs) have low molecule number ($N \lesssim 10^5$), due to inefficient slowing of the source molecular beam and low capture velocity of the MOT. Second, transfer efficiency from these red-MOTs into ODTs is low (typically $\lesssim 5\%$) [53, 111]. This is due to the large size of molecular red-MOT clouds; this in turn is due to sub-Doppler heating from the Type-II transitions

($N_g = 1 \rightarrow N_e = 0$, where $N_g\{N_e\}$ is the rotational angular momentum of the ground {excited} state) required to be driven for rotational closure of molecular optical cycling transitions [65, 67, 110], limiting typical red-MOT radii to $\sigma \gtrsim 1\,\text{mm}$ and temperatures to $T \gtrsim 1\,\text{mK}$ [54, 56, 59, 62, 69, 124] after a compression stage. The temperature can be reduced further to $\lesssim 50\ \mu\text{K}$ by blue-detuned molasses [53–55, 57], but this does not provide any spatial compression of the molecular cloud.

In this chapter, we will detail our efforts to observe collisions in the ODT. In the first part of this chapter, we will detail improvements to the RF MOT number using improved slowing, and the first hints of collisions with an ODT loaded from the RF MOT. In the second part of this chapter, we will detail dramatically improved loading of the ODT using the previously demonstrated blue MOT. Finally, we will detail the analysis method used to determine that what we are seeing is indeed inelastic collisions in the ODT.

8.2 Collisions in an ODT loaded from the RF MOT

In our first ODT paper [111], we had a peak number density $n_0 \approx 2 \times 10^9\ \text{cm}^3$, and a trap lifetime of around 1 s. This made it very hard to distinguish whether we were seeing collisional effects or just one body loss due to imperfect vacuum. We would need around a factor of 5 increase in density in order to fully distinguish these mechanisms within the trap lifetime. One way to achieve this is to increase the number of molecules loaded in the ODT in some way. Our first attempt at this was the new source as outlined in Chapter 6. This did not lead to a noticeable increase in number, although it did lead to longer lived experiment cycles. The next approach is to improve the slowing by switching to $X \rightarrow B$ instead of the previously used $X \rightarrow A$ transition.

8.2.1 The switch to $X \rightarrow B$ slowing

The primary purpose of switching to the new slowing setup is to break the Λ system caused by slowing and repumping through the same excited state [75, 76]. By repumping $v = 1$ through the same excited state, the number of ground states coupled to the excited state is doubled, leading to a reduction in the maximum scattering rate to $\Gamma/7$, as opposed to $\Gamma/4$ if $v = 1$ was repumped through a different excited state. This would significantly increase the slowing force, making the slowing distance shorter, as well as increasing the capture fraction in the MOT. A promising way to achieve this would be to use the $B\Sigma$ state as the primary slowing transition, and keep repumping the $v = 1$ state through the $A\Pi$ state. This is particularly attractive because the Franck-Condon factors for the $X \rightarrow B$ transition are even more diagonal than the $X \rightarrow A$ transition [79, 186], and will likely require less repumping for slowing. Finally, all the CaF molecule cooling experiments currently use the $X \rightarrow B$ transition for slowing, providing further encouragement that this transition works well [141, 187, 188].

A prior concern for switching to $X \rightarrow B$ was the cascade decay through $\left|A^2\Pi_{3/2}, J = 3/2\right\rangle$, $\left|A^2\Pi_{1/2}, J = 3/2\right\rangle$, and, $\left|A^2\Pi_{1/2}, J = 1/2\right\rangle$ states to the $\left|X^2\Sigma, N = 0, 2\right\rangle$ states, whereby the molecules would be lost [76]. This was fueled by the calculation of the branching ratios to the $A\Pi$ states, which indicated a probability of $\sim 2 \times 10^{-3}$ for decay to these states, making cycling of only 10^4 photons possible before the molecules were lost. However, this calculation was wrong because it did not account for the different transition dipole moments for the different states. The

excited state linewidth for a particular transition goes as $\Gamma_{eg} \propto \omega_{eg}^3 \mu_{eg}^2$, where ω_{eg} is the energy difference and μ_{eg} is the transition dipole moment. The branching ratios can then be calculated using the following equation:

$$b_{\nu'\nu} = \frac{q_{\nu'\nu}\omega_{\nu'\nu}^3 \mu_{\nu'\nu}^2}{\sum_{k=0}^{\infty} q_{\nu'k}\omega_{\nu'k}^3 \mu_{\nu'k}^2} \tag{8.1}$$

where $q_{\nu'\nu}$ are the Franck-Condon factors and $b_{\nu'\nu}$ is the branching ratio. Doing this calculation for the $\left|B^2\Sigma, v=0, N=0\right\rangle$ state gives the branching ratios shown in Table 8.1.

$\left\|X^2\Sigma, v=0\right\rangle$	$\left\|X^2\Sigma, v=1\right\rangle$	$\left\|X^2\Sigma, v=2\right\rangle$	$\left\|A^2\Pi_{1/2}\right\rangle$	$\left\|A^2\Pi_{3/2}\right\rangle$
0.9964	3.54e-3	3.02e-6	2.36e-5	1.37e-5

Table 8.1: Calculated branching ratios from the $B^2\Sigma$ state. The Franck-Condon factors for the $A \to B$ transition are taken to be the same as that for the $X \to A$ transition. Other values for transition dipole moments and energies taken from [79].

This clearly shows that cascade decay from the $A\Pi$ states should not be a problem for slowing, and with a single vibrational repump, we should be able to cycle 10^5 photons easily, sufficient for slowing.

Another concern with switching to this new scheme is the wavelength of the transition. For SrF, the $X \to B$ transition sits at 579 nm, and for a long time, high power systems at this wavelength were unavailable. However, new companies have sprung up that use an IR seed laser for a fiber amplifier which can then be either summed with a different frequency, or doubled to produce shorter wavelengths. In particular, Precilasers, a Chinese company, is very proficient with these techniques and offer a wide variety of short wavelengths at high powers. Our 579 nm light is produced from a doubled 1158 nm seed laser. We wanted this laser to be able to output both the while light spectrum needed for slowing, and also individual sidebands needed for a possible future bi-color MOT [110]. To achieve this, we purchased a MOGLABS cateye ECDL system at 1158 nm, which can output up to 80 mW. The laser has some leaked output light that is used to lock it, and also monitor the frequency on a wavemeter. The main output of the laser is coupled to a fiber EOM (Jenoptik PM1158) which can generate the necessary sidebands for both white light slowing and any future MOTs. After the fiber EOM, we can typically get around 10s of mW. This is then coupled to the precilaser fiber amplifier + doubling system (SSHG-579-1.8W-CW), which can produce up to 2 W of 579 nm light. In our prior slowing work with $X \to A$, the white light was produce by driving the fiber EOM at some low frequency ($\sim$ 4 MHz) with very high modulation depths ($\Delta\phi \approx 40$), requiring up to 30 W of RF power. The 1158 nm fiber EOM is unable to handle this high power, so instead the white light is generated by combining two frequencies, 31 MHz and 9.4 MHz, on a splitter and driving these at lower modulation depths ($\Delta\phi \approx 2, 5$ respectively). This has almost the same effect as a single frequency, but it does leave some holes in the frequency spectrum see Figure 8.1.

The final slowing level diagram is shown in Figure 8.2. The effect of the larger scattering rate was seen immediately in the experiment, with an initial pass netting us a factor of 2 win in molecule number already. On optimizing the slowing timing, it was found that slowing needed to only be applied for 14.5 ms; any additional time would lead to over-slowing and we would lose molecules. Furthermore, because of the long slowing distance from the source to the MOT, we also have to delay the start of the slowing pulse by up to 4 ms after laser ablation to allow the molecules to travel some distance before slowing them. We find that we need a very small magnetic field along

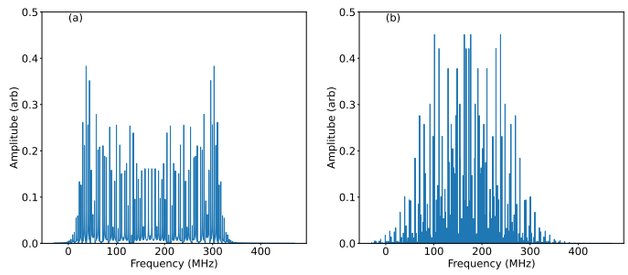

Figure 8.1: X-B slowing profile comparison. (a) slowing profile for a 4 MHz drive with $\Delta\phi \approx 40$ (b) slowing profile for 31 MHz and 9.4 MHz drive with $\Delta\phi \approx 2, 5$ respectively.

the slowing length, and in fact the optimum direction of this field is reversed compared to that needed for the $X \rightarrow A$ slower. It is unclear why this is the case. We almost do not need the $v = 2$ repump, adding that repump only increases the signal by 10%; however, we still keep it in the experiment. We next checked the effect of the slowing power. The MOT number increases dramatically with slowing power up to around 600 mW, after which it seems to saturate and increase slowly with power. This could be because of the long slowing distance making slowing at larger powers not as impactful due to transverse pluming, or it could also be because we have to fiber couple the slowing light to the apparatus, and the fiber has some detrimental effects at higher power. Other CaF experiments are able to make full use of their slowing power by shortening the distance between the MOT and the source, so this could be a possible future improvement. Finally, after optimizing everything, we are able to gain a factor of ~ 10 in molecule number as compared to the $X \rightarrow A$ slower.

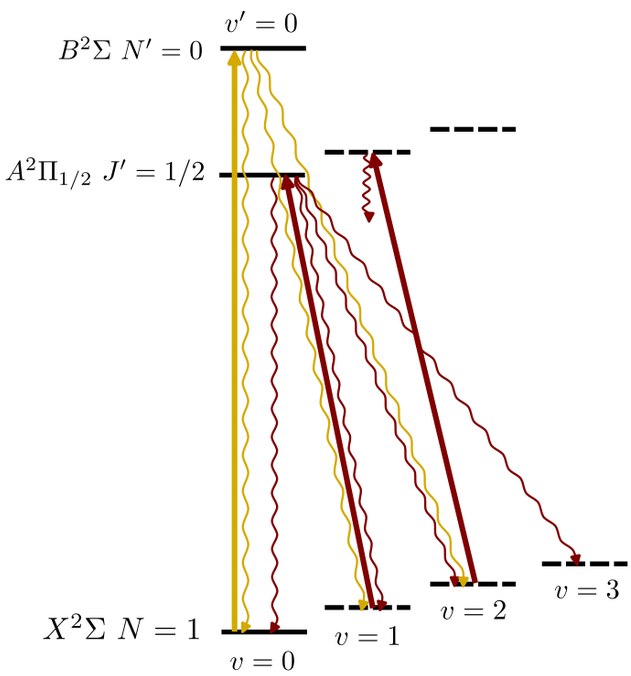

Figure 8.2: $X \rightarrow B$ slowing level diagram. White light slowing is used on the $X \rightarrow B$ transition, along with vibrational repump lasers operating on the $X \rightarrow A$ transition. The $v = 2$ repump is almost not needed, only yielding a gain of around 10% when added.

8.2.2 Λ-cooling at UChicago

To load the ODT, we made use of the fact that Λ-cooling works even in the presence of the large AC Stark shifts in the ODT. At Yale, we were able to achieve temperatures as low as $T \approx 10\,\mu$K (in free space) regularly with Λ-cooling at full power. We could reduce the Λ-cooling power to further cool to $T \approx 5\,\mu$K. We also observed a strong dependence of the temperature on the magnetic field,

with minimum temperatures observed when the Earth's field was roughly canceled out. All of this translated to an ODT temperature of around $T \approx 14\,\mu$K.

However, the same temperatures have not been achieved so far with Λ-cooling at full power at UChicago. We can achieve a free space temperature of $T \approx 30\,\mu$K in the best case, despite all of our efforts at scanning the shim currents and optimizing the Λ-cooling parameters. After the data for this chapter was taken, it was much later found that we were using the wrong set of shim coils, and were applying a very large field (~ 2 G) instead of zeroing the field. Somehow, in spite of this, we were able to observe local minima in the shim scans, indicating there might still be dark states being formed at these large fields. By using the correct coils and zeroing the field, we were able to achieve $T \approx 20\,\mu$K, still not as good as at Yale, but better than before. However, for the all the results in the remainder of this chapter, the old "wrong" shims were used.

8.2.3 ODT setup at UChicago

With this factor of 10 win in molecule number, and with Λ-cooling working decently, we were optimistic that we could get enough density to observe collisions in the ODT, assuming everything else stayed as it was at Yale. We decided to setup the ODT to match the beam size obtained at Yale i.e. around $\omega_0 = 40\,\mu$m. The optical path is the same as before, with the ODT beam being combined with the first pass of the MOT beam using a dichroic mirror. While setting up the ODT, it was found that the Thorlabs high power isolator (IO-10-1064-VHP) used after the ODT fiber had developed a large burn mark on the outer face of the faraday rotator crystal. This was causing the beam to focus to a smaller waist, and also lead to a distortion in the beam shape. We replaced the isolator with a Newport isolator (ISO-FRDY-08-1064-W) rated for much higher intensities. We had to then swap out the 1:5 beam expander to a 1:3 beam expander (made using a telescope with $f_1 = -50$ mm and $f_2 = 150$ mm) to get the right beam waist at the focus. With this new telescope, we profiled the beam again and got the trap depth profile shown in Figure 8.3. This new profile has $\omega_x, \omega_y, \omega_z = 2\pi \times (2641, 2641, 12.5)$ Hz with a Rayleigh range of $z_R = 5.7$ mm.

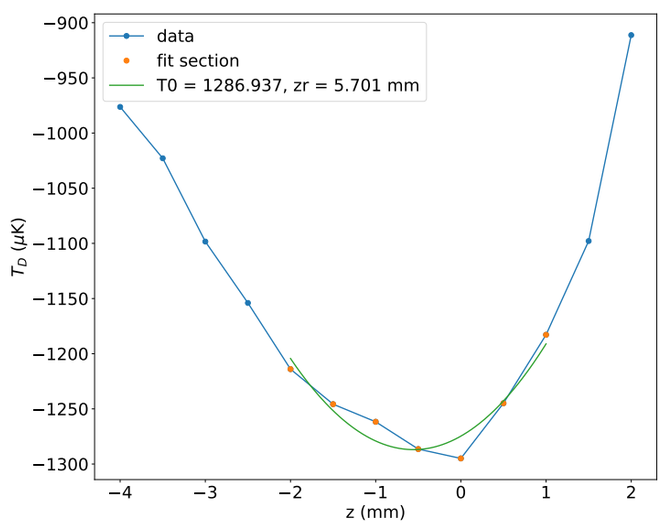

Figure 8.3: ODT beam profile at UChicago. $T_0 = 1.3$ mK and $z_R = 5.7$ mm. This closely matches the old profile at Yale.

The old 1:5 beam expander was very compact, meaning it was comprised of some very short focal length lenses, and so astigmatism was a big problem. The new expander is more forgiving

and makes it easier to align the beam. The beam is still slightly astigmatic, but it was decided to go ahead with this.

8.2.4 ODT optimization

Our biggest win at Yale was the use of the ODT polarization to compensate for the AC Stark shifts caused by the Lambda cooling beams due to the intensity mismatch of the Lambda cooling beams [111]. Specifically, we found the lowest temperature and highest number when the ODT polarization matched the weaker of the Lambda cooling beams it was co-propagating with. This means the ODT polarization wants to be circular for optimal trapping.

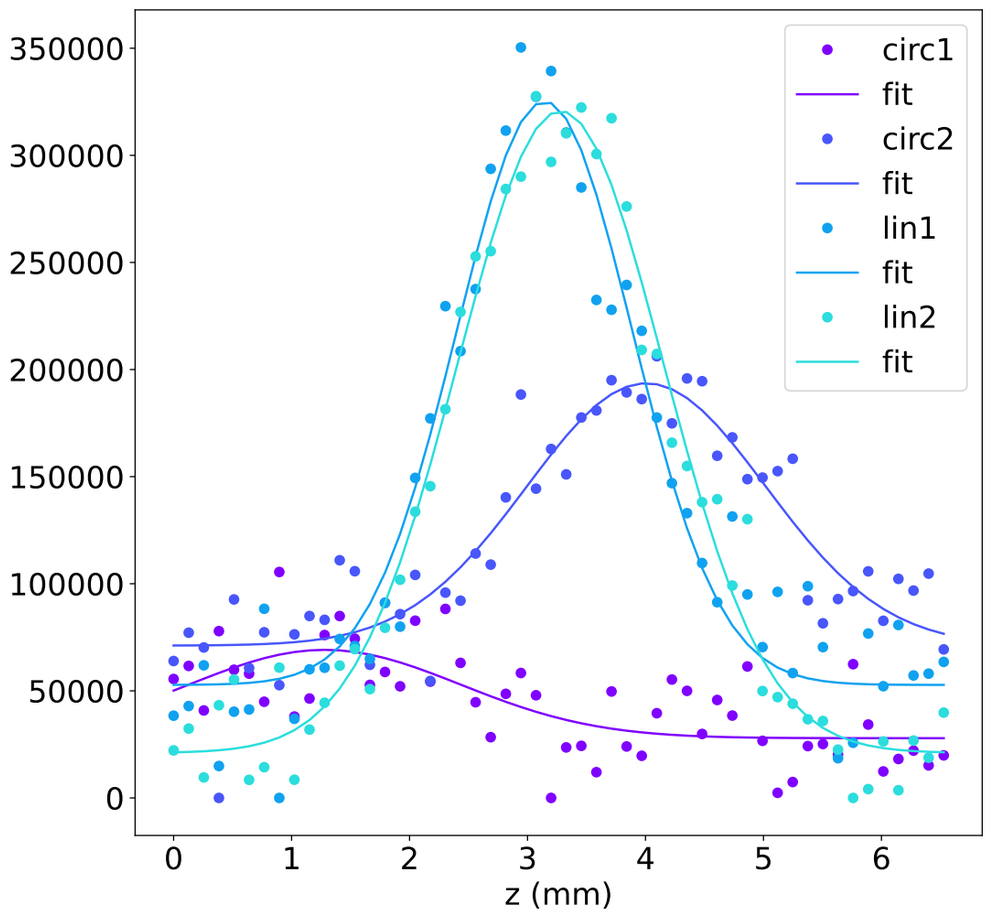

Figure 8.4: Polarization dependence of ODT at UChicago. Optimal trapping is only observed for linear polarization, and circular polarization leads to very hot molecules.

However, here at UChicago, this dependence on polarization seems to have flipped, and we now get optimal trapping at linear polarization (see Figure 8.4). Both linear polarizations work, and one circular polarization is clearly better than the other, but they are both still pretty bad. This could be because our intensity imbalance is completely different, or that our fields were not zeroed well, but so far the source of this change has not been traced. So we set the ODT polarization to be linear and then tried to optimize the loading into the trap. Figure 8.5a shows the results for the temperature vs Raman detuning. We find a few different frequencies where the temperature is minimized, namely around $\delta \approx -2.5, -0.3, +0.1$ and $+2$ MHz. This becomes somewhat clearer when we look at the AC Stark shifts for the two $F = 1$ states used for Λ-cooling, as shown in Figure 8.6. For linear polarization, there are 4 different detunings possible where there will be good dark states both inside and outside the trap, and these are somewhat close to what we measure.

In order to optimize loading to see collisions, we care not only about final temperature, but final number as well. A reasonable metric to look at is the amplitude/width for each detuning, which is a proxy for density. This is shown in Figure 8.5b, with a maximum at $\delta \approx -0.3$ MHz. This is what we will use to load from here on. By optimizing the loading parameters, we can reach a final temperature of around 27 μK.

It is also possible that the optimal loading parameters for Λ-cooling may not be the optimal parametrs to cool the loaded molecules. We see this in free space as well, where the final temperature is coldest for lower laser power. Thus, we also tried to implement a second Λ-cooling pulse, after loading and letting the untrapped molecules fall out of the trap. The second Λ-cooling pulse is applied for 40 ms and does indeed result in some additional cooling. Here, it was found that δ wanted to be the same, Δ wanted to be further blue by around 6 MHz, and the final laser power was optimized at around 35% of the initial power, shown in Figure 8.7.

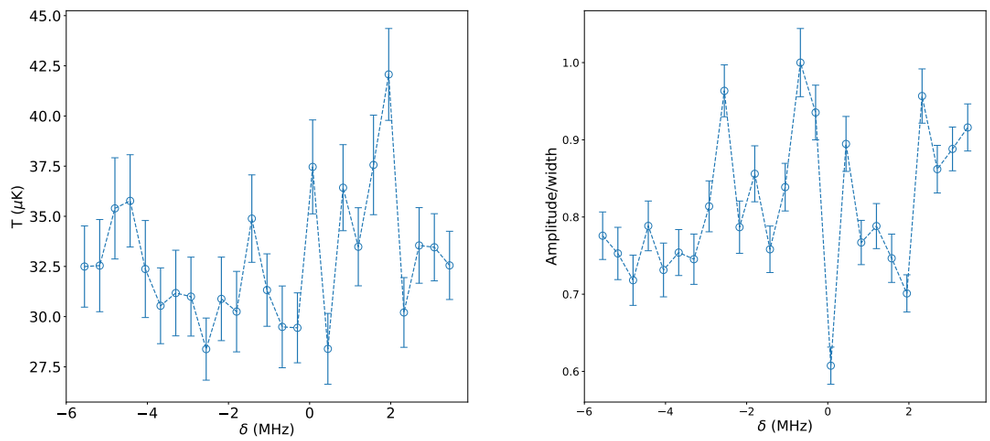

(a) Temperature vs δ (b) Amplitude/width vs δ

Figure 8.5: ODT optimization for the Raman detuning δ. (a) shows that the ODT temperature is minimized for a few different detunings. (b) However, the number/width, which is a measure of density, is maximized for $\delta \approx -300$ kHz.

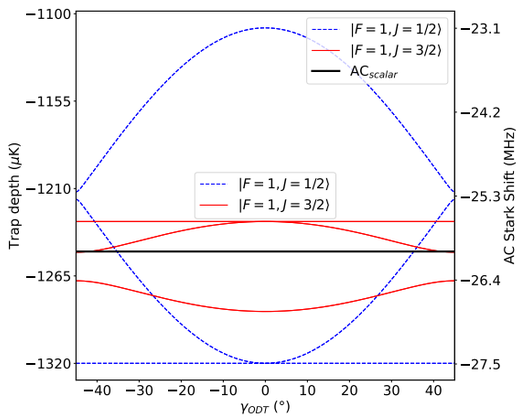

Figure 8.6: AC Stark shifts for the new ODT beam as a function of ODT polarization. The trap depth, and hence the Stark shifts are larger.

We also see that the ODT loading is satured at around 90% of the full ODT power available, and additional power does not really contribute to loading. This is because of the large red MOT cloud size, where beyond a certain point, the large trap depth does nothing to capture the molecules at the edges of the cloud. The only way for those molecules to be captured is to diffuse into the ODT, which is very unlikely for molecules at the edge of the cloud.

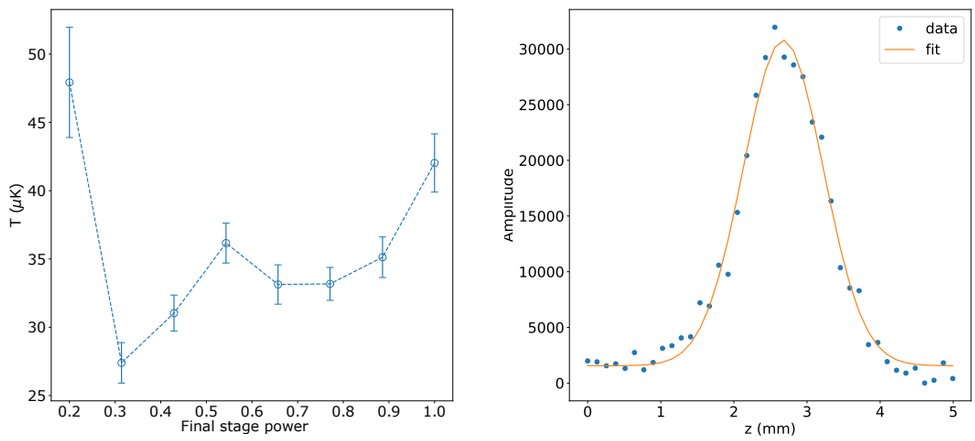

(a) Temperature vs final laser power

(b) ODT profile after second cooling pulse

Figure 8.7: ODT optimization for the second Λ-cooling pulse. The final laser power is around 35% of the initial loading power, and the final width is around 0.55 mm as opposed to the initial 0.62 mm.

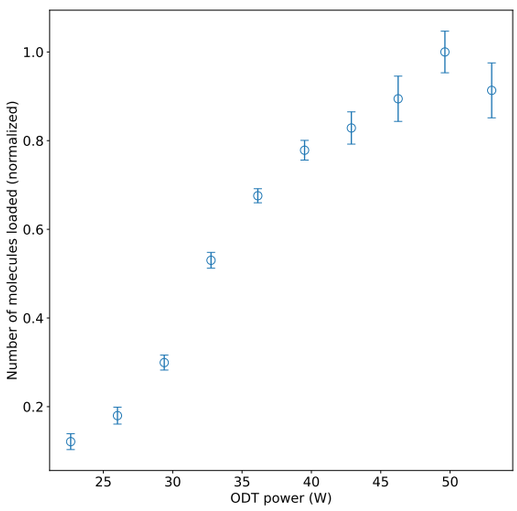

Figure 8.8: Red MOT ODT loading vs ODT power. ODT loading saturates at 90% of the total power available.

Under optimal conditions, we can now load around 4-4.5% of the RF MOT (around 700 molecules) and cool to a final temperature of around $T = 23$ μK in the ODT. This corresponds to a peak density of $n_0 \approx 1 \times 10^{10}$ cm^{-3}.

8.2.5 First hints of collisions

How can we tell if we are seeing collisions between molecules? Since we expect most of these collisions to be inelastic, we can look at a lifetime curve where we can track the number of molecules remaining in the trap vs time. Then, density dependent losses can be modeled using the two-body loss rate equation, with evolution of the number density n given by:

$$\dot{n} = -\frac{1}{\tau}n - \beta n^2, \tag{8.2}$$

where τ is the one-body lifetime and β is the two-body loss rate coefficient. This equation tells us what to expect when looking at lifetime data. If the starting density is high and we are in the collisional regime, we expect to see fast loss, dominated by β, whereas over time, as the molecule number depletes and the density reduces, we should settle to a simple exponential decay with

timescale τ. Thus, if we see two distinct timescales in the lifetime data, we are on the right track!

To convert eq. 8.2 to a number evolution, we assume a Gaussian spatial distribution and define an effective volume ($V_{\text{eff}} = (2\sqrt{\pi})^3\sigma_x\sigma_y\sigma_z$) occupied by the molecules [177]; here z is along the axial direction of the ODT, and x (y) is along the transverse direction in (perpendicular to) the imaging plane. This allows us to integrate over the volume to obtain:

$$\dot{N} = -\frac{1}{\tau}N - \frac{\beta}{V_{\text{eff}}}N^2. \tag{8.3}$$

If the spatial distribution is constant in time, eq. 8.3 has an analytical solution:

$$N(t) = \frac{N_0}{\left(1 + \frac{\beta\tau N_0}{V_{\text{eff}}}\right)e^{-t/\tau} - \frac{\beta N_0 \tau}{V_{\text{eff}}}}. \tag{8.4}$$

Our imaging system cannot resolve the transverse radius (σ_x) of the molecular cloud in the ODT. We also cannot observe properties in the y-direction. We do directly measure the cloud radius along its axial direction (σ_z), as well as the temperatures T_x and T_z. We then infer σ_x using the calculated trap depth, measured ODT beam profile, and value of T_x, and assume $\sigma_y = \sigma_x$ by symmetry. To calculate V_{eff}, we assume that the trap is harmonic and the spatial density is given by:

$$n(\mathbf{r}) = n_0 \exp\left(-\frac{x^2}{2\sigma_x^2}\right) \exp\left(-\frac{y^2}{2\sigma_y^2}\right) \exp\left(-\frac{z^2}{2\sigma_z^2}\right). \tag{8.5}$$

The effective volume is given by $V_{\text{eff}} = \int n(\mathbf{r})d^3\mathbf{r}/n_0 = (2\sqrt{\pi})^3\sigma_x\sigma_y\sigma_z$. We do not have enough resolution to measure the transverse radius σ_x, and our observations give no direct information about σ_y. However, we have measured the ODT laser beam profile, and find a good fit to a Gaussian with $1/e^2$ intensity radius $\omega_0 = 38(3)$ μm. We have also calculated the ODT trap depth to be $U_T = 1.3(1)$ mK [111], and we measure the transverse temperature of molecules in the ODT using time-of-flight (TOF) expansion technique to be $T_x = 25(3)$ μK. From this we deduce the transverse radius $\sigma_x = \sqrt{\omega_0^2 T/4U_T} = 2.5(3)$ μm, and assume $\sigma_y = \sigma_x$.

Figure 8.9 shows an example of one such measurement made from the ODT loaded after the RF MOT. In this experiment, we load the ODT for 130 ms using Λ cooling, after which untrapped molecules are allowed to fall out of the imaging area for 150 ms. This long fall period is necessary for the MOT to recapture only the trapped molecules. After this fall time, we apply a second pulse of Lambda cooling for 40 ms. This was found to further cool molecules in the trap down to around 25 μK. This defines the starting point of the MOT recapture measurement. At this point, we estimated that we have 700 molecules trapped, with a temperature of 25 μK. This leads to $\sigma_x, \sigma_y \approx 2.5$ μm and the measured $\sigma_z = 0.7$ mm and a starting density $n_0 \approx 1 \times 10^{10}$ cm^{-3}.

Here each point is an average of 10 images, and the order of the images was randomized. The images are then fit to a 2D Gaussian and the number is extracted from the fit amplitude, and the errors are from the fits. We can clearly see two different time scales at play here, and a simple exponential decay cannot account for this. Fitting to the two-body decay model with V_{eff} calculated as above, we can extract a two-body loss rate of $\beta = 4(1) \times 10^{-10}$ cm^3/s.

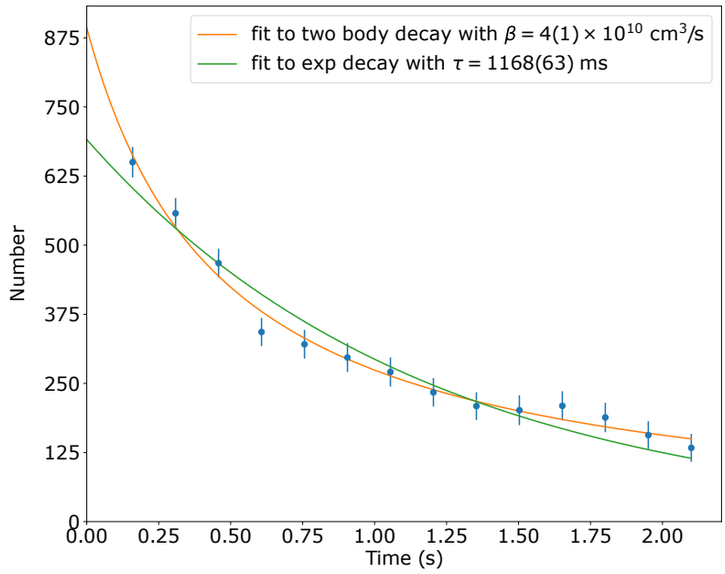

Figure 8.9: First hints of collisions using MOT recapture. One can see two times scales at play here, as is typical for two body processes, thus suggesting some evidence of collisions.

8.2.6 Better lifetime measurement with MOT recapture

At this point, we decided to try and properly calibrate everything to get a good measurement. To that end, we implemented a new data taking procedure. The experiment protocol is the same up to the second Λ cooling pulse. Then, we made the following changes:

1. A list of ODT hold times was chosen, such that we had more points at shorter times and a few points at very long times to get the full picture. The starting point was defined to be $t_0 = 20$ ms after the cooling pulse, which is roughly ω_z^{-1} to allow for any dynamics from the second cooling pulse to complete before the measurement.

2. This list was then randomly shuffled.

3. For each point in this new list, we made two recapture measurements, one at $t = t_0$ and one at $t = t_i$. Each of these measurements is an average of 15 images, and the images for t_0 and t_i were taken in random order within this sequence.

4. Finally, the ratio of N_i/N_0 was compared for all the times in the initial list.

We also made a better measurement of the starting conditions. We find that the capture fraction of molecules at $t_0 = 20$ ms is around 1.5%, which means a starting number of around 700 molecules, which means a starting density of around $n_0 \approx 1 \times 10^{10}$ cm^{-3}.

The lifetime measurements are shown in Figure 8.10. Here, the errors are calculated from the standard error in the fits. Here, one can make two observations. First, there is a fast decay for short times, out to around 600 ms, which is mainly governed by two body loss due to collisions. After this, one can see a long slow decay, extending out to around 4 s, which is the one body limited lifetime. We also decided to fit for β/V_{eff} instead of just β since it is easier to incorporate the uncertainty in the calculated V_{eff}. Based on the previous section, $V_{\text{eff}} \approx 1.9 \times 10^{-7}$ cm^3. From the extracted β/V_{eff} from the fit, this leads to $\beta \approx 3 \times 10^{-10}$ cm^3/s.

This is even more clear in the semilog plot shown in Figure 8.11. In this plot, it is quite clear that the simple exponential fit does not account for the data. The presence of the long slow decay is the key to showing this.

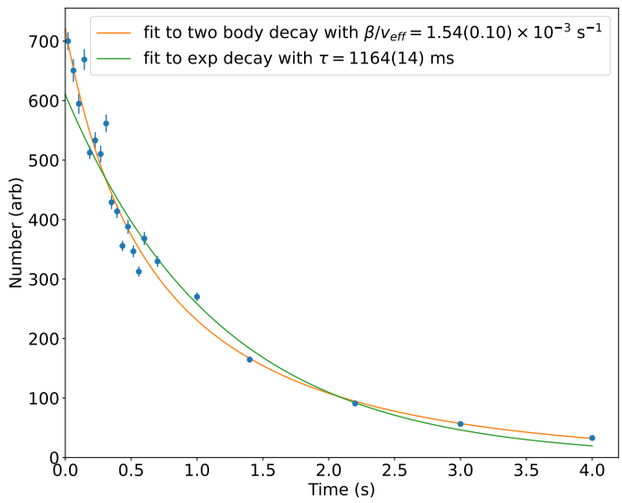

Figure 8.10: Lifetime measurement using MOT recapture showing two-body loss. A fit to only an exponential decay does not account for the fast initial loss seen in the data.

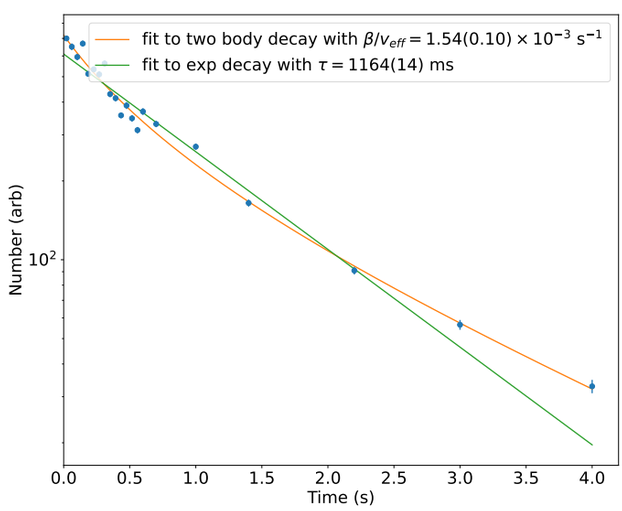

Figure 8.11: Semi log plot of the lifetime measurement shown in Figure 8.10. A fit to only an exponential decay does not account for the fast initial loss seen in the data.

8.2.7 Effect of Helium buildup

We were concerned that the lifetime measurement could be skewed if there was Helium build up in the MOT chamber, especially for the longer time data points. This in theory should not be a problem since we have an in-vacuum shutter to block ballistic Helium, but we decided to test this nonetheless. In order to deduce if this was a problem, we took a lifetime curve following the prescription laid out before. For a few random times in the list, we additionally took data by adding a "wait" column after the image was taken, for a time $t_{\text{wait}} = 4000 \text{ ms} - t_i(t_0)$. The results of this are shown in Figure 8.12. There is a slight difference in the amplitudes at higher times, but it is also hard to keep the experiment well controlled for these long shutter times, especially because of SF6 build up. As a result, it was decided to not use this method for the other lifetime measurements.

8.2.8 Better lifetime measurement and estimating β

For the measurement shown in Figure 8.13, we tried to be thorough with the calibration pre and post data collection, to try and keep the experimental conditions the same throughout. We estimate that we started out the measurement with around 700 molecules in the ODT, but ended the measurement with around 650, mainly due to the ablation degrading over time. We started

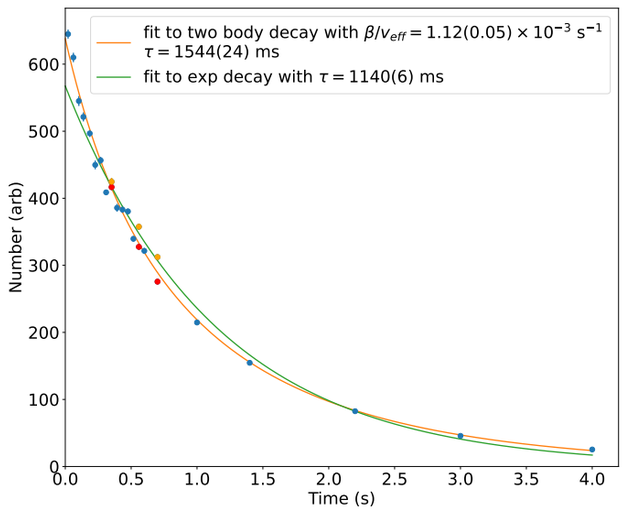

Figure 8.12: Effect of Helium build up. Lifetime data where points in red are without the compensation shutter and points in orange are with the compensation.

out with a temperature of around $T \approx 23 \pm 3$ μK (as determined by TOF), and a starting axial width of around $\sigma_{\text{ODT,measured}} = 0.55 \pm 0.01$ mm, and we already calculated $\sigma_{x,y} = 2.5 \pm 0.3$ μm. This gives $V_{\text{eff}} = 2.3 \pm 0.7 \times 10^{-7}$ cm^3. From the extracted value of β/V_{eff} from Figure 8.13, we find:

$$\beta = 3.2 \pm 0.5 \times 10^{-10} \text{ cm}^3 \text{ s}^{-1}$$

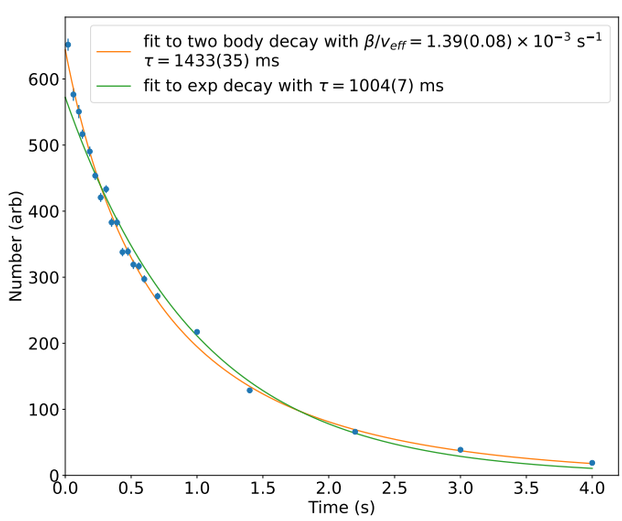

Figure 8.13: Final lifetime curve for the ODT loaded from the RF MOT, with better calibration and fitting for β/V_{eff}.

This value of β is of the same order as that measured for 2 CaF molecules in tweezers, $\beta_{\text{CaF}} = 4.2 \times 10^{-10}$ cm$^3/s$ [183]. The data also gives an estimated value of the bare lifetime of $\tau \approx 1433 \pm 35$ ms, which is not quite as long as the blackbody lifetime of around 3.5 s.

This first hint of collisions is quite encouraging since this suggests that if we can successfully load an ODT from the blue MOT, we should be able to win a lot in density, and thus more easily resolve these collisions.

8.3 Collisions in an ODT loaded from the blue MOT

As demonstrated in Chapter 7, we are able to trap $N \approx 1.2 \times 10^4$ molecules at $T = 200$ μK with a size of around $\sigma_{ax}, \sigma_{rad} = 220, 240$ μm. After this, we have demonstrated that we can cool to

around $T = 30$ μK in free space by applying a Λ-cooling pulse. This is a pretty good starting point for loading the ODT. Based on geometrical arguments alone, we expect to load around $\frac{38\times38}{(210\times260)} \approx 27\%$, i.e. around 4000 molecules in the ODT.

8.3.1 ODT loading from the blue MOT

We start loading the ODT immediately after the blue MOT is fully loaded and compressed, i.e. 30ms after the blue MOT. We use Λ-cooling to load in the trap. There are significant differences between loading from the blue vs red MOT, which are detailed below.

8.3.1.1 Loading timescale

The time taken to load is very different as compared to the red MOT. Because the blue MOT cloud is so much smaller, it takes a much shorter time to load the ODT. We find that the optimal load time is around 20ms (see Figure 8.14), and we can also observe the competition between loading and two body loss by trying to load for longer times. These losses can be either inelastic collisions or light-induced collisional losses, but the fact that we already see losses here is quite encouraging. The point at $t = 0$ corresponds to the initial fraction loaded just by the spatial overlap with the ODT when the beam is turned on.

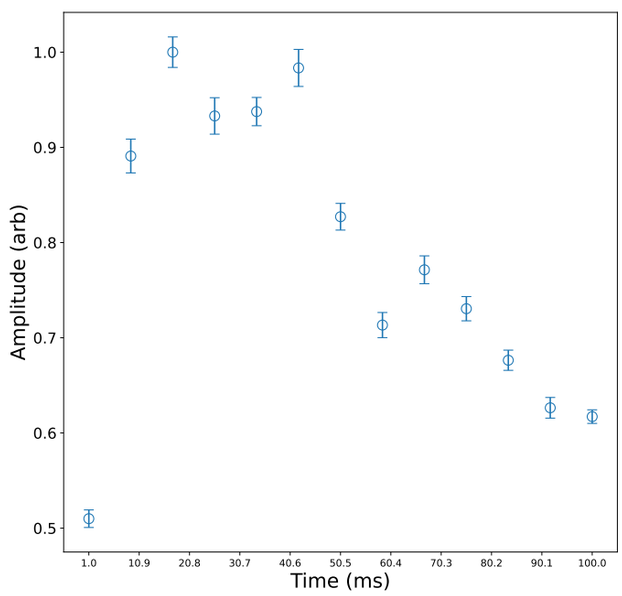

Figure 8.14: Loading the ODT from the blue MOT. The loading is much faster, and is completed in around 20ms. One can also observe the competition between loading and two body loss since longer load times are detrimental.

8.3.1.2 ODT power dependence

The loading shows a linear dependence on the ODT power, with best numbers achieved at maximum ODT power, with no signs of saturation even at the maximum power, very unlike the ODT from the RF MOT, which saturates at 90% power (see Figure 8.15). This suggests that perhaps faster molecules are escaping from the edges of the trap (because of the higher temperature) which the initial Λ-cooling pulse is not able to overcome, i.e. these molecules are lost even before Λ-cooling can dissipate enough momentum to keep them trapped. This would be more important for the tighter blue MOT because of the small MOT trap volume - any molecules that are initially lost have a very low chance of making it back to the ODT within the 20 ms loading window. By

contrast, in the red MOT, the density is low enough that other molecules from the molasses cloud can slowly dissipate into the ODT over the 150 ms loading scale, and thus after a point, higher ODT powers does not really matter. This means that even more molecules can be loaded from the blue MOT by simply increasing the trap power to increase the trap depth.

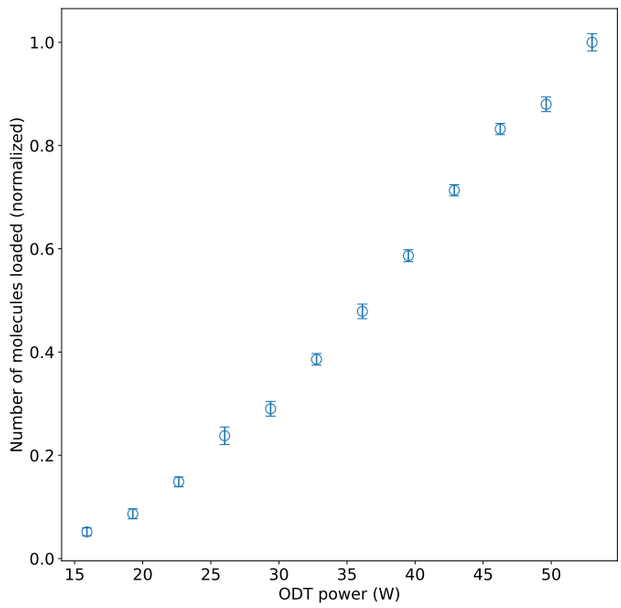

Figure 8.15: ODT loading vs ODT power. Here, unlike loading from the red MOT, the loading is optimized at full power, and does not show any sign of saturation. This indicates that faster molecules could be escaping from the edge of the trap, and that higher loading efficiencies might be achieved if we can boost the power to increase the trap depth.

8.3.1.3 Cooling parameters

We also find that different Λ-cooling parameters are ideal for loading from the blue MOT as compared to loading from the RF MOT. Specifically, the Raman detuning now wants to be very slightly blue detuned (see Figure 8.16). This change is probably because the blue MOT is optimized with the retro being substantially off-center. Thus, the laser intensities/polarization gradients are probably different, which could lead to this difference. We also see a few different minima, as we saw with the RF MOT, with the best loading at $\delta = +2\pi \times 0.1$ MHz. Similar scans show that the loading is optimized at a single photon detuning $\Delta = 2\pi \times 22$ MHz.

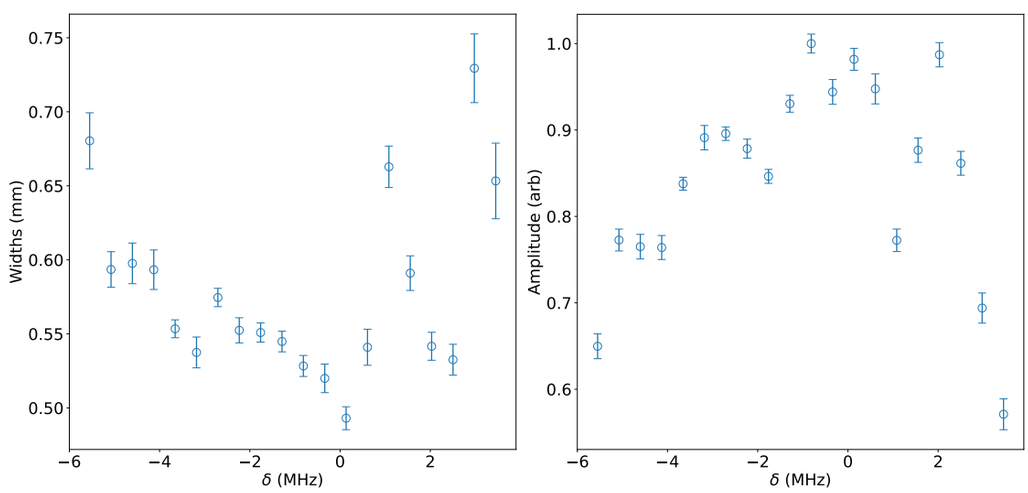

(a) Width vs δ

(b) Amplitude vs δ

Figure 8.16: Raman detuning scan for loading in the ODT. The width is minimized and the amplitude is maximized at $\delta \approx +2\pi \times 0.1$ MHz.

8.3.1.4 ODT alignment

Owing to the small size of the blue MOT, the ODT loading is highly susceptible to the beam alignment and the overlap with the blue MOT. We are trying to align a 40 μm waist beam with a 200 μm cloud, which means the pointing and the overlap need to be exceptionally good. Fortunately, we can use the ODT signal itself to optimize alignment, as shown in Figure 8.17. The goal of the optimization is to minimize the cloud width, which would only be true when the ODT and the blue MOT centers overlap perfectly. If the focus is on either side of the blue MOT center, i.e. if the focus is not perfectly overlapped with the center, we can see an elongation in the ODT cloud profile in that direction. During the loading period, where Λ-cooling is on, the blue MOT profile is essentially frozen in space, but once the light is turned off (to let the untrapped molecules fall out), the molecules are free to move. If the center of the ODT is not perfectly aligned with the blue MOT, then the cloud will be slowly accelerated towards the ODT center, and we see this elongation in the profile. By moving the focus to the "right" in Figure 8.17 we can successively reduce this elongation, until the cloud width is minimized and looks gaussian.

Similarly, the vertical alignment can be adjusted to minimize the width. When the cloud is misaligned vertically, the ODT is overlapped with a gaussian of larger width (corresponding to the edges of the blue MOT), and thus the ODT cloud width appears larger. In practice, one can optimize vertical alignment by looking at images in situ and optimizing for maximum signal/minimum width. It is also important to wait a few shots after adjusting alignment, since the mirrors tend to have a relaxation time, and the ODT signal can fluctuate during this time. This is best done with 2x2 binning, where the blue MOT cloud can be well resolved.

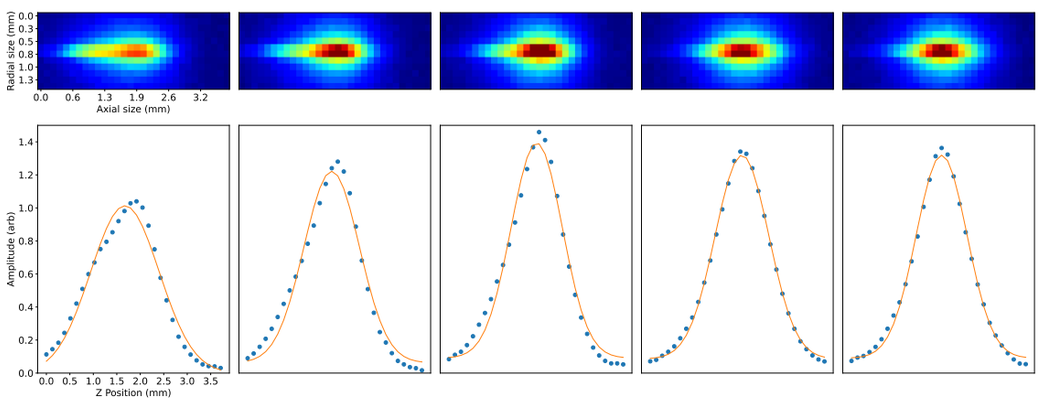

Figure 8.17: ODT loading vs ODT alignment. Series of images from left to right showing how the focus spot of the ODT is optimized by moving the translation stage to the right in the image. When the focus is well overlapped with the blue MOT center, the cloud shows a gaussian profile and has the maximum amplitude.

Under optimal conditions, we load $N \approx 4000$ molecules in the ODT ($\sim 30\%$ capture efficiency), at $T \approx 40$ μK, and $n_0 \approx 3.4 \times 10^{10}$ cm^{-3}. This is an order of magnitude higher efficiency than achieved when loading from type-II red-MOTs [53, 111]. The trap cloud width is around $\sigma_{\mathrm{ODT,measured}} \approx 420$ μm which corresponds to a real width $\sigma_{\mathrm{ODT}} = 595\,\mu$m.

8.3.2 Temperature

Now that we have a lot more molecules in the ODT, we can easily use TOF to measure temperature. We find that the ODT temperature is $T \approx 40$ μK, similar to the free space temperature, and hotter than the ODT loaded from the RF MOT (see Figure 8.19). We also find that the second Λ-cooling

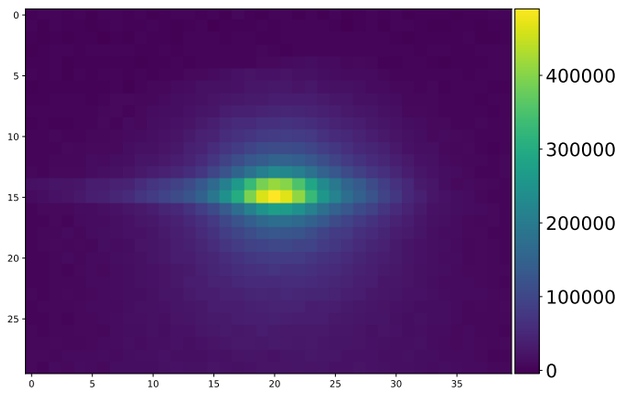

Figure 8.18: Image of the best ODT loaded from the blue MOT. We can load around 30% from the blue MOT into the ODT, and we get around 4000 molecules in the trap.

pulse does not really help and the final temperature remains the same.

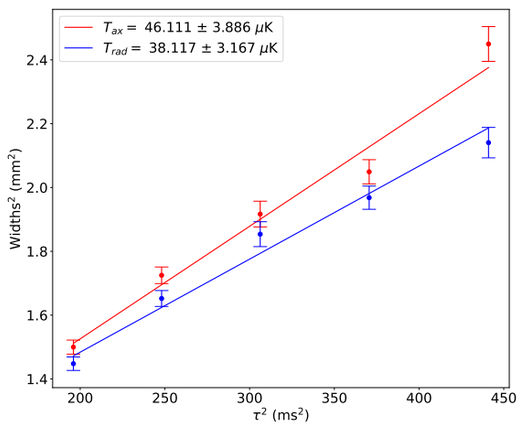

Figure 8.19: Blue MOT ODT TOF. Owing to the larger number of molecules in the ODT, we are easily able to obtain temperatures via TOF, consistent with the measured widths.

This means that our starting density of $n_0 \approx 3.4 \times 10^{10}$ cm^{-3}, is only around 3x higher than the RF MOT, mainly due to the higher temperature.

8.3.3 Lifetime and collisions

With these starting conditions, we look for evidence of inelastic molecule-molecule collisions. To study collisional loss, we perform measurements of the number of molecules remaining in the trap (N) as a function of the hold time (t_h). For all of these measurements, we load the ODT for 20 ms, then let untrapped molecules fall out of the trap by turning off the Λ-cooling light for 32 ms. This defines $t_h = 0$. We then measure the remaining number at time t_h, either by imaging in-situ with the Λ-cooling light (for all points $t_h < 1$ s) [53], or by recapturing in the compressed red MOT and imaging in-situ (for all points $t_h \geq 1$ s). The scattering rate for each method is determined by comparing the fluorescence counts to those from a free space image (2 ms exposure) at $I \sim 170$ mW/cm^2. We assign un-corrrelated uncertainties to each $N(t_h)$ data point by adding in quadrature contributions from fit uncertainties, the shot-to-shot fluctuations in the initial number, and uncertainties in the ratio of the extracted number between the two imaging methods as detailed below.

To get an accurate measurement of the two-body loss rate, it is imperative to get early time data when the molecule number is high and the loss is mostly density dependent, especially in our system where the vacuum limited lifetime is short ($\tau \approx 1.3$ s). As described before, after loading the ODT for 20 ms, we let untrapped molecules fall out of the trap. The earliest time where we

can distinguish the trapped molecules from the untrapped falling molecules is $t_{\text{fall}} = 32$ ms, where the untrapped molecules are not yet out of the field of view of the camera, but have fallen enough that the trapped molecules can be distinguished by an appropriate choice of the region of interest (ROI) (see Figure 8.20). We define $t_h = 0$, the starting point of the lifetime measurement to be after this fall period.

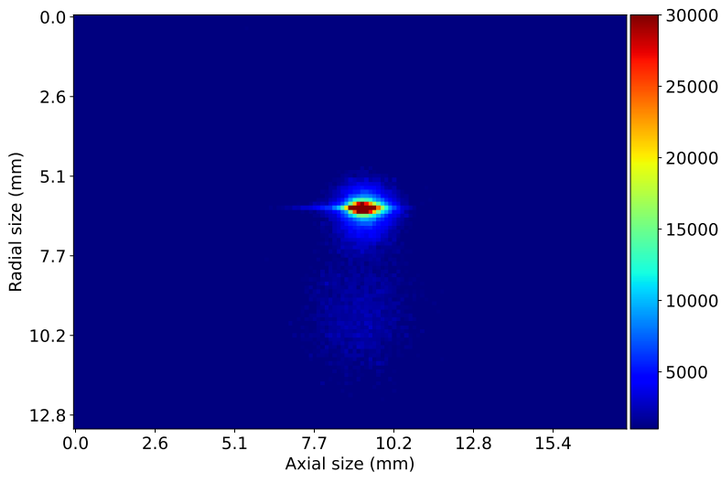

Figure 8.20: 32 ms drop image. Two distinct clouds are visible in the image, one very faint falling cloud at the bottom of the image, and the central ODT cloud. With a 32 ms fall period, the falling molecule cloud is just about distinct from the ODT cloud, and can be appropriately distinguished by choosing a ROI centered on the ODT cloud.

We also divide our data in two chunks. For $t_h < 1$ s, where the number is high and the signal to noise ratio is good, we take a Λ-image of the molecular fluorescence by exposing the camera for 10 ms while Λ-cooling light (at $I \sim 57$ mW/cm^2) is applied to molecules in the ODT. For the rest, where the number is low, we turn off the ODT and recapture the remaining molecules into the compressed red-MOT and take an image of the molecular fluorescence by exposing the camera for 50 ms (at $I \sim 10$ mW/cm^2). This method can only be used when $t_{\text{fall}} \geq 152$ ms, in order to not recapture molecules untrapped by the ODT. For each method, we determine the scattering rate by comparing the fluorescence counts of an image taken immediately after $t_{\text{fall}} = 32(152)$ ms (with maximum molecule number) to the counts from a 2 ms exposure image at laser intensity $I \sim 170$ mW/cm^2, where the scattering rate is measured, Figure 7.6.

The Λ-image data is taken as follows. We first generate a list of hold times, and randomize this list. For each hold time t_i in this list, we take a set of 30 images, 15 each for $t_h = t_i$ and for $t_h = 0$. The order of these 30 images is randomized as well. We then extract the fluorescence counts by integrating the image along the radial direction, then fitting it to a 1D Gaussian plus constant offset; the fluorescence counts are extracted from the Gaussian integral. By comparing the values of N_0 inferred from the $t_h = 0$ data in each set of data at different values of t_i, we are able to measure the shot-to-shot drifts in starting number over the duration of the entire data set. We find that the various N_0 values have a standard deviation $\sim 8\%$, and we add this in quadrature to the other uncertainties in the number at each data point.

For the MOT recapture data, we follow the same procedure, with the additional drop time added. In addition, we take a set of images at a few intermediate t_h using both methods to compare the extracted number and we find that the ratio of the number extracted from MOT recapture to the number from Λ-images is $N_{MOTrecap}/N_{\Lambda-image} = 1.01 \pm 0.10$. We hence include an additional 10% uncertainty in molecule number for all points with $t_h \geq 1$ s, where the MOT

recapture method is used.

In addition to these uncertainties, there is also ambiguity in the determination of the overall scattering rate, because of differing reported values of the branching ratio $\left|A^2\Pi_{1/2}, J = 1/2^+, v = 0\right\rangle \rightarrow \left|X^2\Sigma, N = 1, v = 3\right\rangle$ for SrF molecules [59, 79]. This results in an overall scale factor uncertainty in the molecule number when converting from fluorescence counts to number. To account for this, we use the average of the branching ratios from Refs. [59, 79], and half their difference as its uncertainty. This leads to $\sim 17\%$ uncertainty in the determination of the starting number. We further take into account uncertainties in the calibration of the light collection optics for the imaging setup ($\sim 10\%$) and we find a combined uncertainty $\sim 25\%$ in the overall starting number. We emphasize that this is different from shot-to-shot fluctuations, as this uncertainty affects each data point in the same direction.

We observe a fast initial loss, followed by a long slow decay, as is characteristic of two-body loss processes (see Figure 8.21). As mentioned before, our imaging system cannot resolve the transverse radius (σ_x) of the molecular cloud in the ODT. We also cannot observe properties in the y-direction. We do directly measure the cloud radius along its axial direction (σ_z), as well as the temperatures T_x and T_z. We then infer σ_x using the calculated trap depth, measured ODT beam profile, and value of T_x, and assume $\sigma_y = \sigma_x$ by symmetry. For our blue MOT temperature, we find that $\sigma_x = \sigma_y = \sqrt{\omega_0^2 T/4U_T} = 3.3(4)$ μm.

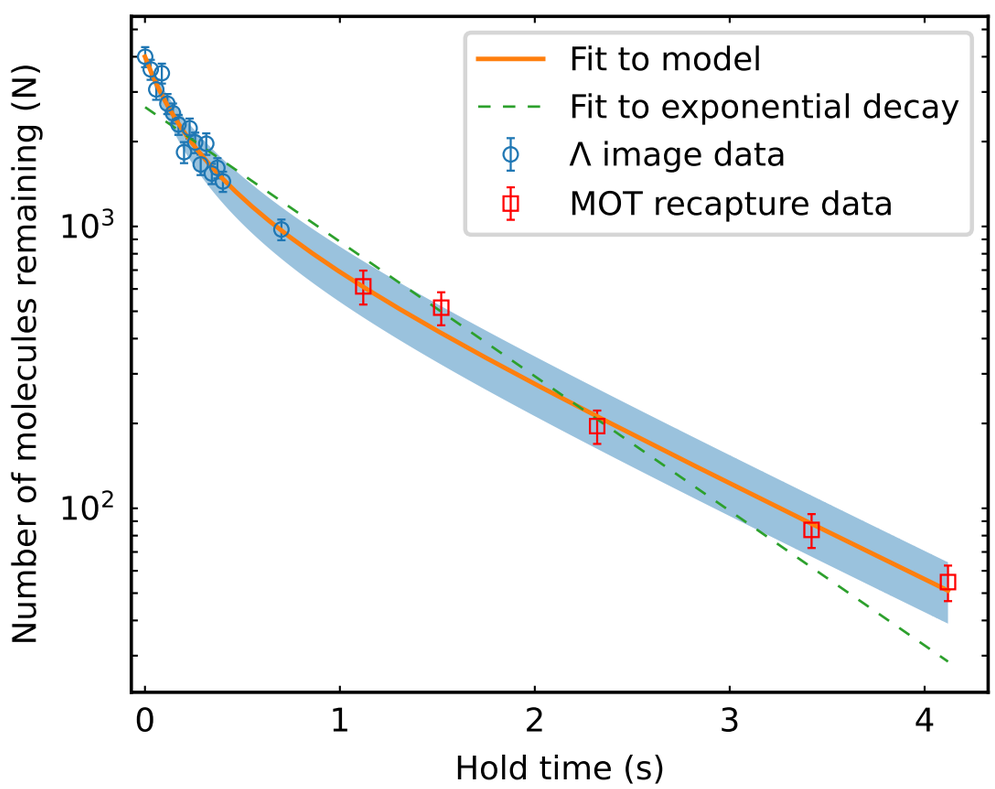

Figure 8.21: Lifetime curve for ODT loaded from blue MOT, showing number of molecules remaining in the trap as a function of trap hold time. Each point is an average of 15 images, and the error bars account for uncertainties as described in the text. Data for $t_h < 1$ s are Λ-images (blue circles) and the rest are MOT recapture images (red squares). The data shows a clear deviation from an exponential decay, a classic signature of two-body loss. By fitting to a model where σ_z is increasing linearly with time, we extract a two-body loss rate coefficient $\beta = 2.7^{+1.2}_{-0.8} \times 10^{-10}$ cm^3 s^{-1}, and a one-body loss rate $\tau = 1.3(1)$ s. The shaded area indicates the uncertainty in β.

We observe that the measured value of σ_z increases from its initial value linearly with hold time. We suspect this results from the ODT beam profile changing due to thermal lensing from

optics along the beam path. In particular, we observe that the trap center is dragged along the axial direction in a non-adiabatic manner as a function of hold time. The low trap frequency along the axial direction ($\omega_z = 2\pi \times 13$ s^{-1}) is not able to support this motion, leading to heating over time. We then use the measured value of σ_z to deduce the axial temperature of the molecules in the ODT, T_z^σ, by following the procedure outlined in [111], where $T_z^\sigma = 2U_T\sigma_z^2/z_R^2$, where z_R is the Rayleigh length of the trap. We find that this inferred axial temperature increase is consistent with the directly measured increase in T_z (Figure 8.22(a)). We do not see any increase in the measured radial temperature; thus we model $\sigma_{x,y}$ to be constant. To account for this behavior, we treat V_{eff} as a function of time in eq. 8.3, with σ_z increasing at the measured rate. We then numerically integrate eq. 8.3 to find values of β and τ that minimize the reduced chi squared (χ^2_{red}) of this model. With $N_0 = 4000$, we find $\beta_{4\text{K}} = 2.7(5) \times 10^{-10}$ cm^3 s^{-1} and $\tau_{4\text{K}} = 1.3(1)$ (with $\chi^2_{\text{red}} = 0.99$, see Figure 8.21), where we incorporate the uncertainty in V_{eff} by adding it in quadrature to the uncertainty determined from the fit.

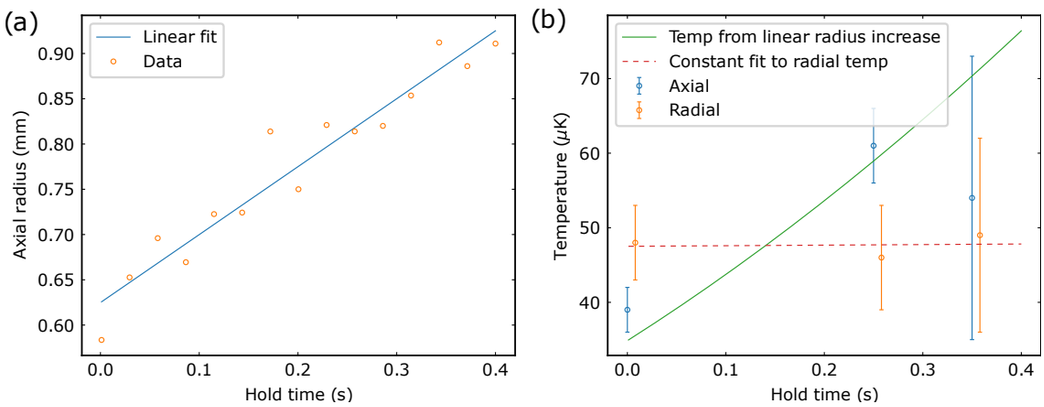

Figure 8.22: (a) Axial radius as a function of hold time, and best fit line with $\sigma_z = mt_h + c$. We find $m = 0.75$ mm/s and $c = 0.63$ mm. (b) Axial and radial temperature versus hold time. Radial temperatures are slightly offset in time for clarity. The axial temperature increases as expected from the axial radius increase, while the radial temperature does not noticeably increase.

The final extracted value of β is strongly dependent on the initial number, so we also consider systematic uncertainties in determining N_0. The scattering rate is affected by uncertainty in the vibrational branching ratio $\left|A^2\Pi_{1/2}, v = 0\right\rangle \rightarrow \left|X^2\Sigma, v = 3\right\rangle$ for SrF [59, 79], and in the calibration of the imaging optics. We estimate a combined uncertainty of 25% in N_0, as mentioned before. We emphasize that this is different from shot-to-shot fluctuations, and instead is a correlated uncertainty for all points, which in turn leads to an uncertainty in the overall normalization of β. To determine the effect of this scale uncertainty, we use the same analysis method with initial numbers $N_0 = \{3000, 5000\}$ (corresponding to the lower and upper bounds given the uncertainty), and numerically integrate eq. 8.3 to find the optimal β for each N_0. The final uncertainty for β is then assigned as the quadrature sum of contributions from this systematic uncertainty and from the fit error for $N_0 = 4000$. Finally, we find $\beta = 2.7^{+1.2}_{-0.8} \times 10^{-10}$ cm^3 s^{-1} and $\tau = 1.3(1)$ s. See § 8.3.4 for the detailed statistical analysis.

As a cross-check, we also fit the data to the analytical solution (eq. 8.4) by following the prescription from Ref. [184]. That is: we first extract τ by fitting a pure exponential decay to only late-time ($t_h \geq 1$ s) data points, and find $\tau_1 = 1.2(2)$ s. Then, we extract β by fixing this value of τ and fitting only to early-time data points ($t_h < 250$ ms) where the axial radius change is less than 15%, such that V_{eff} can be treated as a constant; we use the average V_{eff} for $t_h < 250$ ms, $V_{\text{eff}} = 3.4(9) \times 10^{-7}$ cm^3. Throughout, we perform the same error analysis as before, and find

$\beta_e = 2.7^{+1.4}_{-1.0} \times 10^{-10}$ cm^3 s^{-1} (with $\chi^2_{red} = 1.20$), consistent with results from the more complete model.

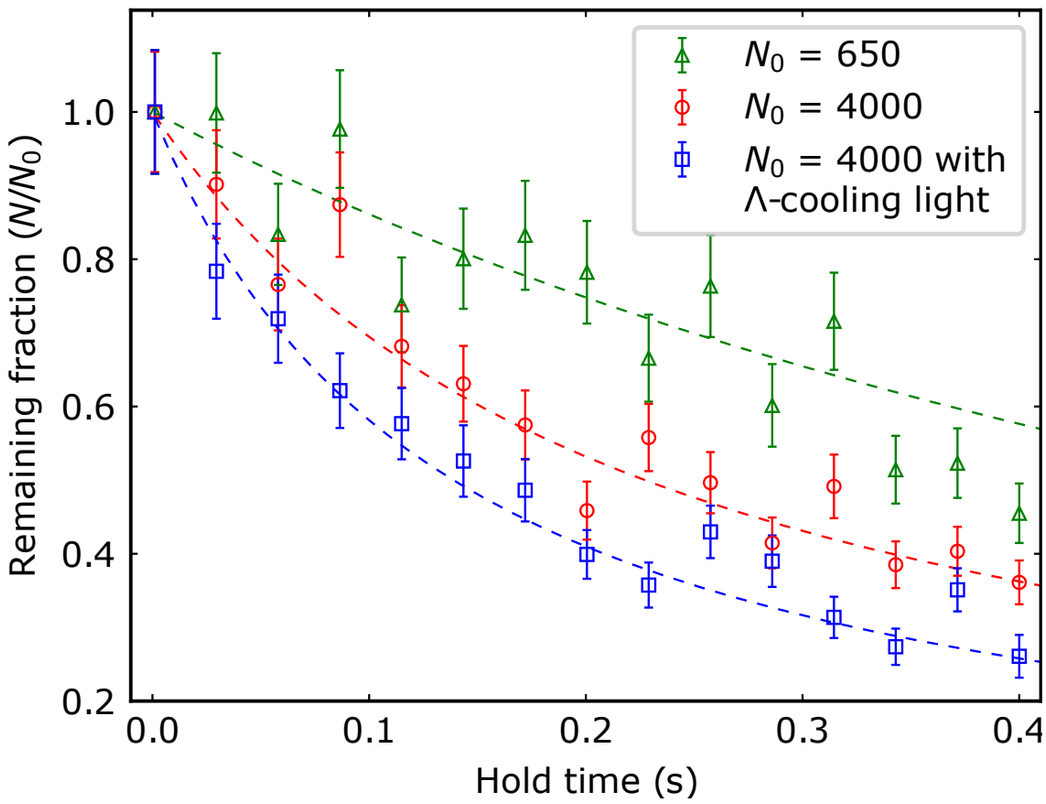

Figure 8.23: Short-time evolution of trap population for different starting conditions. Dashed lines are fits for the first 9 points to the two body loss rate model with fixed $\tau = 1.3$ s and the average V_{eff} for $t_h < 250$ ms. Data with initial ODT number $N_0 \approx 650$ (green triangles) has a slower initial loss as compared to the case with $N_0 \approx 4000$ (red circles), clearly demonstrating the density dependent loss. The presence of Λ-cooling light leads to additional two-body loss (blue squares) due to light-assisted collisions, while the one-body loss rate remains the same (as seen in longer-time data, not shown).

To further verify the presence of a density-dependent loss, we load the ODT with lower initial number, $N_0 \approx 650$, while keeping the temperature and trap depth the same, thereby reducing the starting density by a factor of 6. This is done by reducing the length of the slowing pulse after laser ablation from 14.5 ms to 9.3 ms. We then perform the same sequence of measurements, and see that the short-time loss rate is reduced (see Figure 8.23), as expected since the initial collision-induced loss time scale $\tau_c \propto 1/\beta n_0$. We can analyze this as follows. The contribution of collisions to the loss rate is most pronounced when the density is highest, so we look at the slopes of the curves in Figure 8.23 at $t = 0$. The slope at $t = 0$ in Figure 8.23 for the case of $N_0 = 650$ is 3 times smaller than that for $N_0 = 4000$. However, the effect of the one-body loss is not negligible at $t = 0$, especially for the case of $N_0 = 650$ (low initial density). To determine the pure two body loss time constant in this data, we fit the initial points ($t_h < 250$ ms) to a single exponential decay to find the total loss time constant, then we find the pure two body loss time constant by subtracting the one body loss time constant as follows

$$\frac{1}{\tau_{2b}} = \frac{1}{\tau_{tot}} - \frac{1}{\tau_{1b}}, \tag{8.6}$$

where τ_{tot} is the fitted total loss time constant, $\tau_{1b} = 1.3(1)$ s is the one body loss time constant

measured in the experiment, and τ_{2b} is the extracted two body loss time constant. We find

$$\frac{\tau_{tot}(N_0 = 650)}{\tau_{tot}(N_0 = 4000)} = 2.9 \pm 0.9, \tag{8.7}$$

which yields

$$\frac{\tau_{2b}(N_0 = 650)}{\tau_{2b}(N_0 = 4000)} = 6.0 \pm 2.9. \tag{8.8}$$

This is consistent with the expected value $4000/650 = 6.2$.

There are numerous possible loss channels in our experiment. The molecules are all in the rotational $N = 1$ states, and rotational quenching to $N = 0$ can lead to large inelastic losses [73]. They also occupy all of the many individual sublevels in the $N = 1$ manifold of hyperfine and spin-rotation states; this opens up p-wave and f-wave collision channels that would be absent if all the (bosonic) molecules were in the same quantum state. In addition, colliding pairs of SrF molecules can undergo the barrierless chemical reaction SrF + SrF $\rightarrow$ SrF_2 + Sr [189]. Finally, "sticky collisions" between the molecules can also lead to losses [72, 74, 143, 174–178].

We next compare our measured value of β to some experimental and theoretical benchmarks. The molecules in the trap are at temperatures above the p- and d-wave barriers (≈ 5 μK and ≈ 30 μK respectively) determined by the van der Waals C_6 coefficient for $N = 1$ states in SrF. The universal loss rate model [142], which assumes that colliding molecules are lost if they reach short range, i.e., if they do not reflect off the van der Waals (vdW) + centrifugal potential, has proven consistent with observed experimental loss rates [72, 74, 143, 174–178]. We use the generic solutions from [144], which are valid for systems where the temperature is above the p- and d-wave barriers. We find a thermally- and ensemble-averaged loss rate constant $\beta_{\text{univ}} = 2.6 \times 10^{-10}$ cm^3 s^{-1}. We also calculate the maximum allowed loss rate constant by summing the maximum inelastic cross-sections for each partial wave [190] and find $\beta_{\text{max}} = 2.8 \times 10^{-10}$ cm^3 s^{-1} [191]. The close match indicates small reflection probabilities, as expected for $T = 40\,\mu$K. Our experimental measurement of β is consistent with both calculations.

The closest case to our conditions was in the observation of collisions between a pair of CaF molecules in a mixture of $N = 1$ states in an optical tweezer trap [13]. There, the molecules were at $T \approx 80\,\mu$K, above the p-wave, but below the d-wave barrier for CaF ($\approx 20\,\mu$K and $\approx 105\,\mu$K respectively). The reported loss rate coefficient is $\beta_{\text{CaF}} = 40 \times 10^{-10}$ cm^3 s^{-1}, $\approx$ 10x larger than the predicted universal value $\beta_{\text{univ,CaF}} \approx 3.25 \times 10^{-10}$ cm^3 s^{-1}, in contrast to our results, which match the model.

We also explore light-assisted collisions due to the Λ-cooling light. For this, we turn on the Λ-cooling light at $t_h = 0$. We observe that the one-body lifetime is unaffected by the light, however, the two-body loss rate coefficient is increased due to light-assisted collisions. Under these conditions, we find $\beta_{\text{tot}} = 4.9^{+1.7}_{-1.2} \times 10^{-10}$ cm^3 s^{-1}. This is two orders of magnitude lower than previously reported using a pair of CaF molecules in an optical tweezer [13]. The combined loss rate coefficient sets an upper bound on the peak density achievable by loading an ODT using Λ-cooling. Given the typical loading time (20 ms) from the blue-MOT, this bound is $n_0^{\text{max}} \sim 10^{11}$ cm^{-3}. While the peak densities we achieve are lower than n_0^{max}, it may be possible to reach it if larger numbers of molecules [110], lower temperatures [68, 111], and/or deeper traps can be achieved. We also note that the observed value of β_{tot} is low enough such that light-induced losses during the in-situ

imaging with Λ-cooling light do not substantially affect the extracted values of $N(t_h)$.

8.3.4 Statistical analysis

Here, we will describe the error analysis for the two-body loss rate coefficient for the no light case with full loading (i.e. $N_0 = 4000$), but the analysis is general and same for the other cases described before. In general, β here has 4 sources of uncertainty, for the way in which we took data (basically, taking ratios of $N(t)/N(0)$ and multiplying by our measurement of $N(0)$:

- Uncorrelated uncertainties for each point: These are the errors in determining the total camera counts at each point, and stem from shot-to-shot fluctuations, gaussian fit uncertainties, as well as the calibration from MOT recapture to Λ-imaging for the later points.
- Systematic variation in $N(0)$: This comes from two sources, uncertainty in how we convert imaging counts into number (primarily from uncertainty in scattering rate measurement due to uncertainty in b_{03} branching ratio, as well as uncertainty in our exponential decay fits to loss rates taken with and without $v = 3$ repump light), and uncertainty in those counts over time (here parameterized by difference between image at beginning of scan and at end of scan, as well as variation in the $N(0)$ images with time). This would be a correlated uncertainty for all points.
- Uncertainty in V_{eff} . This primarily comes from uncertainty in our temperature and trap depth, which inform our estimate of the radial widths $\sigma_{x,y}$. This can be added in quadrature at the end, as this does not affect any of the data analysis. This is around 20% for our temperature and trap depth measurements.
- Uncertainty in the fit to β. Though our fit routines do return uncertainties, it is possible they are underestimates; it will be helpful here to look at χ^2 plots.

Here we will look at each of these in turn.

8.3.4.1 Errors in number determination

First let us look at errors that can affect the fits, i.e. errors in the number determination. These are as follows:

1. Uncorrelated uncertainties: These stem from the errors in the gaussian fits, and shot-to-shot uncertainties for the early time points, and are shown in Figure 8.21. For the MOT recapture points, we took additional data at a few intermediate points with both methods to compare the derived number. We find a ratio $N_{\text{MOT recap}}/N_{\Lambda\text{-image}} = 1.01 \pm 0.10$, and so an additional 10% uncertainty is added to the MOT recapture points.
2. Scattering rate error: As seen in § 7.2.4, there is a roughly 17% uncertainty in determining the molecule number from the scattering rate. This is a correlated uncertainty.
3. Deviations in N_0: By virtue of the way we take the data, we can track the variations in N_0 throughout the data taking run, and we see that these have a standard deviation $\sim 8\%$. This is a correlated uncertainty.
4. Camera calibration errors: We estimate a 10% uncertainty in the camera calibration (i.e. conversion between the camera counts to photon number). This is a correlated uncertainty.

The combined correlated error for each point is thus around $\sim 25\%$. We need to account for this, as we shall see later in this section.

8.3.4.2 Reduced chi-squared analysis

Before we account for correlated errors, we first need to make sure our fits with the uncorrelated errors are good, and extract a meaningful uncertainty from these fits. That is, we would like to know if the fits to the data are good or bad, and which model best captures what is happening. A useful metric to do this statistically is the reduced chi-squared of the model, which is defined as the chi-squared per degree of freedom:

$$\chi^2_{\text{red}} = \frac{\chi^2}{\nu} \quad ; \quad \chi^2 = \sum \frac{(O_i - M_i)^2}{\sigma_i^2} \tag{8.9}$$

where $O_i(M_i)$ are the observation (model) points, and σ_i is the error of each observation. ν, the degrees of freedom, is defined as $\nu = m - n$ where m are the number of data points being fit to, and n are the number of fit parameters. A fit is said to be good when the reduced chi-squared $\chi^2_{\text{red}} \sim 1$. If $\chi^2_{\text{red}} \gg 1$, the error bars are too small, and if $\chi^2_{\text{red}} \ll 1$, then the error bars are too large.

For this analysis, we will fix $N_0 = 4000$, with the uncertainty for each point being the uncorrelated uncertainty. In our case, we can fit to two models, the analytical solution as in eq. 8.4 or the ODE as in eq. 8.3. Given that we see an expanding width σ_z, we will first try to fit to eq. 8.3, with β and τ as the free variables. We numerically integrate eq. 8.3 for a given set (β, τ) and compute the χ^2_{red}. By iterating over different sets of (β, τ), we can find the optimal (β, τ) values that minimize χ^2_{red}. Then, we can also find the errors in (β, τ) by finding values of (β, τ) that satisfy the following:

$$\chi^2_{\text{red,err}} = \frac{\chi^2 + 1}{\nu} \tag{8.10}$$

The result of fitting to all the points is shown in Figure 8.24. We see a minimum $\chi^2_{\text{red}} = 0.99$ for $\beta \approx 2.7 \times 10^{-10}$ and $\tau \approx 1.3$ s, with a contour of low χ^2_{red} around this central value. In this case, $\nu = 21 - 2 = 19$, and the contour shows the points for which $\chi^2_{\text{red,err}} = \chi^2_{\text{red}} + 1/19$. Since χ^2_{red} is so close to 1, it means we have correctly assigned our uncorrelated uncertainties, and the model fits the data well. We can also extract the errors to the fit from this contour directly, to find $\beta = 2.7(2) \times 10^{-10}$ cm^3 s^{-1} and $\tau = 1.30(05)$ s.

As mentioned before, we can also try to fit to the analytical solution eq. 8.4, by following the procedure from [184]. To do this, we first fit a simple exponential decay to the later time points, to get $\tau = 1.2$ s as the one-body lifetime. Then, we can fix $N_0 = 4000$ and $\tau = 1.2$ s in eq. 8.4 and again iterate over β to find the optimum, as shown in Figure 8.25. In this case, we only have 1 fit parameter, β, and we are only fitting to the initial points $t_{\text{hold}} < 250$ ms, i.e $\nu = 9 - 1 = 8$. This gives us $\beta = 2.7(0.6) \times 10^{-10}$ cm^3 s^{-1}.

8.3.4.3 Final error budget

The χ^2_{red} analysis does not account for the uncertainty in V_{eff} or the systematic variation in N_0. We know that the combined correlated uncertainty is 25%, and so we can account for that as follows. We can repeat the same χ^2_{red} analysis from the previous section and find the values of β that are optimal for $N_0 = 3000, 5000$, corresponding to the 25% range around $N_0 = 4000$. These will give

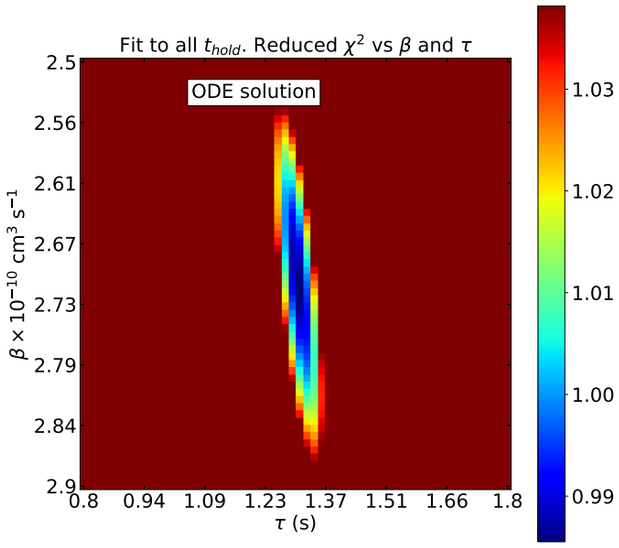

Figure 8.24: χ^2_{red} for the ODE model for $N_0 = 4000$. A contour of low $\chi^2_{\text{red}} \sim 1$ is seen, around $\beta \approx 2.7 \times 10^{-10}$ cm^3 s^{-1} and $\tau \approx 1.3$ s.

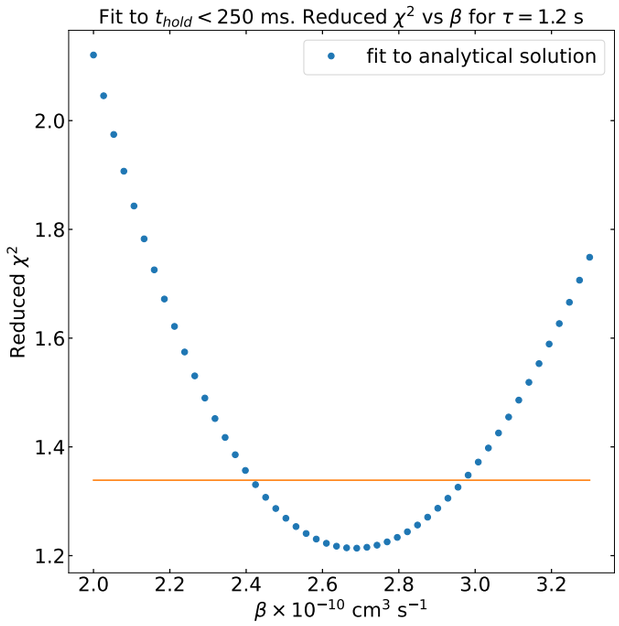

Figure 8.25: χ^2_{red} for the ODE model for $N_0 = 4000$ and $\tau = 1.2$ s. We fit only to the first 9 points to determine β, which is why χ^2_{red} is a bit higher. The orange line shows the points for which $\chi^2_{\text{red,err}} = \chi^2_{\text{red}} + 1/\nu$.

the upper and lower bound on β, which we can then add to the actual uncertainty in the central value of β found in the previous section. Note that here, τ is unaffected since we are scaling the entire data by the same factor. Because the two-body loss is strongly dependent on initial number, this analysis leads to asymmetric bounds. Our data is a normalized set of molecule numbers, and thus the two-body loss is more pronounced for $N_0 = 3000$ than it is for $N_0 = 5000$.

To account for the uncertainty in V_{eff}, we add the 20% uncertainty ($\sigma_x = \sigma_y = 3.3(4)\,\mu$m) we determined from our temperature and trap depth measurement uncertainties to the error bars from the previous analyses combined. This nets us the result described before, i.e. $\beta = 2.7^{+1.2}_{-0.8} \times 10^{-10}$ cm^3 s^{-1} and $\tau = 1.3(1)$ s.

8.3.5 Calculation of van der Waals C_6 coefficient

To determine the centrifugal barriers for two-body SrF scattering, we need to compute the van der Waals (vdW) C_6 coefficient that arises from second-order dipolar coupling. Throughout, we

assume the molecules are in the ground vibronic manifold $\left|X^2\Sigma^+, v=0\right\rangle$, and for simplicity we ignore electronic and nuclear spins.

For rotational ground state $|N=0\rangle$ molecules, the rotational wavefunction is spatially isotropic and only a single rotational sublevel is occupied. This leads to the well-known result [192] $C_6^g = -\left[1/(4\pi\epsilon_0)^2\right] \cdot d_0^4/(6B_0)$, where $d_0 = 3.47$ Debye is the ground state permanent dipole moment of SrF and $B_0 = 2\pi\hbar \times 7.5$ GHz is the ground state rotational constant of SrF. However, our molecules are in an incoherent mixture of $|N=1\rangle$ states, which (in the absence of external fields, and ignoring effects due to spin-rotation and hyperfine couplings) comprise a nine-fold degenerate subspace in the space of $|N=1\rangle$ two-body states. Following the approach of [191], we apply second-order degenerate perturbation theory on the intermolecular potential operator $\hat{V}_{\mathrm{AB}}$ in order to obtain the $|N=1\rangle$ C_6 coefficients.

In a body-fixed (BF) frame where the orientation of the vector between the two molecules is fixed, the resultant C_6 coefficients, as a function of d_0 and B_0, are listed in Table 6 of [191]. We used these values in eq. 8.17 to compute the centrifugal barrier heights. Due to the anisotropic nature of the vdW interaction, the C_6 coefficients including the relative angular motion of the molecules in the space-fixed (SF) frame must be accounted for. To find C_6 values in the SF frame, we numerically compute matrix elements of the second-order degenerate perturbation operator $\hat{W}_{\mathrm{AB}}^{\mathrm{SF}}$ associated with $\hat{V}_{\mathrm{AB}}$, as given by eq. 82 of [191]. By incorporating the ℓ^{th} partial wave $|\ell, m_\ell\rangle$ into that matrix element, and subsequently diagonalizing the combined potential $\hat{W}_{\mathrm{AB}}^{\mathrm{SF}} + \hat{\ell}^2/(2\mu R^2)$ (where R is the intermolecular separation) in the subspace of $|N=1\rangle$ two-body states, we obtain the intermolecular potential curves shown in Figure 8.26. The resultant barrier heights that we obtain from these curves are in nearly exact agreement with those obtained from using the BF calculation results. Hence, for computational ease we use the analytically determined BF centrifugal barriers in the rest of this work.

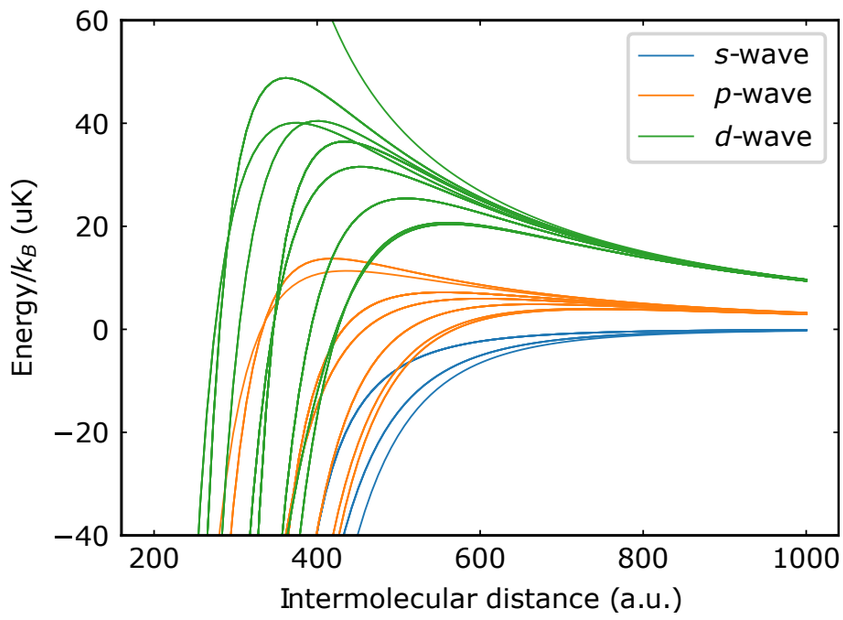

Figure 8.26: Numerically computed potential energy curves in a single channel model for two-body SrF scattering, where colliding molecules are in the $|N=1\rangle$ manifold. Curves up to the $\ell = 2$ partial wave are shown.

8.3.6 Unitarity Limit Calculation

Here, we compute the unitarity limit for two-body scattering of SrF molecules in an incoherent mixture of $|N=1\rangle$ states. The unitarity limit corresponds to the maximum possible loss allowed by quantum scattering theory. Since our molecules are not in a single quantum state, all partial waves (indexed by ℓ) contribute to the unitarity limit scattering cross section σ. In this limit, the scattering cross section of the ℓ^{th} partial wave is:

$$\sigma_\ell = \frac{\pi(2\ell+1)}{k^2}. \tag{8.11}$$

As usual in two-body scattering, we work in the center-of-mass frame of the two-body system. Therefore, the wavevector $\vec{k}$ relates to the collisional energy E_r and relative velocity $\vec{v}_r$ of the two particles via $E_r = \hbar^2 k^2/2\mu$ and $\vec{v}_r = \hbar\vec{k}/\mu$, respectively, where μ is the reduced mass. For two-body SrF scattering, $\mu = M/2$ where $M \approx 107$ amu is the mass of SrF.

We assume that SrF molecules in our trap obey a Maxwell-Boltzmann (MB) thermal distribution in their absolute velocities and energies. Since the convolution of two Gaussians is another Gaussian, it follows that the probability density function $f_r(v_r, T)$ of relative speeds for two-body SrF scattering is also a normalized MB distribution:

$$f_r(v_r, T)\, dv_r = \sqrt{\frac{2}{\pi}} \left(\frac{\mu}{k_B T}\right)^{\frac{3}{2}} v_r^2 \exp\left\{\left(-\frac{\mu v_r^2}{2k_B T}\right)\right\} dv_r, \tag{8.12}$$

where $\int_0^\infty f_r(v_r, T)\, dv_r = 1$ and T is the initial temperature of molecules in our trap.

We define a thermally averaged two-body loss rate coefficient $\beta_{\text{th}}(\ell, T)$ for the ℓ^{th} partial wave by considering which collisional speeds will allow for unitary loss. We make the simplifying assumption that unitary loss occurs with unit (zero) probability when E_r is greater (less) than the centrifugal barrier of the intermolecular potential, i.e. we assume that colliding SrF molecules cannot tunnel through the centrifugal barrier. By denoting the collisional speed associated to the barrier height as v_b, we have:

$$\beta_{\text{th}}(\ell, T) = \int_{v_b}^{\infty} f_r(v_r, T)\sigma_\ell(v_r) v_r\, dv_r. \tag{8.13}$$

By introducing the dimensionless parameter $x \equiv E_r/k_B T$ and substituting in the expressions for $f_r(v_r, T)$ and $\sigma_\ell(v_r)$, the expression for $\beta_{\text{th}}(\ell, T)$ can be simplified as:

$$\beta_{\text{th}}(\ell, T) = \frac{\pi\hbar^2(2\ell+1)}{2\mu k_B T}\sqrt{\frac{8k_B T}{\pi\mu}} \int_{T_b(\ell)/T}^{\infty} e^{-x}\, dx. \tag{8.14}$$

We identify $\lambda_{\text{th}} = \sqrt{2\pi\hbar^2/\mu k_B T}$ and $\bar{v}_{\text{th}} = \sqrt{8k_B T/\pi\mu}$ as the thermal de Broglie wavelength and average speed, respectively, of a particle with mass μ in an ensemble at temperature T that obeys MB statistics. Therefore, we conclude:

$$\beta_{\text{th}}(\ell, T) = \lambda_{\text{th}}^2 \bar{v}_{\text{th}}(2\ell+1)e^{-T_b(\ell)/T}/4, \tag{8.15}$$

where $k_B T_b(\ell)$ is the height of the centrifugal barrier experienced by the ℓ^{th} partial wave.

Each distinct vdW C_6 coefficient corresponds to a distinct value for $T_b(\ell)$. For the i^{th} C_6 coefficient, denoted $C_{6,i}$, we relate it to the i^{th} barrier $T_{b,i}(\ell)$ as follows. If we neglect quadrupole-quadrupole interactions, the leading order terms in the SrF $|N=1\rangle$ + SrF $|N=1\rangle$ intermolecular

potential lead to the potential:

$$V_i(r) = \frac{\hbar^2 \ell(\ell+1)}{2\mu r^2} + \frac{C_{6,i}}{r^6}. \tag{8.16}$$

If $C_{6,i} > 0$, the vdW interaction is repulsive and no barrier exists. We treat this as meaning the molecules never reach short range, and so the contribution to $\beta_{\rm th}(\ell, T)$ here is zero. If $C_{6,i} < 0$, the vdW interaction is attractive and there will exist a maximum in $V_i(r)$ at $r_{\rm b}$, corresponding to the centrifugal barrier. By only considering cases where $C_6 < 0$, we analytically compute the barrier height of the $\ell^{\rm th}$ partial wave to be:

$$T_{b,i}(\ell) = \frac{V(r_{\rm b})}{k_B} = \left(\frac{\hbar^2 \ell(\ell+1)}{\mu}\right)^{\frac{3}{2}} \left(\frac{1}{54|C_{6,i}|}\right)^{\frac{1}{2}} \frac{1}{k_B}. \tag{8.17}$$

For each two-body eigenstate, we compute their barrier heights up to the h-wave contribution ($\ell = 5$). We neglect summation over all partial waves with $\ell > 5$ because their contribution to $\beta_{\rm th}(\ell, T)$ is increasingly exponentially suppressed. We thus obtain the total $\beta_{\rm th,i}(T)$ for the $i^{\rm th}$ two-body eigenstate:

$$\beta_{\rm th,i}(T) = \begin{cases} \sum_{\ell=0}^{5} \lambda_{\rm th}^2 \bar{v}_{\rm th} (2\ell+1) e^{-T_{b,i}(\ell)/T}/4 & \text{if } C_6 < 0 \\ 0 & \text{if } C_6 > 0. \end{cases} \tag{8.18}$$

We finally obtain the overall two-body loss rate coefficient in the unitarity limit for an incoherent mixture of $|N=1\rangle$ states, denoted as $\langle \beta_{\rm th}(T) \rangle_{|N=1\rangle}$, by taking a statistical average over all nine possible $|N=1\rangle$ two-body eigenstates. Here, we assume a uniform probability distribution over all possible states, i.e. $p_i = 1/9 \; \forall i \in \{1, 2, ..., 8, 9\}$. So, we have:

$$\langle \beta_{\rm th}(T) \rangle_{|N=1\rangle} = \sum_{i=1}^{9} p_i \, \beta_{\rm th,i}(T). \tag{8.19}$$

Carrying out this computation, we find $\langle \beta_{\rm th}(T) \rangle_{|N=1\rangle} \approx 2.8 \times 10^{-10}$ cm^3/s at $T = 40\,\mu$K, the temperature of SrF molecules in our ODT.

We note that the intermolecular potential for $|N=1\rangle$ states contains a first-order contribution from the quadrupole-quadrupole interaction, of the form C_5/r^5 [191]. Here, $C_5 \sim \theta_{zz}^2/(4\pi\varepsilon_0)$, where $\theta_{zz} = 8.95\, e a_0^2$ is the quadrupole moment of SrF [189]. This extra term affects the barrier heights only minimally; by including the quadrupole-quadrupole interaction, we found that $\langle \beta_{\rm th}(T) \rangle_{|N=1\rangle}$ is changed by only $\approx 3\%$, and hence its effect is negligible.

8.4 Conclusions

In this chapter, we demonstrated that an ODT can be successfully loaded from a blue MOT. We achieve record high transfer transfer efficiencies ($\sim 30\%$), allowing us to trap around $N \sim 4000$ molecules in the ODT. This enabled us to reach the high densities needed to observe molecule-molecule collisions in the ODT. This was the first such observation of collisions in a bulk gas of a directly laser-cooled molecule. From this data, we were able to extract a two-body loss rate coefficient around 3 times smaller than the corresponding unitarity limit. This suggests that the loss rate could be further lowered if the molecules were pumped in the rovibrational ground state, to get rid of the additional loss channels in the current experiment. Finally, this opens the path

to evaporative cooling in a bulk gas towards degeneracy.

Chapter 9

Present and future directions

9.1 Introduction

In Chapter 8, we demonstrated high density loading of an ODT using a blue MOT and we measured the inelastic loss rate coefficient for SrF molecules in the $N = 1$ manifold. This was the first observation of collisions in a bulk gas of a directly cooled molecules. However, several questions remain to be addressed. The molecules are all in an incoherent mixture of internal levels, and thus quantum control over the molecular ensemble remains to be achieved. Furthermore, by preparing the molecules in the absolute rovibrational ground state, a lot of the collisional mechanisms can be distilled out, making the physics simpler to understand. Ultimately, the goal of the experiment is to make a BEC, and so techniques to shield these inelastic collisions need to be developed and implemented.

In this chapter, we describe the ongoing work to achieve coherent quantum control over the SrF molecules in the ODT. This includes optical pumping to a single quantum state in the $N = 1$ manifold, and then microwave transfer to the rovibrational ground state. We also describe work towards microwave shielding of SrF molecules [80, 183]. We describe in brief a next generation apparatus, and the improvements expected from improved slowing and trapping techniques. Finally, we outline future directions that may be pursued, including co-trapping Rb for sympathetic cooling and direct evaporative cooling.

9.2 Coherent quantum control

Here we describe our current efforts towards achieving coherent quantum control of the SrF molecules in the ODT.

9.2.1 Optical pumping into $N = 1, F = 0$

The first step is to prepare the molecules in a single quantum state. There are different states we can pump the molecules into, to probe different things. For example, if we pumped all the molecules into $|N = 1, F = 0\rangle$, we can get rid of p-wave and f-wave collisions, and by measuring the loss rate here, we can probe the combined effect of the other loss channels. Further, by pumping molecules into one of the Zeeman sublevels of $|N = 0, F = 1\rangle$, we can eliminate rotational relaxation and odd-l collisions, and thus probe the effect of hyperfine relaxation and sticky collisions. Finally, by pumping the molecules in $|N = 0, F = 0\rangle$, we can be rid of hyperfine relaxation as well, and directly

probe sticky collisions. Thus, the experimental path we will follow is similar to the above: first optimize the optical pumping into $|N = 1, F = 0\rangle$, then optimize the transfer of these molecules to $|N = 0, F = 0\rangle$ and finally measure the lifetime in the ODT here.

9.2.1.1 Simulations

Optical pumping (OP) into $F = 0$ is likely to be most efficiently done with low power resonant lasers. Our starting guess for the easiest way to accomplish this was to use V_{00}^m to address the $|F = 2, 1 \uparrow\rangle$ levels, while V_{00}^d can be resonant with $|F = 1 \downarrow\rangle$ with no sidebands so that no laser light addresses the $|F = 0\rangle$ state. Due to the current experimental setup where V_{00}^M and V_{00}^d are combined on a PBS before the MOT AOM, we are constrained to use 2x power in V_{00}^M than in V_{00}^d. T. Langin performed OBE simulations to determine optimal parameters for OP into $F = 0$ given these conditions, with the results shown in Figure 9.1.

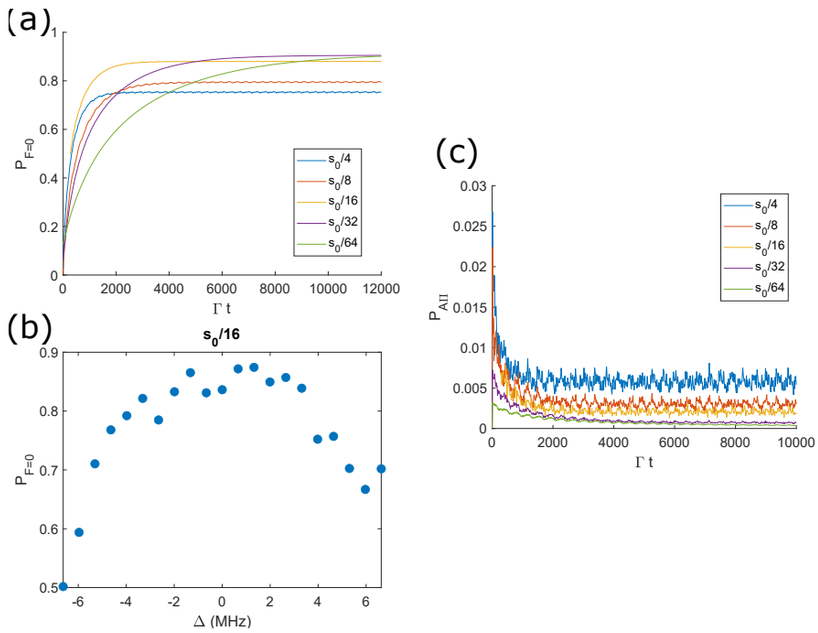

Figure 9.1: Simulations of optical pumping efficiency into $F = 0$. (a) $F = 0$ population vs time for different V_{00}^d intensities. $\approx 90\%$ pumping is achieved for $s_0/32$ (roughly 80 μW of V_{00}^d power). Higher powers lead to depletion from $F = 0$ due to the $|F = 1 \downarrow\rangle$ state being only 7.5Γ away. (b) OP vs Δ. This wants to be on resonance, as expected. (c) Excited state population vs time. We can determine how many photons are expected to be scattered during the pumping process by integrating this curve over time. Figure obtained from T. Langin.

We find that having the $|F = 1 \uparrow\rangle$ intensity any higher than $s_0/32$ reduces the OP efficiency (Figure 9.1(a)). The pumping is very fast, with a timescale roughly $8000\Gamma t = 200\,\mu$s, and around 90% of the molecules are pumped into $|F = 0\rangle$. The pumping efficiency falls off with detuning, as expected (Figure 9.1(b)). We also find that around 10 photons are absorbed in this process (Figure 9.1(c) - these curves are integrated to get total value of $N_{scat} = 2\int_0^{t_{end}} P_e \Gamma dt \sim 20$, where t_{end} is set by $P_{F=0} = 99\%$ of its maximum value). Each photon absorption and emission gives 2 velocity kicks of $\hbar k/m$; so along each axis the kick is rouhgly $\hbar k/\sqrt{3}m$. This gives a total random walk velocity along each dimension $v_{rand} = (\sqrt{N_{kick}/3})\hbar k/m$, and in temperature units, this is $T_{rand} = 2.7\,\mu$K. Thus we would expect slight heating of the cloud - however, given our 40 μK ODT temperatures, this should not be a problem.

9.2.1.2 Experimental protocol

Figure 9.2 shows the OP and detection scheme we use in the experiment. Lasers L1, L2 and L3 serve to cycle photons and pump into $F = 0$. We can use two independent methods to detect the efficiency of this OP scheme.

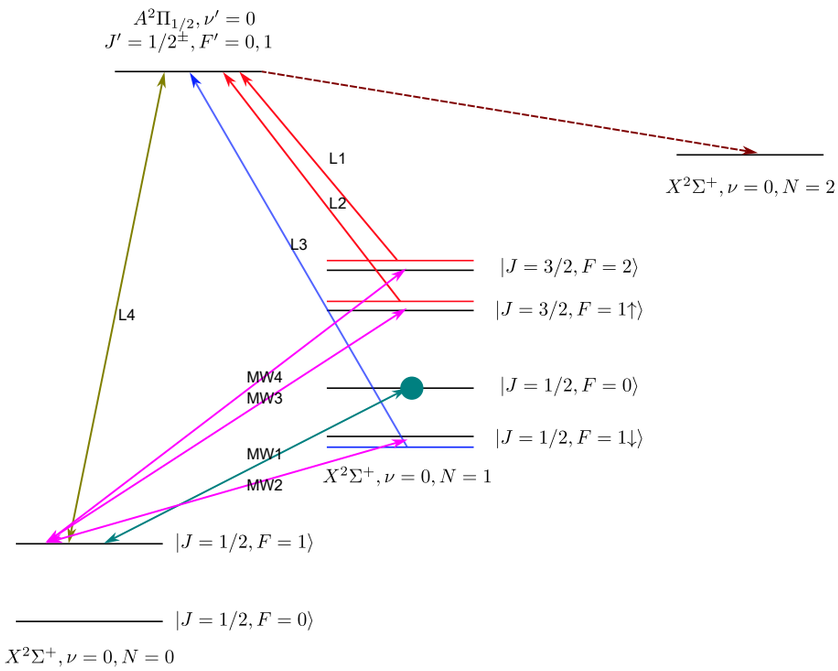

Figure 9.2: Level diagram showing the optical pumping and readout scheme. Lasers L1, L2 and L3 are used to transfer molecules into $|N = 1, F = 0\rangle$. To measure the efficiency of this transfer, either microwaves MW1, or microwaves MW2, MW3 and MW4 are used to transfer molecules to $|N = 0, F = 1\rangle$. The pumpout laser L4 transfers these molecules into either the $N = 2$ manifold, or into $|N = 0, F = 0, 1\rangle$, both of which appear as loss.

- First, we can use microwave MW1 to continuously transfer molecules from $|N = 1, F = 0\rangle \leftrightarrow |N = 0, F = 1\rangle$. By shining a pumpout laser (L4) on the $|N = 0, F = 1\rangle \rightarrow \left|A^2\Pi, J = 1/2^-\right\rangle$ transition, these molecules are pumped to either the $N = 2$ manifold, or go back to $N = 0$, both of which appear as loss. If we do this for long enough, we can deplete all the molecules in $|N = 1, F = 0\rangle$. By measuring the remaining population in $N = 1$, we get a measurement of the number remaining in the other hyperfine levels.

- We can also use microwaves MW2, MW3 and MW4 to continuously transfer molecules from $|N = 1, F = 2, 1\uparrow, 1\downarrow\rangle \leftrightarrow |N = 0, F = 1\rangle$. By similarly applying the pumpout laser (L4), we can get rid of these molecules. Then, by measuring the remaining molecules in $N = 1$, we get a measurement of how much we pumped in $F = 0$.

For the purposes of optimizing the OP efficiency, we use the first method, i.e. deplete molecules out of $F = 0$ and measure the remaining number. In this case, we want to minimize the signal, with the idea that lower the number remaining, higher the OP efficiency, and ideally, we can completely deplete $F = 0$, meaning 100% OP.

The experiment protocol we use is the following:

1. Trap molecules in the MOT.

2. Apply a Λ-cooling pulse for 5 ms to cool the molecules.

3. Wait for V_{00}^M shutter to open[1].

[1]This takes around 8 ms. This is a normally closed shutter and we have found empirically that these are slow to perform a open-close-open cycle.

4. Switch laser powers and frequencies and apply OP pulse.

5. Turn off all OP light and pump out of $F = 1$ using MWs and L4. Here, to turn off the OP light, we use the RF switch on the AOM to turn off the RF.

6. Measure remaining population.

There are a number of parameters that can affect the OP efficiency. The detunings of L1, L2, and L3, the duration of the OP pulse, the background magnetic field and the overall laser power (here we are constrained to a 1:2 power ratio between L3 ($V_{00}^{\dagger}$) and L1+L2 (V_{00}^{M}); however the ratio between L1 and L2 can be changed by tuning the EOM) can all matter. In the following, we will show the scans for each of these parameters, and finally show the best OP efficiency achieved so far.

To start off the scans, we set the OP parameters to be the optimal parameters from the simulation: V_{00}^{M} single pass laser power around 100 μW, L1, L2 and L3, all close to resonance, and the OP light on for 400 μs. We also used a 5 ms microwave + pumpout laser pulse, during which time the OP light was nominally off. We shall see the effect of this "nominal" label later. Also note that in our setup, L1 and L2 are derived from V_{00}^{M}, where L1 is the carrier frequency and L2 is the sideband. L3 is derived from $V_{00}^{\dagger}$. Both of these are then combined together on an AOM, which gives us the overall power and timing control. We also note that all the scans were randomized to avoid any drifts, and we averaged 6 shots per point.

The microwave setup is shown in Figure 9.3. The microwaves are introduced along the slowing path at a slight angle to the slowing beam. The nominal polarization is assumed to be along the horizontal direction, however, as seen in § 9.2.2.1, all three polarizations are present in the chamber, presumably due to microwave reflections off the chamber walls.

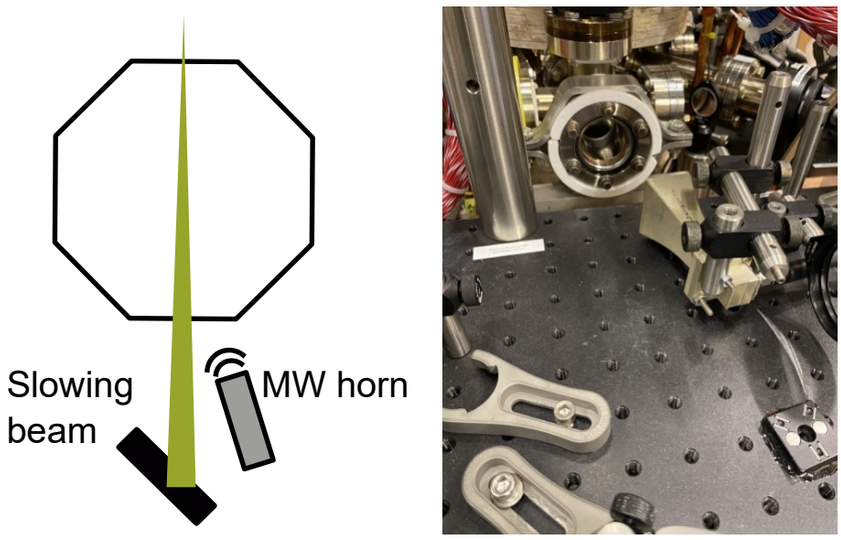

Figure 9.3: Microwave horn setup for optical pumping. The microwaves are introduced from along the slowing path, tilted at a slight angle with respect to the slowing propagation. Left shows a cartoon of the setup, while right shows the actual implementation in the experiment.

9.2.1.3 Frequency scans

Figure 9.4 shows the results of the frequency scans for each L1, L2 and L3. We find that L3, i.e. $V_{00}^{\dagger}$ wants to be a few MHz blue of the resonance, and in fact is quite close to the detuning used for the blue-MOT. It turns out to be around 6-7 MHz blue of resonance. Similarly, L1 i.e. V_{00}^{M} wants to be very close to the detuning used for the red-MOT, and is around 6-7 MHz red of resonance.

The sideband on V_{00}^{M}, i.e. L2 is not super sensitive to frequency, but wants to be somewhat low amplitude, with the best efficiency at a carrier to sideband ratio of 2:1 or so.

This makes sense: the lasers all want to be as far away from the $F = 0$ state as possible to avoid any off-resonant excitation, while still being able to scatter photons from their respective hyperfine states. These optimal detunings are indicated in Figure Figure 9.2.

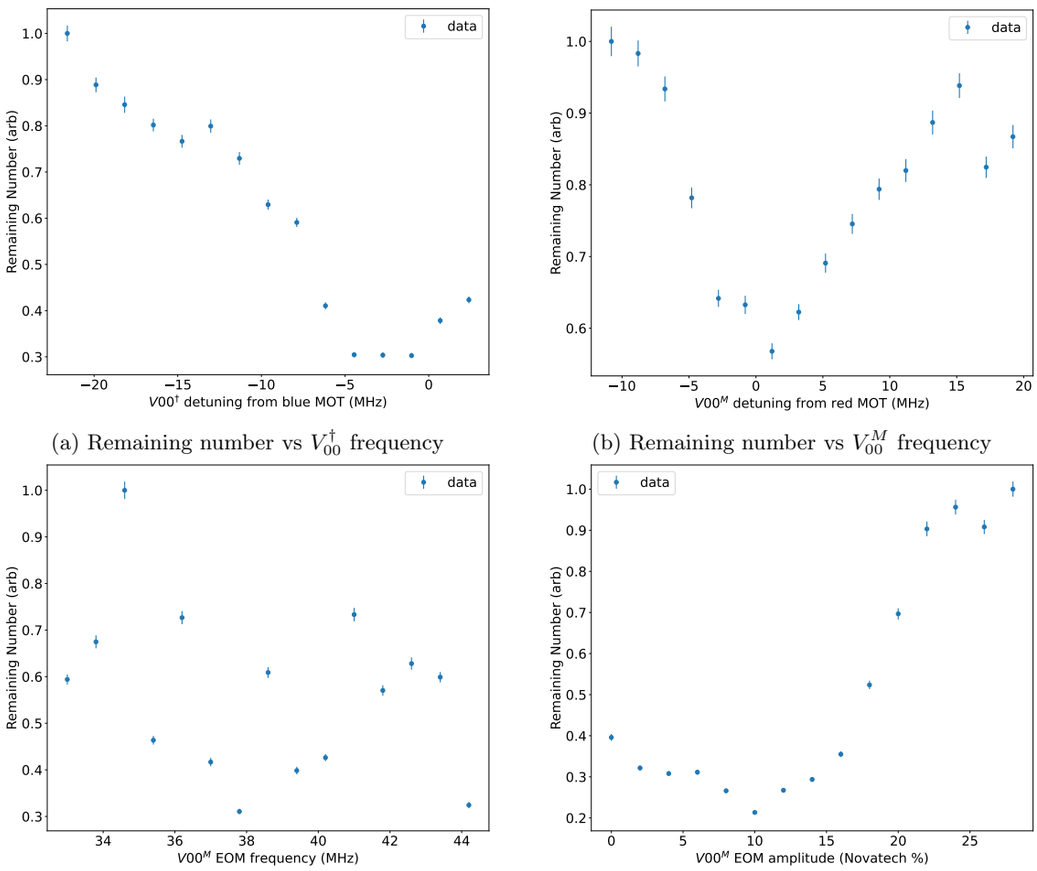

(a) Remaining number vs $V_{00}^{\dagger}$ frequency

(b) Remaining number vs V_{00}^{M} frequency

(c) Remaining number vs V_{00}^{M} EOM frequency

(d) Remaining number vs V_{00}^{M} EOM amplitude

Figure 9.4: Optimal OP parameters for the lasers. All the plots are normalized to their highest value, and these are not an indication of OP efficiency.

9.2.1.4 OP power and timing scans

Next, we look at the OP power and timings. Figure 9.5 shows these scans. We see that the OP efficiency saturates at around a 800 μs pulse timing. We also see that the OP efficiency decreases with increasing overall laser power (L1+L2+L3), as higher power induces more off-resonant scatters from $F = 0$. We also see from this plot that we might benefit from going to even lower powers. Currently, we are running into the limitation that the main program is running into the minimum resolution, and we need to make a new lookup table at these incredibly low powers in order to get reliable performance. Note that we are using 0.01% of the maximum power in these lasers, and we want to go even lower.

9.2.1.5 Shim scans

Next we look at the effect of having a magnetic field during OP. Figure 9.6 shows the scans for all 3 shims coils currently being used, along with the nominal values used for Λ-cooling. We see that the optimal shim values are different from the Λ-cooling values, but also are not very sensitive around this value.

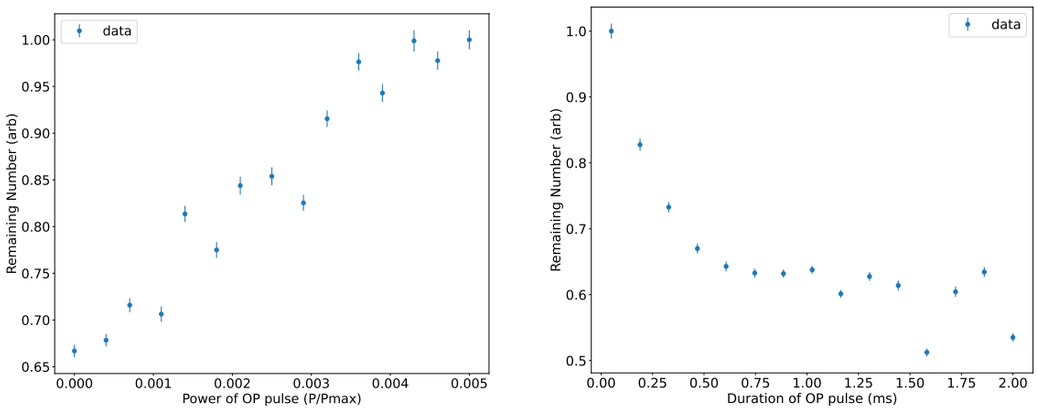

(a) Remaining number vs overall laser power

(b) Remaining number vs OP duration

Figure 9.5: OP power and timing scans. All the plots are normalized to their highest value, and these are not an indication of OP efficiency.

9.2.1.6 Final OP efficiency

Finally, after optimizing these parameters, we can measure the OP efficiency by both the methods mentioned before. For these measurements, ideally we would have the same conditions, and measure the number remaining with and without the microwaves applied, which will tell us the efficiency. However, for experimental reasons, we instead change the microwave detuning to be very far detuned such that effectively there are no microwaves.

Figure 9.7a shows the efficiency measured using the first method, where we remove everything from $F = 0$ and measure the remaining number. We see that the final remaining number is around 18%, meaning that the OP efficiency is around 100-18 = 82%. Figure 9.7b shows the same measurement, but now with the microwaves applied to the other three hyperfine levels. In this case, this is a direct measurement of the number remaining in $F = 0$, and we see that the final remaining fraction is around 83%, which agrees with our other measurement.

9.2.1.7 Effect of leaked light

As mentioned in the experiment protocol, we used the RF switch on the AOM to turn off the OP light during the pump out period. There are a few subtleties here in our implementation that became clear over the course of this measurement.

First off, the way we control our AOM is by using the 2-frequency method, where we have a resonant frequency used for the primary power control, and an off-resonant frequency used to match the thermal load during changes to the resonant power. Thus, there is an inverse relationship between the RF/optical power applied for the resonant and off-resonant frequencies: when we want max optical power, the resonant RF is at max power and the off-resonant is at min power, and vice versa.

Now, during the pump out period, we want no OP light since that can induce additional cycling which is undesired. Our initial approach was to just set the resonant RF power to its lowest value, and we did not turn off the RF switch either. We saw that this led to a slow decay vs pump out time, and we thought that maybe the Rabi frequency is just very low. But we saw that given long enough (around 10-15 ms), we could totally deplete everything from $F = 0$, which seemed too good to be true.

We then measured how much power we were getting out of the fiber, by just setting the resonant RF power to its lowest value, and we found that we were coupling a few 100 μW. This was all

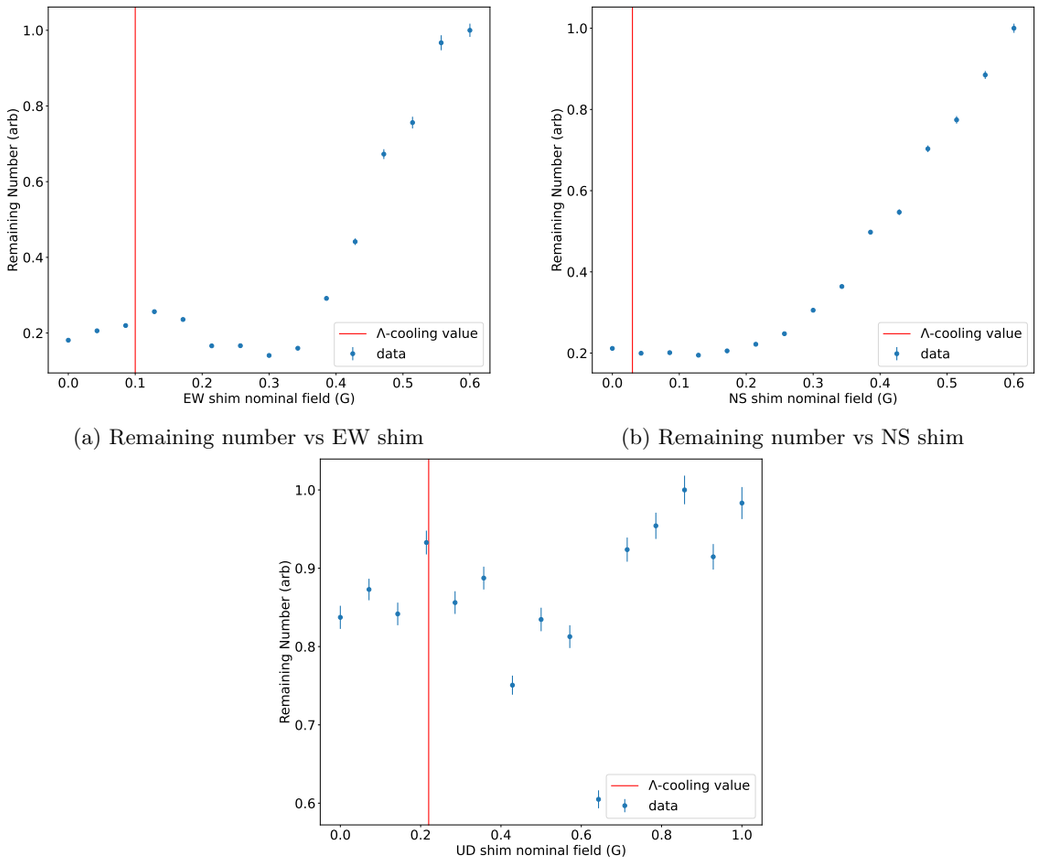

(a) Remaining number vs EW shim

(b) Remaining number vs NS shim

(c) Remaining number vs UD shim

Figure 9.6: Optimal OP parameters for the shim coils. The red lines show the nominal values used for Λ-cooling. All the plots are normalized to their highest value, and these are not an indication of OP efficiency.

mostly the off-resonant light, since it is now at max power, and is close enough in frequency to the resonant that some of it is fiber coupled. This also explains the slow decay, because this off-resonant light can now cause additional cycling into $F = 0$ which are then pumped out with the microwaves.

We then tried the opposite, i.e. setting the RF power on the resonant to max, but also turning off the RF switch, and saw that a few 100 nW of light was coupled now. We took the above scans with this setting, but we should note that this small amount of light still caused some leakage, albeit at a much slower rate.

To get a completely clean measurement, we should fully extinguish the OP light using a mechanical shutter. However, these shutters typically take 1.7 ms to close. We implement this technique in the ODT OP results shown next.

9.2.1.8 OP in ODT

A concern in implementing OP in the ODT is the differential Stark shifts between the states. This could lead to inefficient scattering for example making in OP worse. For this, we follow the same procedure as in the free space case, except we load the ODT after the blue MOT, and we also shutter the MOT light entirely during the detection period. We performed the same sets of frequency scans but this time in the ODT, and the results are shown in Figure 9.8. First of all, we see, as expected, that both the V_{00}^{d} and V_{00}^{M} carriers are blue-shifted by around 20 MHz. This is expected as the scalar shift for our ODT is around 20-25 MHz, with all the states being shifted down. We also see that there is still not a huge dependence on V_{00}^{M} sideband, as was seen in the

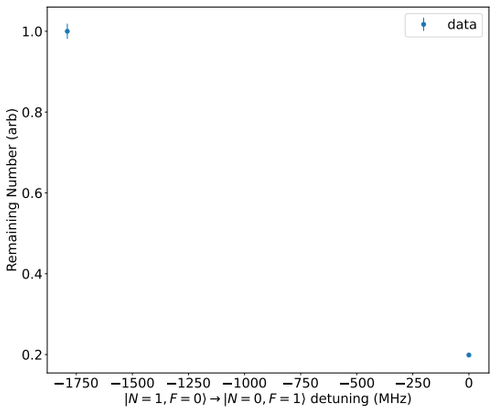

(a) Remaining number vs MW1 detuning

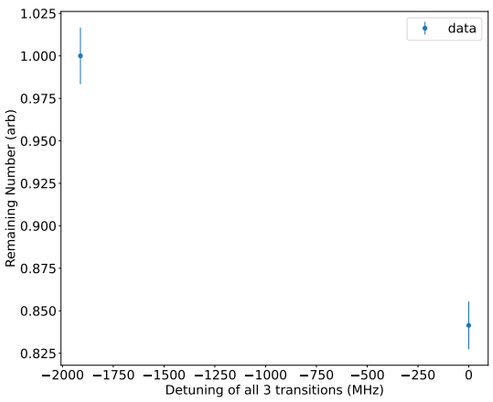

(b) Remaining number vs MW2/3/4 detuning

Figure 9.7: OP efficiency measurements. For the same conditions, the microwave detuning is varied from far detuned to resonance to extract OP efficiency.

free space case as well. We see that the other OP parameters are not really affected by the ODT, and remain more or less the same. Ultimately, we are able to obtain around 83% OP efficiency in the ODT as well, very much in line with what we expect.

With this result, we now have around 83% of the molecules in our ODT prepared in the $|N=1,F=0\rangle$ state. We would now like to obtain coherent control between the rotational states in the molecules, and ultimately transfer the molecules into the absolute rovibrational ground state. This will be discussed in the next section.

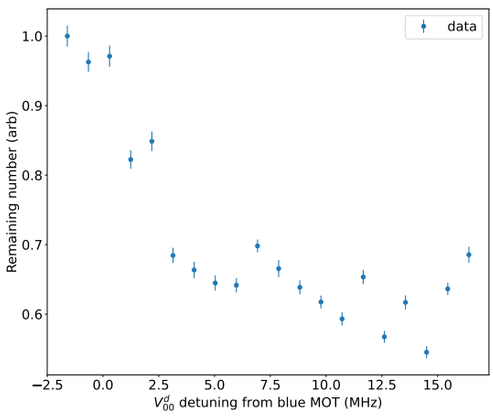

(a) Remaining number vs $V_{00}^{\dagger}$ frequency

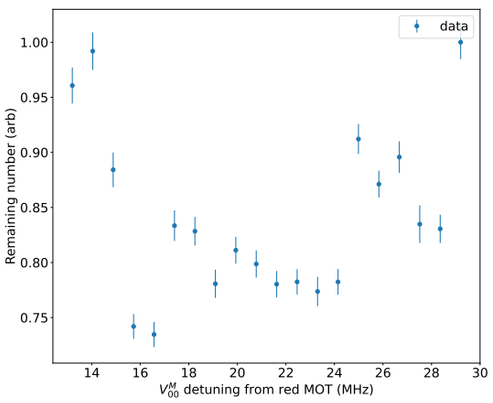

(b) Remaining number vs V_{00}^{M} frequency

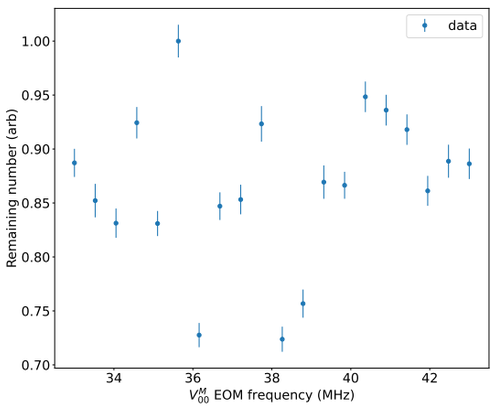

(c) Remaining number vs V_{00}^{M} EOM frequency

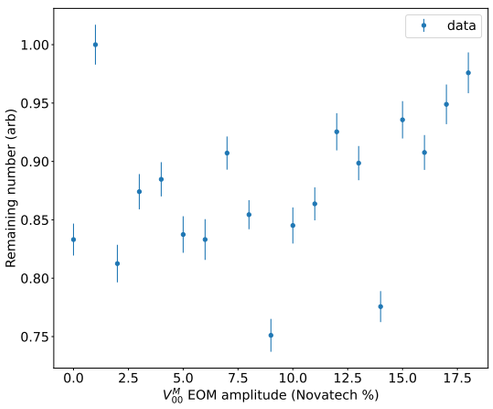

(d) Remaining number vs V_{00}^{M} EOM amplitude

Figure 9.8: Optimal OP parameters for the lasers in the ODT. All the plots are normalized to their highest value, and these are not an indication of OP efficiency.

9.2.2 Coherent control

Now that the most of the molecules are in a single quantum state, we next aim to drive Rabi oscillations between the rotational states. In this section, we will use the notation $|N, F, m_F\rangle$ to label all the states. We shall start by driving Rabi oscillations between the $|1,0,0\rangle$ and $|0,1,0\rangle$ states. This is a strong dipole allowed transition and so is easy to drive.

We closely follow the treatment from [193]. We consider a molecule that is subject to a microwave electric field of the form $\vec{E} = \frac{1}{2}\mathcal{E}_0\vec{\epsilon}e^{-i(\omega t+\phi)} + \text{h.c.}$ with angular frequency ω close to the resonance of the transition at frequency ω_0. $\mathcal{E}_0$ is the amplitude of the electric field, $\vec{\epsilon}$ is the polarization, and ϕ is the phase. If the two states in question are well separated from any other states, then we can restrict ourselves to a two-level system, where, upon transforming to the rotating frame and using the rotating wave approximation, the evolution of the states is given by the Schrodinger's equation:

$$i\hbar\frac{d\,|\psi(t)\rangle}{dt} = H_{RF}\,|\psi(t)\rangle \tag{9.1}$$

where H_{RF} is the Hamiltonian in the rotating frame

$$H_{RF} = -\frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}(\cos\phi\sigma_x + \sin\phi\sigma_y) \tag{9.2}$$

where $\Omega = -\vec{d}\cdot\vec{\epsilon}\mathcal{E}_0$ is the Rabi frequency, $\Delta = \omega - \omega_0$ is the detuning from resonance, ϕ is the phase of the microwave drive, and $\sigma_{x,y,z}$ are the usual Pauli matrices. The general solution to the Schrodinger equation is then

$$|\psi(t)\rangle = U_{RF}(t,\Omega,\Delta,\phi)\,|\psi(0)\rangle \tag{9.3}$$

where the propagator is

$$U_{RF}(t,\Omega,\Delta,\phi) = \cos\left(\frac{\Omega_{tot}t}{2}\right)\mathbb{1} + i\sin\left(\frac{\Omega_{tot}t}{2}\right)\left[-\frac{\Delta}{\Omega_{tot}}\sigma_z + \frac{\Omega}{\Omega_{tot}}\left(e^{-i\phi}\sigma_+ + e^{i\phi}\sigma_-\right)\right] \tag{9.4}$$

where $\Omega_{tot} = \sqrt{\Omega^2+\Delta^2}$ is the generalized Rabi frequency, $\mathbb{1}$ is the identity operator, and $\sigma_\pm$ are the Pauli raising and lowering operators.

9.2.2.1 Rabi driving in free space

We start with trying to drive Rabi oscillations between $|1,0,0\rangle \rightarrow |0,1,0\rangle$ in free space. For this, we first roughly find the center frequency by scanning the microwave frequency in the presence of the pump-out laser. We clearly resolve the three components corresponding to the three transitions from $|1,0,0\rangle \rightarrow |0,1,(0,\pm1)\rangle$ under a magnetic field, as seen in Figure 9.9. Even though the horn is nominally only emitting π polarization, the molecules experience all three components due to reflections from the chamber walls.

We choose the microwave frequency for our transition of interest and try driving Rabi oscillations. This is a bit of an iterative process, because we do not yet know the center frequency well enough to be perfectly on resonance, and we will not get the full contrast if we are not perfectly on resonance. Thus, our initial goal is to find what the Rabi frequency is for a given microwave power, close to resonance. Our molecules start in $|1,0,0\rangle$, and ultimately, we are also measuring the population in $|1,0,0\rangle$. Thus, based on eq. 9.3, we see that the number remaining in $|1,0,0\rangle$

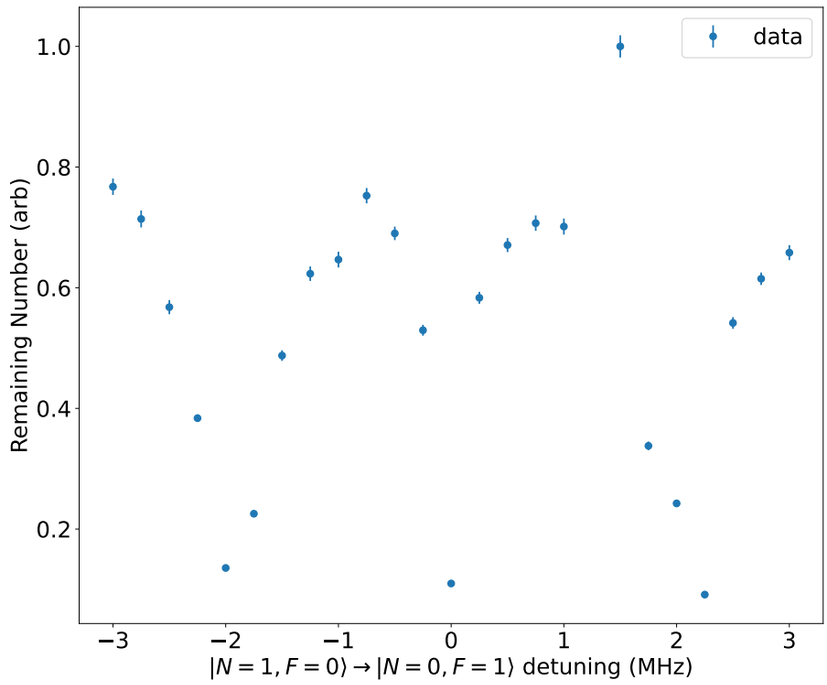

Figure 9.9: Frequency scan of $|1,0,0\rangle \rightarrow |0,1,(0,\pm 1)\rangle$ microwaves in the presence of the pumpout laser. The three transitions are clearly resolved, and show a difference of around 2 MHz, corresponding to a residual field around 1 G. Additionally, this implies that all microwave polarizations are experienced by the molecules.

after a time t is:

$$f(t,\Omega,\Delta) = P_{meas} = |\langle 1,0,0|U_{RF}(t,\Omega,\Delta,\phi)|1,0,0\rangle|^2 \tag{9.5}$$

We also see decoherence in our system, and to account for this, we can fit to a generalized function of the form:

$$P_{meas} = \left(a + b\cos^2\left(\frac{\Omega_{tot}t}{2} + \phi\right)\right)e^{-t/\tau} + c \tag{9.6}$$

where $a = \frac{\Delta^2}{\Omega_{tot}^2}$, $b = \frac{\Omega^2}{\Omega_{tot}^2}$, τ is the coherence time, and ϕ, c are fit constants.

An initial attempt at driving Rabi oscillations is shown in Figure 9.10, where we see low contrast of the oscillations, but we are able to extract a π pulse time of $\tau_\pi \approx 11\,\mu$s. We also see very fast decoherence of around 30 μs. Regardless of this, we can use the π pulse time to perform a detuning scan for a Rabi frequency $\Omega = \frac{\pi}{\tau_\pi}$. In this case, we see that the remaining population in $|1,0,0\rangle$ is

$$f\left(\tau_\pi, \frac{\pi}{\tau_\pi}, \Delta\right) = 1 - \frac{1}{1+\frac{\tau_\pi^2\Delta^2}{\pi^2}}\sin^2\left(\frac{\sqrt{\left(\frac{\pi}{\tau_\pi}\right)^2+\Delta^2}}{2}\tau_\pi\right) \tag{9.7}$$

and we can fit the remaining population to a generalized function of the form $P_{meas} = a + b\,f(\tau_\pi, \frac{\pi}{\tau_\pi}, \Delta)$.

Such a π pulse spectroscopy scan is shown in Figure 9.11 for a π pulse duration of $\tau_\pi = 70\,\mu$s, corresponding to a Rabi frequency $\Omega \approx 2\pi \times 7$ kHz. We can deplete almost 80% of all molecules in the trap, which, when combined with the OP efficiency, means that the π pulse can deplete around 94% of the molecules from $|1,0,0\rangle$ to $|0,1,0\rangle$.

9.2.2.2 Shim coil calibration

Upon using this center frequency to drive Rabi oscillations, we still saw fast decoherence, on the order of 30 μs. Our initial guess for this low coherence time was that perhaps the shim coil fields

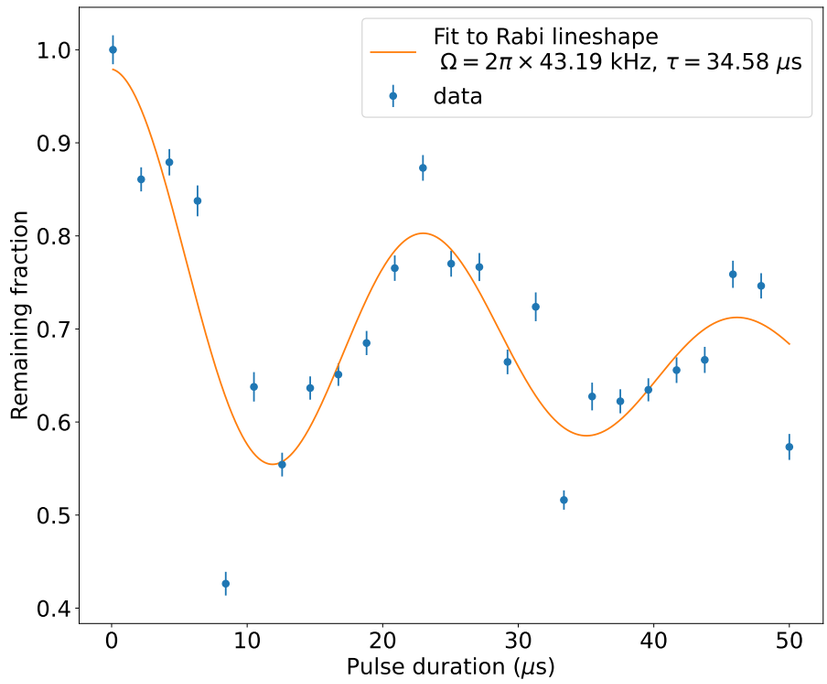

Figure 9.10: Initial attempts to drive Rabi oscillations between $|1,0,0\rangle \rightarrow |0,1,0\rangle$. The contrast is low but we are able to extract a π pulse time. We find a Rabi frequency $\Omega \approx 2\pi \times 43$ kHz. Rapid decoherence is also seen with a $1/e$ decoherence time of 35 μs.

are drifting shot to shot. We also did not actually know what was the field we were applying for a given current, as we had not calibrated the shims carefully before.

We decided to use the Rabi spectroscopy technique to first calibrate our shims and then apply a known shim field. This calibration is again done following the procedure laid out in [193]. Here, we look at the microwave transition from the $|1,0,0\rangle \rightarrow |0,1,1\rangle$. The $|0,1,1\rangle$ state experiences a Zeeman shift in the presence of a magnetic field, and the center frequency for this transition shifts as a result. Our initial attempts at driving this transition discovered that the center frequency changes vs time after the shim currents are turned off, due to eddy currents in the chamber. This is mapped out in Figure 9.12, where we see that the $1/e$ timescale is around 1 ms, with the frequency having fully settled by around 3 ms. We can also extract the changing B-fields from the center frequency shifts by solving the Zeeman Hamiltonian. We see that the B-field decay is a bit faster ($1/e$ timescale 0.6 ms), as expected since the frequency shift for this transition is quadratic in B. We see that the center frequency is $\omega_0 = 14890.241726$ MHz.

As a result, to calibrate the shims, we always change the current, and wait 3 ms before doing spectroscopy. For each shim coil, in turn, we measure the transition probability vs microwave frequency for a number of shim currents and extract the center frequencies. Here, we need to be careful to avoid the zero-field crossing where all three Zeeman substates in the $|0,1\rangle$ manifold are degenerate and the two-level model breaks down. Away from this region, the shift in the center frequency under an applied current in the shim coils is:

$$\omega_0 = \omega_0^0 - \frac{\mu_B}{\hbar}\sqrt{(\alpha_i I_i + B_i^0)^2 + B_0^2} \tag{9.8}$$

where ω_0^0 is the frequency of the $|1,0,0\rangle \rightarrow |0,1,0\rangle$ transition, α_i is the shim current to field calibration that we want to measure, I_i is the applied shim current, B_i^0 is the residual field along the shim axis, and B_0 is the residual field along the other two axes. The results are shown in Figure 9.13, with $\alpha_1 = 0.228$ G/A, $\alpha_2 = 0.212$ G/A and $\alpha_3 = 0.538$ G/A. We can see that the lab

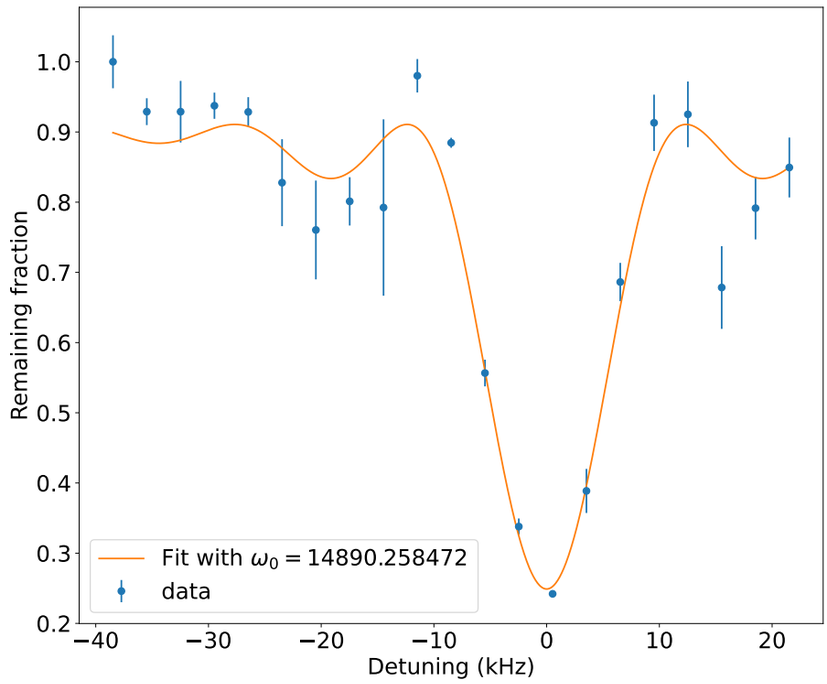

Figure 9.11: π pulse spectroscopy on the $|1,0,0\rangle \rightarrow |0,10\rangle$ in free space for a π pulse time of $\tau_\pi = 70\,\mu$s corresponding to Rabi frequency $\Omega \approx 2\pi \times 7$ kHz. Under ideal conditions, around 75-80% of the molecules can be depleted, correponding to a π pulse efficiency $\approx 94\%$.

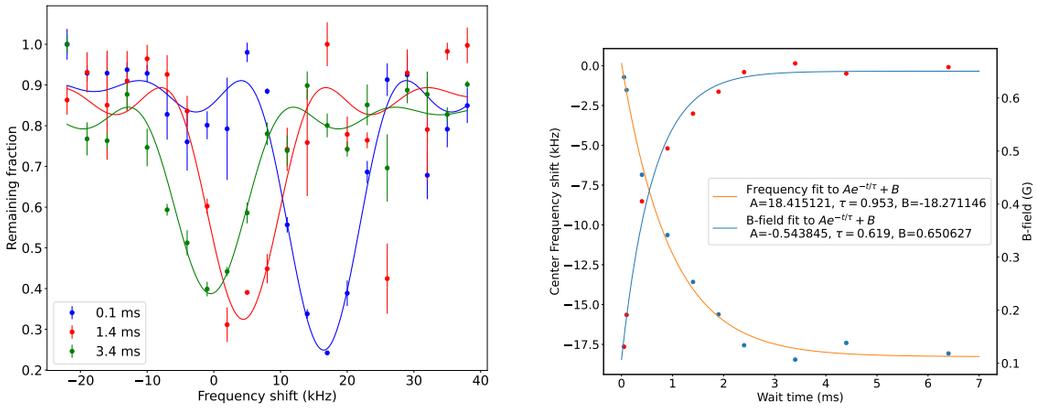

(a) Frequency shift vs time

(b) Frequency shift vs time

Figure 9.12: Shim current decay timescale. (a) π pulse spectroscopy on the $|1,0,0\rangle \rightarrow |0,1,0\rangle$ transition as a function of wait time showing the changing fields over time. (b) (left y-axis) Extracted center frequency shift vs time fit to an exponential decay with time constant around 1 ms. (right y-axis) Extracted B-field decay (from solving the Zeeman Hamiltonian) vs time fit to an exponential decay with time constant around 0.6 ms.

fields along the horizontal axes are around 0.2 G each and are largest along the up-down direction around 0.4 G. We also note that the shim currents needed for optimal Λ-cooling indeed cancel out earth's field, as expected.

After this shim calibration, we realized that the shim driver circuit had a major drawback in that it did not feedback on the current, and was thus susceptible to drifts over time. The old driver circuit used a MOSFET with a DAQ controlling the gate voltage. Here, the current was set by making a look-up table of the current vs gate voltage. However, it was discovered that in a normal experiment run, the current would slowly rise as the MOSFET heated up and the characteristics changed. We switched to a new current control circuit that used a high-power OPAMP that fed back on the voltage drop across a sense resistor, which made the current drifts go away.

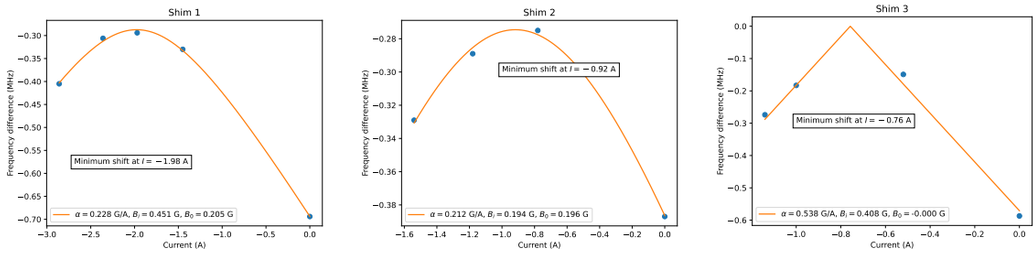

(a) Shim 1 calibration (b) Shim 2 calibration (c) Shim 3 calibration

Figure 9.13: Shim calibrations. Measured frequency shift of the $|1,0,0\rangle \rightarrow |0,1,1\rangle$ transition relative to the $|1,0,0\rangle \rightarrow |0,1,0\rangle$ transition as a function of the shim current for each shim. (a) Shim 1 (b) Shim 2 and (c) Shim 3. The points are the lineshape centers estimated from the fits.

9.2.2.3 Rabi decoherence

We now returned to our transition of interest, the $|1,0,0\rangle \rightarrow |0,1,0\rangle$ and tried Rabi driving again. However, in spite of these changes, and even with more stable shims, we still were not able to increase our coherence times beyond the 30 μs seen previously. In fact, we saw a marked dependence of the coherence time on the Rabi frequency, where a higher Rabi frequency led to faster decays. One example of a high Rabi frequency is shown in Figure 9.14, where the coherence time is only around 15 μs.

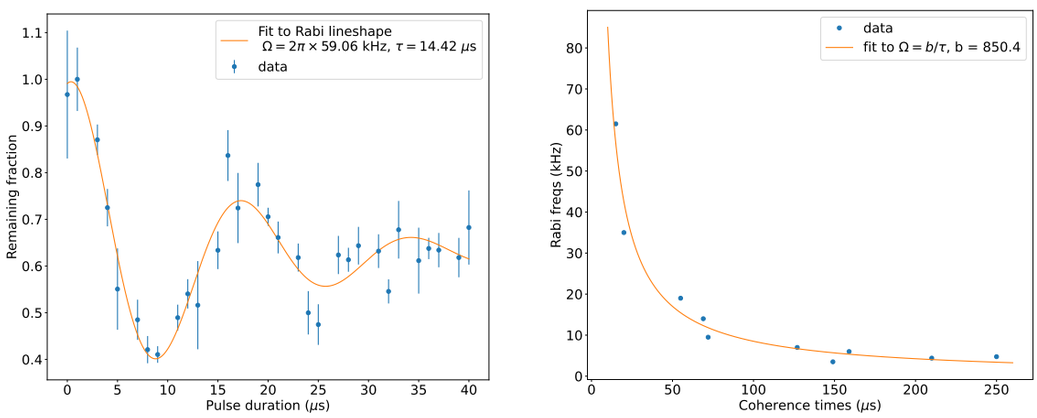

(a) High Rabi frequency (b) Rabi frequency vs coherence time

Figure 9.14: Rabi frequency vs coherence. (Left) High Rabi frequency with a short coherence time. (Right) Rabi frequency vs coherence time where the data is fit to an inverse relationship $\Omega = b/\tau$.

We then mapped out the dependence of the coherence time on the Rabi frequency, as shown in Figure 9.14b and find that the coherence time is inversely proportional to the Rabi frequency.

9.2.2.4 Ramsey spectroscopy

We need to figure out whether the microwaves are the problem or whether we do not have sufficient control over the external fields in the lab that are causing the fast decoherence. For this, we turned to the Ramsey technique. This is a very common technique used in precision metrology experiments and measurements of fundamental constants [17, 21, 194]. In such an experiment, the molecules are all initially prepared in a single quantum state; in our case the $|1,0,0\rangle$. Then, a $\pi/2$ pulse of duration τ_R and detuning Δ creates a super-position of the two states of interest, in our case the $|1,0,0\rangle$ and $|0,1,0\rangle$ state. This state is allowed to evolve for a time T before a second identical $\pi/2$ pulse is applied to bring the population back to the initial state. We perform these experiments in free space with the state preparation as described above. We collectively model the population measured in the $|1,0,0\rangle$ state at the end of the sequence, including decoherence mechanisms, using

the following equation:

$$f(\Delta, T) = \left(a + b\,\cos^2\left(\frac{\Delta(T)\,T}{2} + \phi\right)\right) e^{-T/\tau} + c \tag{9.9}$$

where $\Delta(T)$ is the detuning, and we allow this to be a function of the evolution time to account for the changing fields shown in Figure 9.12, T is the free evolution time, τ is the $1/e$ decoherence time and a, b, c, ϕ are fit constants.

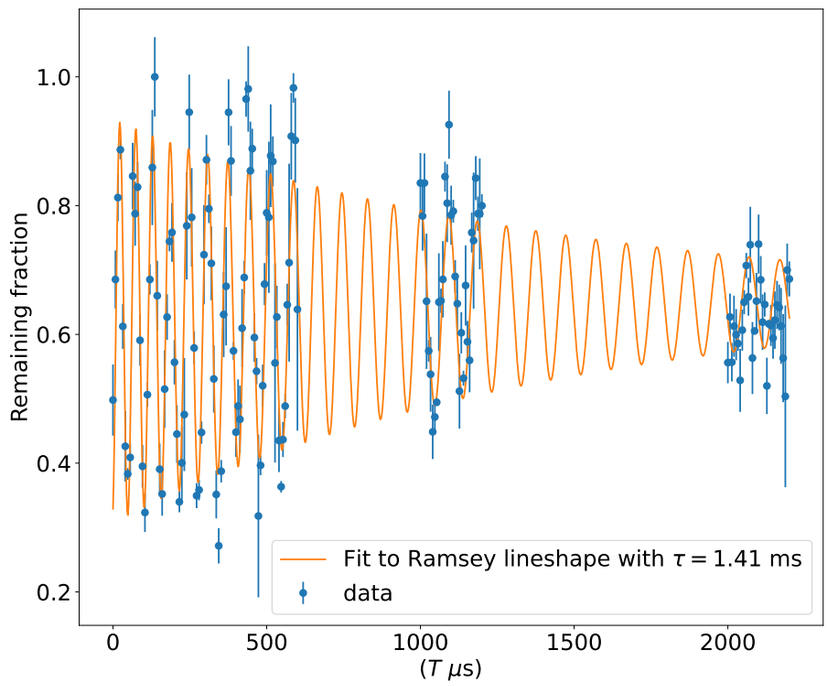

Figure 9.15: Ramsey oscillations in free space for molecules prepared in a coherent superposition of $|1,0,0\rangle$ and $|0,1,0\rangle$. The $\pi/2$ pulses have duration of 4.5 μs, and the free evolution time T is scanned. The detuning changes as a function of time according to Figure 9.12 and is accounted for in the fit. The initial detuning is around +20 kHz and it ends at around +8 kHz detuning. We find a decoherence time of around 1.41 ms, much longer than the Rabi decoherence time. The main decoherence mechanism is probably leaked light from the MOT AOM.

We perform such an experiment as shown in Figure 9.15 which was performed before the shim current decay timescale was discovered. In this case, the detuning is allowed to change as per Figure 9.12. The initial detuning was set to be around +20 kHz, and as seen in Figure 9.15, the detuning chirps and ends at around 8 kHz, and is accounted for in the fit. The $\pi/2$ pulse duration is 4.5 μs. We find a $1/e$ decay timescale of around 1.41 ms, orders of magnitude higher than the Rabi decoherence rate. The main decoherence mechanism is probably the leaked light from the MOT AOM, but we have not traced this yet.

This means that the fields in our chamber are well controlled shot-to-shot and are not the cause of the fast Rabi decoherence we see in the experiment. Current work is being pursued to determine the cause of this fast decoherence, and whether it is important for us.

9.2.2.5 Transfer to the rovibrational ground state

Ultimately, our goal is to transfer molecules to the rovibrational ground state. However, the transition $|1,0,0\rangle \rightarrow |0,0,0\rangle$ is nominally forbidden by angular momentum selection rules. A magnetic field can, however, induce mixing of the $|0,0,0\rangle$ state with the neighboring $|0,1,0\rangle$ state.

Under second order perturbation theory, the mixed wavefunction can be estimated as:

$$\begin{aligned} |0,0,0\rangle &\simeq |0,0,0\rangle + \frac{\mu_B B g_S}{E_{0,1,0} - E_{0,0,0}} \langle 0,1,0|T_0^1(\mathbf{S})|0,0,0\rangle \, |0,1,0\rangle \\ &\simeq |0,0,0\rangle + \frac{\mu_B B}{\Delta_{hfs}} |0,1,0\rangle \end{aligned} \tag{9.10}$$

where $\Delta_{hfs} = 107$ MHz is the hyperfine splitting in the ground state, B is the applied magnetic field, and μ_B is the Bohr magneton. Thus, the Rabi frequency for the $|1,0,0\rangle \rightarrow |0,0,0\rangle$ transition under an applied magnetic field is

$$\Omega \simeq \frac{\mu_B B}{\Delta_{hfs}} \tilde{\Omega} \tag{9.11}$$

where $\tilde{\Omega}$ is the Rabi frequency for the $|1,0,0\rangle \rightarrow |0,1,0\rangle$ transition. Using $\mu_B = 1.4$ MHz/G, we see that

$$\Omega \simeq \frac{1}{75} B \tilde{\Omega} \tag{9.12}$$

where B is the applied field in Gauss. Our shims have a max field capability of (shim 1, shim 2, shim 3) = ($\pm$1.5,$\pm$1.9,$\pm$4) G. Thus, if we only apply the field along one of the shims (to keep the analysis simple), we see that the best Rabi frequencies we can achieve are

$$\Omega \simeq \frac{1}{20} \tilde{\Omega} \tag{9.13}$$

which would equate to around 10s of kHz Rabi frequencies for our current setup. This might be enough to directly drive a π pulse to transfer population. Otherwise, we will have to resort to something like a Landau-Zener sweep, as shown in [141]. This remains to be explored.

9.3 Preparing for microwave shielding

9.3.1 Estimates of microwave parameters

As we saw in Chapter 8, our molecules undergo rapid inelastic collisions in the trap, and we need to figure out how to suppress these. There are two complementary approaches to this. The first one, called electric field shielding was briefly alluded to in § 2.6, and relies on a large static electric field to bring neighboring rotational states close to each other to tune the scattering length [82, 192, 195]. For SrF, this field requirement is around 15 kV/cm, a non-trivial field to generate, and so we shall not be pursuing this direction. This technique is, however, quite successful, and was demonstrated to work for KRb molecules [42, 83, 196] and is being pursued for other molecules [197, 198].

The other approach is called microwave shielding [80, 81, 199–201], and this relies on a blue detuned microwave field to induce dipole-dipole interactions which in turn induces a repulsive barrier between colliding molecules. This barrier is repulsive in all three dimensions, and was shown to provide excellent suppression of inelastic collisions [43, 44, 183, 184, 202] and even led to the production of degenerate gases. We decided to pursue this since it is technically easier to produce high power microwaves.

We briefly review the principle of microwave shielding here. Consider a molecule in the ground rotational state ($|g\rangle = |N = 0, m_N = 0\rangle$) that is dressed with circularly polarized and blue detuned microwaves with the state $|e\rangle = |N = 1, m_N = +1\rangle$ to make two states $|+\rangle = \frac{1}{\sqrt{2}}[|g\rangle + |e\rangle]$ and

$|-\rangle = \frac{1}{\sqrt{2}}[|g\rangle - |e\rangle]$. The uncoupled states $|N=1, m_N=-1, 0\rangle$ remain as spectator states $|0\rangle$. These dressed states have a dipole moment and thus undergo dipole dipole interactions (DDI) that scale as $\pm 1/R^3$. When the molecules are far apart, these states are separated by the dressing energy $\hbar\Omega$ and the dipoles follow the strong AC field of the microwave, thus exhibiting an effective dipolar interaction at long range. However, as the molecules approach each other, the DDI starts becoming stronger and it starts reorienting the molecules along the internuclear axis rather than the AC field. In the crossover regime where the DDI is stronger than the drive, molecules prepared in $|+\rangle$ undergo a repulsive interaction, and rise in energy to enter the $|++\rangle$ branch; we shall call this the shielding state. Molecules in $|-\rangle$ undergo an attractive interaction and drop in energy; we shall call this the anti-shielding state. There is an avoided crossing between these states at the crossover point, with a separation determined by the microwave drive Rabi frequency $\hbar\Omega$.

Based on the above discussion, we find that there are 3 main requirements for this scheme to work well:

1. **High Rabi Frequency**: For microwave shielding to be effective, we want that the molecules dressed in $|+\rangle$ actually end up in the repulsive branch ($|++\rangle$ branch) as they approach each other, i.e. adiabaticity is required. Furthermore, we do not want to drive unwanted transitions to the other dressed rotational states $|N=1, m_N=-1, 0\rangle$ as these are all attractive branches. This means we need a high Rabi frequency for the largest separation between the states.

2. **Pure circular polarization**: The avoided crossing coupling between the dressed states and the DDI potential states depends on the orientation of the polarization relative to the intermolecular axis. For example, linear polarization can couple to a non-avoided crossing and thus to unshielded states that will cause inelastic loss. This can be partially mitigated by electric fields that push the π-component away from resonance but we should aim for circularity [199, 200].

3. **Extremely low phase noise**: We also need to ensure molecules are always prepared in $|+\rangle$, and this means we need spectral purity, i.e. very low phase noise. Specifically, the phase noise at a frequency offset Ω from the carrier is important as this can cause $|+\rangle$ to be driven to $|-\rangle$ or the other non-shielded states.

In the following, we estimate the microwave shielding requirements for SrF. As we saw, in order for the microwave shielding to be effective, one needs a high Rabi frequency Ω and a positive detuning Δ. Furthermore, if we want to avoid microwave-induced loss due to the hyperfine structure, we need $\Omega \geq \gamma$ ($\gamma = 2\pi \times 75$ MHz for SrF) where γ is the spin-rotation constant. Thus, ideally we should aim for $\Omega \sim 2\pi \times 100$ MHz for effective shielding. The Rabi frequency is related to the dipole moment and the applied electric field as

$$\Omega = \frac{dE_0}{\hbar} \tag{9.14}$$

where E_0 is the peak electric field. The transition dipole moment of interest for SrF for the shielding transition is $d \approx 1$ D. If we aim for $\Omega = 2\pi \times 50$ MHz, we get that we need a microwave intensity of $I \sim 13$ W/cm^2. The $X^2\Sigma, N=0 \rightarrow N=1$ transition in SrF is at 15 GHz frequency, which corresponds to $\lambda = 2$ cm. For a diffraction limited Gaussian microwave beam with a focus

spot $w_0 = 1$ cm (i.e. beam diameter is 2 cm), we find that we need a microwave power of $P = 20$ W to achieve a peak intensity of 13 W/cm^2.

In practice, achieving this tight of a focus spot at the exact location of the molecules will be highly improbable, and if we assume that we can get a focus spot $w_0 = 2.5$ cm at the molecules, then we would need a power of $P \approx 100$ W to achieve the same intensity. Thus, we need microwave equipment that will support generation of high powers up to 100 W. The Doyle group also saw some shielding at $\Omega \approx 2\pi \times 20$ MHz, so we can be conservative and aim for this as the worse case scenario, and this would require around 50 W of power.

Furthermore, we want a very low phase noise microwave source in order to avoid 1-body losses caused by absorption of microwave photons in different sidebands than the carrier frequency, as we shall see in § 9.3.3. Our initial design used a SynthHD Pro to generate the microwaves, however it had a high phase noise which would limit the trap lifetime.

We next outline how we tackle each of these issues. First we address the generation of high power circularly polarized microwaves. Next, we discuss how to suppress the phase noise further by using microwave filters. We then show how we can focus the microwaves in the chamber to realize the high Rabi frequencies. Finally, we summarize the results of microwave testing.

9.3.2 Generating high power circularly polarized microwaves

The basic shielding mechanism works well only if the microwaves are circularly polarized (σ_+). There are ways to get around with not having perfect circular polarization (such as adding an external electric field to make the π transitions out of resonance) but it will greatly simplify things if we have as circular polarization as possible. Here we outline our technique to produce circular polarization using a low phase noise microwave source.

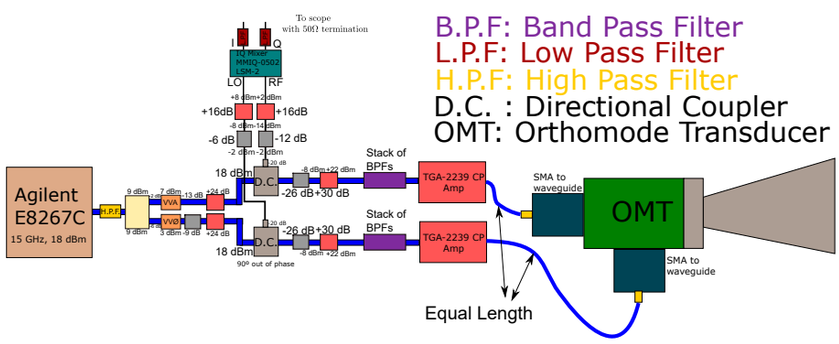

Figure 9.16: Schematic of high power low phase noise microwave generation at 15 GHz. The Agilent E8267C will be used to generate microwaves at 15 GHz which will be split into two arms, one of them being 90° out of phase with the other. These are then amplified separately and combined with an orthomode transducer to produce circularly polarized microwaves out of the horn. Additionally, directional couplers on each arm of the setup allow mixing of the two signals and thus monitoring the phase on an oscilloscope. We use iris-coupled band pass filters before the final amplifiers to further suppress the phase noise.

The method is shown in Figure 9.16 and is fairly simple. We will use a Agilent E8267C to generate microwaves at 15 GHz, 18 dBm. The output of the Agilent will then be split into two arms, with a voltage variable phase (VVϕ) controller on one of the arms and a voltage variable attenuator (VVA) on the other arm. The VVϕ controls the relative phase of the two arms such that they are 90° out of phase. The VVϕ has a certain insertion loss, and the VVA will be used to balance the power on the other arm. These two arms then go through some attenuators and preamps to adjust the powers. We use a stack of iris-coupled band pass filters (XAHX

microwave HXLBQ-DTA641) along each path to further suppress the phase noise, before they are independently amplified with two amplifiers (TGA-2239 CP) to produce 50 W on each arm of the setup. One of the arms of the setup goes to one input of the OMT and the other arm of the setup to the other input of the OMT. The OMT has inputs that are orthogonal linear polarizations, essentially like a polarization beam splitter for microwave frequencies. Thus, combining two beams with orthogonal linear polarizations and 90° out of phase should produce circular polarization out of the OMT, which will then be emitted from the microwave horn and sent to the experiment.

For diagnostics, we also have a directional coupler on each arm of the setup. The two samples from the two arms are combined on an IQ-mixer. IQ mixers are two separate mixers with a phase delay between the mixers. For a normal mixer with just one output port, one can determine only the magnitude of the phase shift but not the direction (i.e. 45° is the same as 135°). However, by using two such mixers with a phase delay between them, we can be rid of this ambiguity[2]. Basically, by measuring the voltage on the I and Q arms of the mixer, we can determine the phase

$$\theta = \tan^{-1}\left(\frac{V_I}{V_Q}\right) \tag{9.15}$$

This is very simple to implement and we can make a lookup table for voltages versus phase. It should be noted that each arm of the setup has nominally the same path length. This is to ensure that no additional phase shifts are incurred on account of different cable lengths. The parts are also chosen to be the same or with identical dimensions for both arms.

9.3.3 Suppressing the phase noise

The dressed state lifetime depends linearly on the phase noise of the microwaves sent into the experiment, as evidenced by some data from recent experiments:

- Doyle group experiment saw a lifetime of around 500-600 ms and their phase noise was around -154 to -155 dBc/Hz based on their function generator model (an Agilent E8247D) [183]
- The Bloch group showed 3 data points with 3 function generators [43, 203].
  - R&S SMF100A - 300 ms - phase noise around -148 dBc/Hz
  - Agilent E8267D - 600 ms - phase noise around -152 to -153 dBc/Hz
  - R&S SMA100B - 1.8 secs - phase noise around -157 dBc/Hz
- The Bloch group also mentions in their paper that they see a linear scaling with worse microwave sources [43].

Thus, every 3 dB suppression in phase noise leads to around a factor of 2 gain in the lifetime. This can be understood as follows. The separation between the dressed states is Ω and the phase noise can induce undesired transitions between these states if there is some available microwave intensity at this frequency. This is precisely the phase noise of the source. From [183], the population in the undesired dressed state is given by:

$$P = \frac{1}{2}\left[1 - \exp\left(-\frac{\pi^2}{2}\Omega^2 S_\phi(\Omega)t\right)\right] \tag{9.16}$$

[2]See `https://www.markimicrowave.com/blog/all-about-mixers-as-phase-detectors/` for more details

where Ω is the Rabi frequency and $S_\phi(\Omega)$ is the phase noise at an offset frequency equal to Ω from the carrier. The dressed state lifeimet is thus

$$\tau_D = \frac{2}{\pi^2 \Omega^2 S_\phi} \tag{9.17}$$

This means that to achieve a 3 s dressed state lifetime for a 50 MHz Rabi frequency, we would need a phase noise as low as $S_\phi = -166$ dBc/Hz at a 50 MHz offset. In the following, we look at the different components in our microwave generation that can add phase noise, and how we can suppress it.

9.3.3.1 Phase noise from the amplifier

Due to the stringent requirements, we need to know if any of the amplifiers in our microwave circuit adds phase noise and if so, how can we mitigate it. Amplifiers mainly contribute to the phase noise in two forms: white noise and 1/f noise (see Figure 9.17) [204, 205]. Close to the carrier, the amplifier has 1/f noise, which is the straight line part of the graph, and far away from the carrier, the amplifier has white noise, which is independent of frequency.

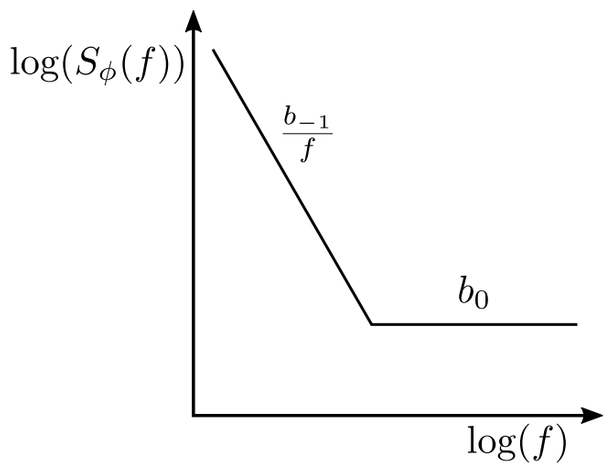

Figure 9.17: Phase noise vs frequency for a typical amplifier on a log-log plot. Close to the carrier, the amplifier has a 1/f noise curve, while far from the carrier the amplifier has white noise which is independent of frequency

Together, this can be modeled as:

$$S_\phi(f) = b_0 + \frac{b_{-1}}{f} \tag{9.18}$$

The point where the noise spectrum goes from being 1/f noise to white noise is called the corner frequency, which is around 10-100 kHz for most amplifiers. Thus, for our case of interest where we are looking at ~ 10 MHz away from the carrier, we are only interested in the white noise component of the phase noise.

The white noise contribution is easy enough to calculate if you know the noise figure of the amplifier:

$$b_0 = \frac{F k_B T_0}{P_0} \tag{9.19}$$

where F is the noise figure of the amplifier, k_B is the Boltzmann constant, T_0 is the temperature at which the noise figure is specified (and also the operating temperature, usually room temperature or 290K), and P_0 is the input power to the amplifier. A useful relation is $k_B T_0(290\text{ K}) = 4\times10^{-21}$ J =

-174 dBm/Hz

Unfortunately for us, and for most power amplifiers, the manufacturer does not specify a noise figure, which makes this calculation hard. Both the amplifiers used in the Doyle [183] and Bloch [43] experiments do not specify any noise figure either. Typical amplifiers have noise figures in the 1-10 dB range, so we can be conservative and choose a noise figure of 10 dB. Note that it is possible to measure this with a spectrum analyzer, but we just assume a 10 dB noise figure. Another note of caution is that the noise figure usually increases a fair bit if the amplifier is saturated (or hot), and we will be close to saturation when we actually operate.

The input power we will send to the amplifier when operating at full power is $P_0 = 22$ dBm. Our amplifier heats up a fair bit at full power, and if we operate the amplifier at 50°C, we get $k_B T_0 \approx -173$ dBm/Hz. Then we see that:

$$b_0 = 10 \text{ dB} - 173 \text{ dBm/Hz} - 22 \text{ dBm} = -185 \text{ dBc/Hz} \tag{9.20}$$

quite low. The main benefit here is because we are operating the amplifier at such high powers, that the white noise component really diminishes *relative* to the carrier.

An added benefit is that white noise reduces linearly when you add more amplifiers in parallel (and increases linearly if amplifiers are in series). In our case, we will be running 2 amplifiers in parallel to produce our circular polarization. As a result, our ideal white noise limit would be somewhere close to $b_0 = -188$ dBc/Hz. From this analysis, we conclude that we need not worry about phase noise contributions from the amplifiers.

9.3.3.2 Suppression of source phase noise

Our function generator is a Keysight E8267C. The phase noise from the datasheet is shown in Figure 9.18. We can be conservative here and say that we can expect a phase noise of around -153 dBc/Hz, similar to what the Doyle group saw with their Keysight.

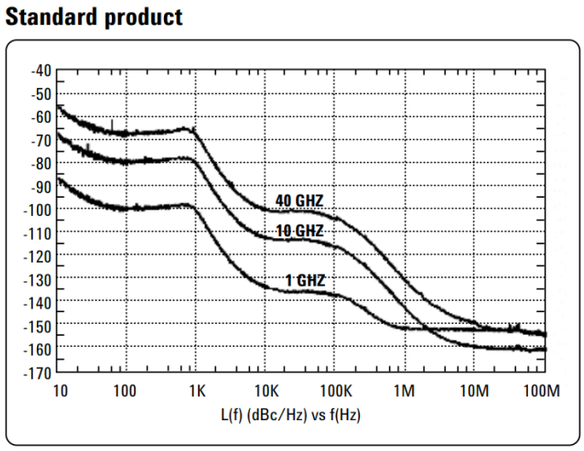

Figure 9.18: Phase noise vs frequency for our function generator.

By itself, this means we will get a dressed state lifetime of around 500-600 ms, clearly not enough to do much except see some suppression of inelastic collisions. From the data above, we would need to suppress the phase noise by at least 15-20 dB. One option to get better phase noise is to buy a better microwave source, but this is prohibitively expensive, with better sources costing in the $70-80k range. The other option is to buy or make a very narrow band-pass filter centered

around the frequency of interest, such that the sidebands are suppressed.

Such filters are commercially available and are called iris coupled bandpass filters. The mechanism here is that the filter consists of various sections, each of which is a rectangular half-wavelength resonator, and the sections are separated by irises. The more sections you have, the narrower the bandwidth can be, but the longer and bulkier the filter will be as well. These filters can typically provide a bandwidth of around 0.1-1 % of the carrier frequency. A company in China[3] makes such filters, and they provided us with a simulation of such a filter centered at 14.975 GHz.

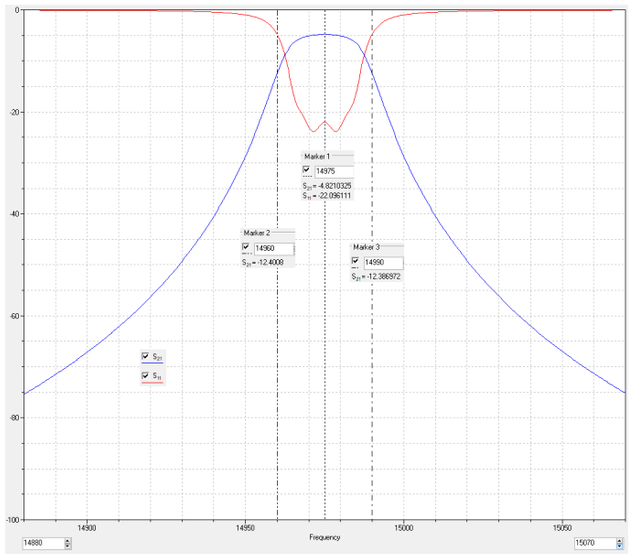

Figure 9.19: Simulated performance of an iris coupled bandpass filter centered at 14.975 GHz.

The attenuation of the filter 20 MHz away from the carrier is around -22 dB, however, relative to the carrier attenuation, it is actually only around -17 dB. A graph showing relative attenuation vs relative frequency is shown below (see Figure 9.20). Thus, one filter alone is probably not good enough for our purpose, and we probably want to use 2 of these back to back in order to get better suppression.

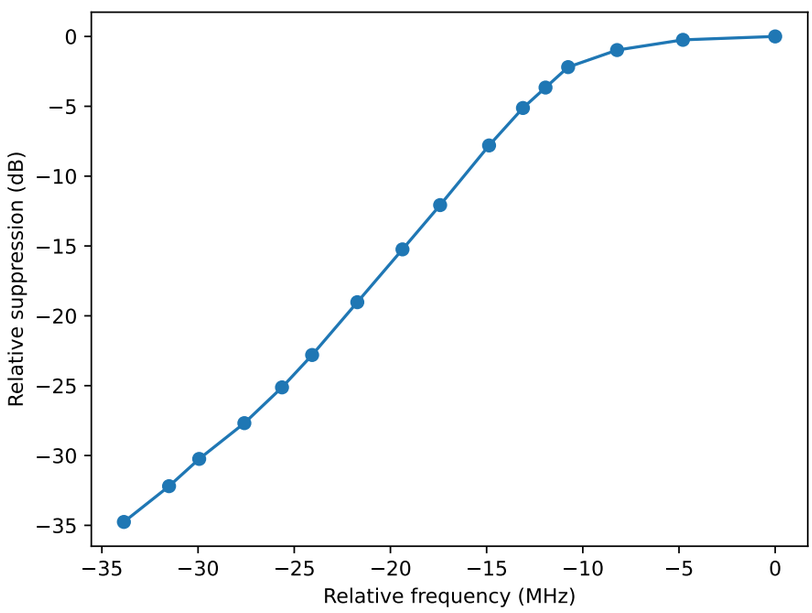

Figure 9.20: Relative attenuation vs relative frequency based on the simulated performance of an iris coupled bandpass filter centered at 14.975 GHz.

Another alternative would be to design the filter in such a way that the carrier is also not

[3] xahxmicrowave.com

"centered" on the curve, but lies 5 MHz away from the center. Then the carrier suppression would be around 0.5 dB, whereas the sideband suppression at 25 MHz would be -24.5 dB relative to the carrier thus winning us some more suppression.

9.3.4 Sending microwaves into the chamber

Once we have generated our microwaves, we need to send them in the chamber. This is not so easy because of geometric constraints caused by the viewports and the small aperture of the in vacuum coils. However, there exist quasi-optical techniques to focus microwave beams using metallic mirrors machined to be sections of an ellipse [206]. Here, we outline the design of such a mirror for our geometry.

9.3.4.1 Gaussian beam basics

A Gaussian beam is completely determined by the waist radius w_0 and wavelength of the beam:

$$z_c = \frac{\pi w_0^2}{\lambda} \tag{9.21}$$

$$R(z) = z + \frac{z_c^2}{z} \tag{9.22}$$

$$w(z) = w_0 \left[1 + \left(\frac{z}{z_c}\right)^2\right]^{0.5} \tag{9.23}$$

where $R(z)$ is the radius of curvature of the beam, $w(z)$ is the waist at location z and $z = 0$ is defined at the minimum waist location (i.e. at w_0).

9.3.4.2 Estimating the output beam shape of our conical horn

We are planning to use the LB-CNH-WC69-15 microwave horn from A-INFO. Figure 9.21 shows the dimensions of this horn. The waist radius and location of the radiated beam from such a conical horn can be estimated by [207]:

$$w_0 = 0.768r \left[1 + \left(\frac{1.85r^2}{\lambda L}\right)^2\right]^{-0.5} \tag{9.24}$$

$$\Delta = L \left[1 + \left(\frac{0.54\lambda L}{r^2}\right)^2\right]^{-1} \tag{9.25}$$

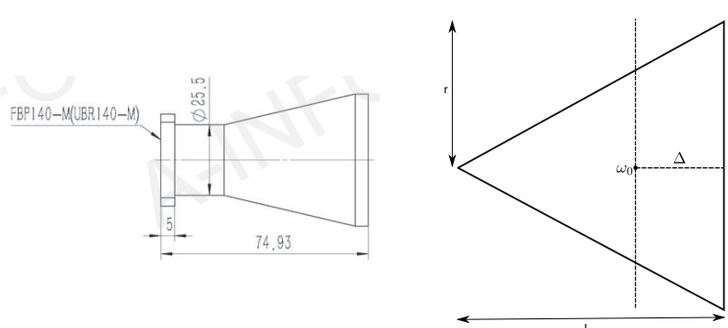

Figure 9.21: Estimating the output beam shape of our conical horn. (left) Actual dimensions of the microwave horn (LB-CNH-WC69-15) in mm. (right) Approximating the output of a conical horn as a Gaussian beam of waist w_0 at a distance Δ from the horn aperture.

Thus, the beam waist lies inside the microwave horn and is displaced from the aperture as shown in Figure 9.21. For our microwave horn, the parameters are $r = 23.9$ mm, $L = 74.93$ mm and $\lambda = 2$ cm. This gives a beam waist of $w_0 = 1.5$ cm at a distance $\Delta = 2.5$ cm from the front

aperture of the horn. Thus, for the purposes of calculations for the ellipsoid mirror geometry, we will assume that the horn is emitting a Gaussian beam of waist w_0 and remember that finally we have to add (or subtract) the distance Δ to determine where the aperture of the horn should lie.

9.3.4.3 Theory and design of the ellipsoidal mirror

The concave surface of an ellipsoidal mirror is described by an ellipse rotated about its major axis i.e. $\frac{x^2+y^2}{b^2} + \frac{z^2}{a^2} = 1$ where $2a$ is the length of the major axis and $2b$ is the length of the minor axis. The ellipse is defined by two foci on the major axis, F_1 and F_2, such that for any point P on the ellipse, the lengths R_1 and R_2 corresponding to $\overline{F_1P}$ and $\overline{F_2P}$ satisfy $R_1 + R_2 = 2a$. An incident Gaussian beam passing through F_1 and having a radius of curvature R_1 will get reflected from the point P and will be transformed into an output beam passing through F_2 with radius of curvature R_2. The focal length of the mirror is given by $f = \frac{R_1 R_2}{R_1+R_2}$. Furthermore, by specifying the angle of incident θ_i, we can determine the eccentricity of the ellipse by

$$e = \frac{(R_1^2 + R_2^2 - 2R_1R_2\cos 2\theta_i)^{1/2}}{R_1 + R_2} \tag{9.26}$$

thus fully defining the ellipse (Figure 9.22).

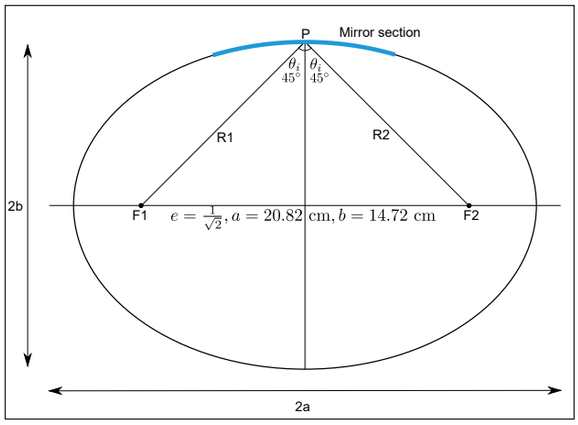

Figure 9.22: Geometry of ellipsoidal reflector. A Gaussian beam with radius of curvature R1 passing through the focus F1 gets reflected with radius of curvature R2 passing through the focus F2. The blue section corresponds to the part of the ellipse used as a mirror.

There are many physical constraints along the path of the microwaves into the chamber as shown in Figure 9.23. Going up from the ellipsoid mirror, we have the bottom MOT mirror (which is a 2" mirror but is at a 45° to the horizontal), the bottom MOT QWP, the window to the vacuum chamber (this is a 2.75 CF window), the long cylindrical protrusion on the vacuum flange, and finally the RF MOT coils. Because of the diffraction limit, the microwaves can only be focused down to a waist of $w_{0,min} = 1$ cm, but since the microwave beam is a gaussian beam, it will diverge very quickly after this waist, and a lot of the power will be lost on hitting the various objects in between the ellipsoidal mirror and the molecules. Thus, there is some optimization to be done as to where the waist should be located and the size of the waist.

The goal of the optimization process is to get the maximum Rabi frequency at the molecule position. This means that we want to get the maximum possible power with the least possible waist (since $\Omega \propto P/w^2$), given all the above constraints. For the purposes of determining the optimal width and location, these constraints can be modeled as a long cylinder of diameter 4.06 cm and length 14.76 cm (which encompasses the MOT mirror, the MOT QWP, the window and the

flange), and the RF MOT coils can be modeled as a cylinder with diameter 3 cm (corresponding to the opening on the RF MOT coil board). Then, for a given waist w and location z, we can calculate what fraction of power is transmitted at the beginning of the bottom cylinder and at the beginning of the RF-MOT cylinder. Fraction of power transmitted is estimated by taking the ratio of the area of the waist at each aperture to the area of that aperture i.e. $P_{lost} = \frac{w_{bottom}^2}{2.03^2} \times \frac{w_{RFMOT}^2}{1.5^2}$ and if the waist is less than the aperture radius, then all of the power is transmitted. We only care about how much power is transmitted at the beginning of each aperture, since once the microwaves are inside the apertures, they are just propagating in free space.

The parameter space for w and z is small enough that this can be done by brute force and we find that the most power with the smallest waist at the molecules occurs at a waist of $w_0 = 1.5$ cm and location $z = 13.9$ cm. In this case, the waist at the molecules is $w_{molecules} = 1.845$ cm, and the power transmitted through is $\sim 10\%$ of the case where all the power is transmitted.

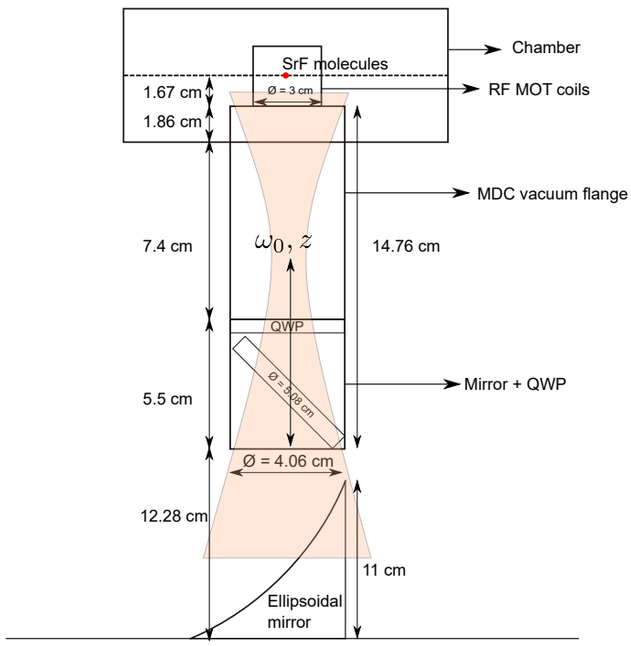

Figure 9.23: Geometric constraints along the microwave propagation path into the chamber. We can model the restraints as a long tube of diameter 4.06 cm which takes into account the open apertures from the mirror, the MOT QWP and the 6" vacuum flange. The RF MOT coils can be modeled as a tube of diameter 3 cm. Microwaves, with waist ω_0 and location z, are shown in beige.

From here, we can backtrack and calculate the ellipsoid parameters. Our microwave horn emits a beam with waist $w_0 = 1.5$ cm, thus the magnification needed is $M = \frac{w_{0,out}}{w_{0,in}} = 1$. To get the distance from the waist to the ellipsoidal mirror, we added 6.3 cm to the optimal z obtained before (6.3 cm is roughly half the distance from the bottom of the MOT mirror to the table.) Thus, the distance from the ellipsoidal mirror to the waist is around $z = 20.2$ cm and this gives $R_2 = z + \frac{z_c^2}{z} = 20.82$ cm. From ref [206], Eq 3.45, we see that this corresponds to a focal length of

$$f = z_c \frac{1 + (d_{in}/z_c)^2}{2d_{in}/z_c} = 10.41 \text{ cm} = \frac{R_2}{2} \tag{9.27}$$

Then, we have $R_1 = R_2 = 20.82$ cm, and for $2\theta_i = 90°$, the eccentricity of the ellipse is $e = \frac{1}{\sqrt{2}}$. With this, we know everything about the ellipse, and the final parameters for the ellipse are $a = 20.82$ cm and $b = 14.72$ cm. The waist for the microwave horn needs to be at a distance $z = 20.2$ cm from the mirror which means that the aperture of the microwave horn needs to be at a distance of 17.7 cm from the mirror.

The final optimized ellipsoidal mirror is shown in Figure 9.24. The mirror has grooves along the sides and at the back to allow for correct positioning. The vertical height of the mirror is around 11 cm, which allows for enough length to position the MOT mirror and QWP. Because of space restraints, the diameter of the mirror is around the same size as the waist at the mirror, which could cause some inefficiency in reflection.

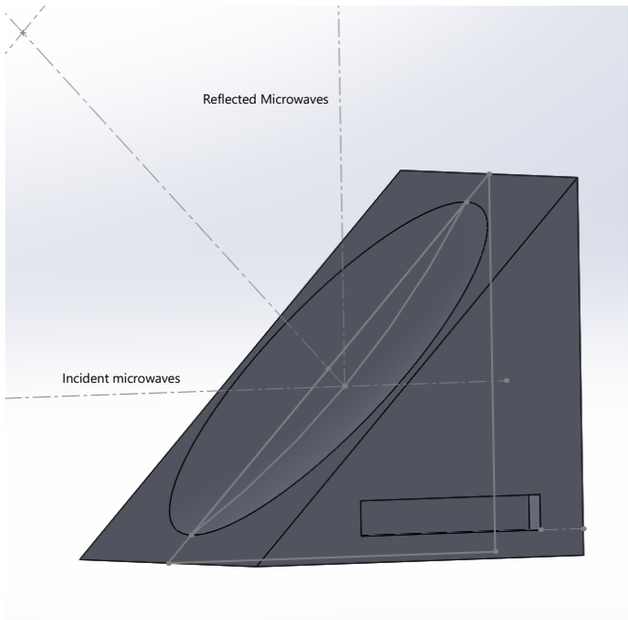

Figure 9.24: Optimized ellipsoidal mirror design with incident and reflected microwaves shown. The mirror will also have grooves along the sides and back such that it can be clamped to the table.

9.3.5 Need for π-polarization or electric field

We note that a recent demonstration of microwave shielding [44] used an additional π polarized microwave to get rid of so-called field linked resonances [208] that were causing enhanced three-body loss. These field linked resonances arise due to an attractive well in the long-range intermolecular potential through the microwave dressing. These occur if the Rabi frequency is sufficiently large, which is needed for shielding. While they offer a fantastic tool to control the scattering length of the molecules, at low temperatures, these states can be long lived, leading to three-body recombination, and thus extra loss. This was found to be the limiting problem in direct evaporation of molecules [184, 202].

One way to counteract this issue is by the use of π polarized microwaves. As we discussed before, this leads to worse shielding if the Rabi frequency is low. However, in the presence of these resonances, the π polarization serves to get rid of the attractive part at long range, while leaving the repulsive barrier mostly unaffected at short range. This was key to making the first molecular BEC [44]. This requires a π polarization Rabi frequency that is similar in magnitude to the σ^+ Rabi frequency. In our case, this will be hard to produce using another horn + focusing mirror set up as there is not enough space around the apparatus to accommodate this.

One potential solution is to turn to microwave patch antennas. An example of this is shown in Figure 9.25, where the copper tab on top serves as the antenna. The substrate material can be any dielectric that has low loss for our frequency (e.g. MACOR) and the desired resonant frequency

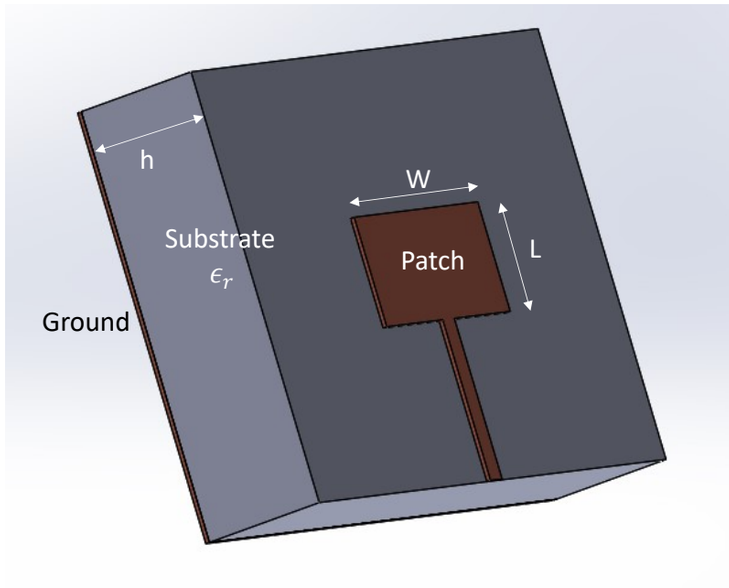

Figure 9.25: Microwave patch antenna to generate π polarization. The antenna can be placed in vacuum, and close to the molecules to generate high Rabi frequencies.

determines the size of the antenna by the following relations:

$$w = \frac{c}{2f_0\sqrt{\frac{\epsilon_R+1}{2}}} \quad ; \quad \epsilon_{eff} = \frac{\epsilon_R+1}{2} + \frac{\epsilon_R^{-1}}{2}\left[\frac{1}{\sqrt{1+12\left(\frac{h}{w}\right)}}\right] \tag{9.28}$$

$$l = \frac{c}{2f_0\sqrt{\epsilon_{eff}}} - 0.824h\left(\frac{(\epsilon_{eff}+0.3)\left(\frac{w}{h}+0.264\right)}{(\epsilon_{eff}-0.258)\left(\frac{w}{h}+0.8\right)}\right) \tag{9.29}$$

By using a MACOR substrate, this whole assembly can be made UHV compatible, and thus the antenna can be placed very close to the molecule cloud to generate high Rabi frequencies with π polarization.

The other possibility is to use a static electric field to cancel the field linked resonances. This was first proposed in [201] where they show that an optimal field value for this kind of shielding is when

$$\frac{dE_{dc}}{B} = 1 \tag{9.30}$$

where d is the dipole moment and B is the rotational constant. For SrF, $d \approx 3.47$ D, $B = 2\pi \times 7.5$ GHz, and thus we find that $E \approx 4$ kV/cm. This is not too large to be impractical, and this is something we are considering in the new apparatus.

9.3.6 Microwave testing

We have performed thorough testing of each of the required characteristics of the microwave setup, and these results will be presented in G. Zheng's thesis. Here, we briefly summarize the results. We have built the high power microwave setup and tested the power generation capabilities of the amplifiers. These work as expected - both the Qorvo amps can output around 47 dBm power.

We have also tested the phase noise characteristics by building a delay-line-discriminator and using a spectrum analyzer to determine the phase noise 10 MHz away from the carrier. We find that without the filters, the phase noise of the source roughly matches the phase noise of the generator. However, we see some pretty large spurs at intermittent frequencies, so we need to be careful to avoid those, see Figure 9.26. We also see that the band pass filters provide additional suppression of the phase noise, but not as much as would be predicted based on the datasheet.

Regardless, we can stack a few of these filters and increase the lifetime.

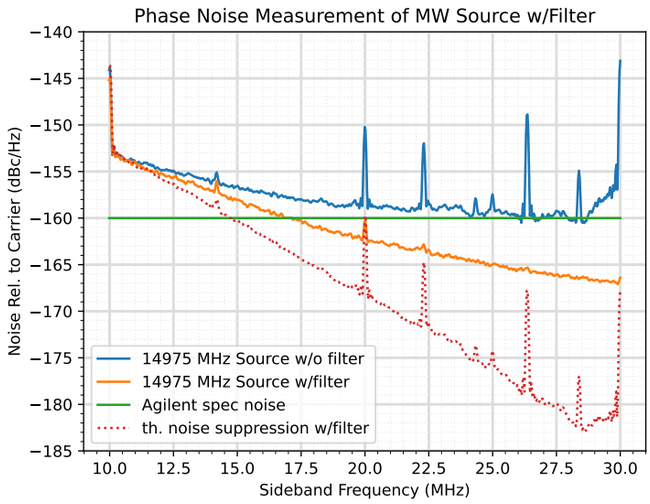

Figure 9.26: Phase noise measurement using a delay line discriminator. Far enough away from the carrier, the measured phase nosie matched the spec of the Agilent. However, we also see spurs at different frequencies. The phase noise with the bandpass filter is also shown. It suppresses the phase noise, but not as much as the datasheet predicts. Figure obtained from G. Zheng.

We have also measured the circularity of the output from the microwave horn by looking at the pickup from a rectangular waveguide held at 45° to the propagation direciton. In this case, the waveguide samples the polarizations along

$$x' = \frac{1}{\sqrt{2}}(x+y) \quad ; \quad y' = \frac{1}{\sqrt{2}}(x-y) \tag{9.31}$$

We can then vary the phase of one of the arms and map out the full polarization curve. In addition, we use the IQ mixer to directly determine the phase as well, and find that these measurements indicate we can achieve around 95% circularity.

We are currently addressing the redesign of the microwave mirror to focus the beam into the chamber. The old design was tested, and found to be too small, as we feared. There was significant power loss around that mirror. G. Zheng has redesigned the mirror to be around 3 times larger, with similar ellipse parameters, and it is under testing. The larger size will be paired with the new apparatus to enable easy microwave access.

With these improvements, we expect to reach $\Omega \approx 2\pi \times 40$ MHz in the chamber - a good starting point for shielding experiments.

9.4 New apparatus

The current apparatus is over 15 years old, and was not designed to implement the new directions that the experiment is undertaking. Q. Wang has designed an entirely new apparatus with a brand new chamber with more optical and microwave access, and the details of this will be documented in Q. Wang's thesis. The new apparatus is under construction currently. We briefly summarize the capabilities and expected gains here.

9.4.1 Transverse cooling, push beam and two-color MOTs

The new apparatus will have a flange dedicated to transverse cooling along both axes. This was proposed by T. Langin [110], where, with the new approach, it is expected to lead to large gains in molecule number. Such a transverse cooler was demonstrated for YbF molecules [209] where large gains were observed. We plan to follow the approach laid out in [110] where the two transverse cooling lasers polarizations are oriented parallel to the beam direction. This can lead to a factor of 10 gain in molecule number in simulations.

Another big improvement is expected by adding a push beam from the back of the slowing region. In the white light slower used currently, there is no spatial force dependence, and thus a significant portion of the molecules are stopped before they reach the MOT region. A simple and elegant solution is to add a push beam from the back, as simulated by T. Langin in [110], which would be parked close to the resonance at 0 velocity. This beam would guide the molecules that are stopped before the MOT into the trapping region. This is again expected to provide a factor of 10 gain in molecule number.

Finally, we also plan to implement a two-color MOT, where we use both the $X \rightarrow A$ and $X \rightarrow B$ transitions for the MOT, as simulated by T. Langin in [110]. The higher g-factor of the B state is expected to provide larger trapping forces, thereby increasing the capture velocity of the MOT and allowing us to capture even more molecules. The new apparatus will have all new viewports and waveplates that can accommodate both these wavelengths.

9.4.2 Three beam MOT

The power hungry nature of molecular MOTs caused us to use a single beam circulating around the experiment. This causes pretty large intensity imbalances, and also does not provide enough degrees of freedom to fine tune the alignment. However, with the new Precilasers systems, we now have over 2 W of power at both the $X \rightarrow A$ and $X \rightarrow B$ transitions. Furthermore, we have also built and tested double pass AOMs for the $X \rightarrow A$ transition that will allow us to independently address the 4 hyperfine levels instead of relying on EOMs. Together, this means we can now make a three beam MOT where each axis has its own independent beam that is passed and retro-reflected. We expect this will improve our instability issues, and provide more robust trapping.

9.4.3 Better microwave access and mirror

As shown in Figure 9.23, the current chamber has a lot of apertures that will hinder the delivery of the full power microwave power for the most effective shielding. In the new apparatus, the top flange will be a 8" CF viewport which will offer a roughly 6" clear aperture. In combination with a larger focusing mirror, we expect to deliver a large chunk of the 100 W power into the chamber. Ultimately, we hope to achieve upwards of $2\pi \times 40$ MHz Rabi frequencies for shielding.

9.4.4 Higher resolution imaging

As we saw in Chapter 5 our current imaging setup is unable to resolve the transverse axis of the ODT. We are designing a new imaging system that will magnify the ODT and allow us to resolve all axes. We are also purchasing a new and faster camera (Andor ZL41 Wave 5.5) that will allow us to take multiple images in the same experimental cycle. This will eliminate some of the systematic

issues we have with number fluctuations, as we can now normalize the signal in each cycle with the MOT.

9.4.5 Co-trapping Rb and sympathetic cooling

One of the big challenges in molecular evaporation is the low molecule numbers in the trap. A possible way to mitigate this is by co-trapping another abundant species (such as Rb) and using sympathetic cooling [29, 210]. We had prepared for this at Yale, where we demonstrated a RF and DC MOT of Rb [92]. We perform quick loading of Rb using a 2D MOT and trap around 1 billion atoms. The new apparatus will have a dedicated flange for the Rb 2D MOT so that we can co-trap Rb with SrF. Recent studies of Rb-CaF collisions [73] have shown that in the ground rotational state of the molecules, the inelastic collisional loss rate with atoms is very low. The elastic collision rate was not measured. Theoretical studies of Rb-CaF collisions show that there is a high probability of finding Feshbach resonances between Rb and CaF [211] which would allow for fine control of the scattering lengths. Furthermore, studies with assembled bialkali molecules have shown that the parent atoms can be used for sympathetic cooling of the molecules [212–214]. These results bode well for the possibility of sympathetic cooling of SrF with Rb.

The new chamber will have AR coated viewports and waveplates for Rb trapping as well. In addition, we plan to use a different camera mounted to a 1.33 CF mini flange to image the Rb and SrF MOTs such that it will be easy to co-align them.

A lot of exploring needs to be done here to determine if this is a viable route to sympathetic cooling or not. Theoretical studies of SrF + Rb collisions have shown that favorable ratios of the elastic to inelastic collision rates can be obtained for certain large magnetic fields [215]. The bare elastic rate, as well as the B-field dependence and the existence of any Feshbach resonances needs to be measured. Currently, the ability to generate 1000 G fields does not exist; however, as seen in [211], there is a high likelihood of finding resonances at substantially lower fields, which are within reach.

Furthermore, as seen in Figure 9.27, the AC Stark shift/trap depth for Rb is around 6 times higher than for SrF for a 1064 nm ODT. Thus, if we tried to evaporate Rb by lowering the trap depth, as is conventional, then we would end up preferentially evaporating SrF instead, and lose all the SrF. This can be mitigated by adding a laser that is blue detuned of the Rb transition, but red detuned of the SrF transition. This would be anti trapping for Rb and thus reduce the trap depth, while still providing trapping for SrF. This new laser has to lie between 663 nm ($X \rightarrow A$ for SrF) and 780 nm (D_2 line for Rb). Care must also be taken that this laser is far enough away from both transitions so that the off-resonant scattering rate is low to avoid excess heating of either species. Preliminary calculations indicate that a second laser at 700 nm roughly heats both SrF and Rb equally, and will also lower the trap depth sufficiently for Rb that it now lies above SrF thus making direct evaporation feasible. This laser has not yet been acquired.

9.5 Path to degeneracy

Once we implement microwave shielding, we would like to try direct evaporation of SrF molecules towards degeneracy. Here, we shall look at some estimates of what we may be able to achieve with shielding and other contraints, and trying to optimize evaporation. We shall mostly be following

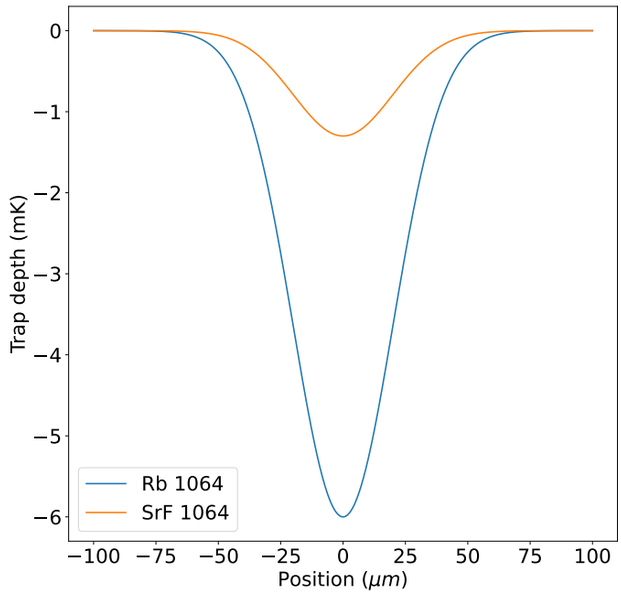

Figure 9.27: Trap depth comparison between SrF and Rb. For a 1064 nm trap, the Stark shift for Rb is around 6 times higher than that for SrF. This would preferentially evaporate SrF while leaving Rb in the trap.

the treatments from Refs [44, 216]. We treat the whole process classically, which is fine until we are in the neighborhood of degeneracy. We assume that $\eta = U/k_BT$, the ratio of the trap depth to the temperature, is kept constant during the evaporation. We further assume that we have independent control over the trap frequencies and trap depth, and so the dependence of the geometric mean trap frequency on the trap depth ($\bar{\omega} \propto U^{\nu}$) can be controlled with ν fixed during evaporation. Lower trap frequency results in lower density and evaporation rate but adiabatic expansion also leads to cooling. BEC is achieved in a 3D harmonic trap when the phase-space density (PSD) reaches 1.202. Thus, the optimization of η and ν to produce a BEC with maximum molecule number is our goal here.

As always, we treat our trap to be harmonic with the potential

$$U(\vec{r}) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2). \tag{9.32}$$

The particle number density obeys a boltzmann distribution

$$n(\vec{r}) = n_0 \exp\left(-\frac{U(\vec{r})}{k_BT}\right), \tag{9.33}$$

where $n_0 = N\bar{\omega}^3\sqrt{\frac{m}{2\pi k_BT}}^3$ is the peak density, $\bar{\omega} = (\omega_x\omega_y\omega_z)^{1/3}$ is the geometric mean trap frequency and N is total particle number. The energy density (kinetic and potential) distribution is

$$e(\vec{r}) = \left(\frac{3}{2}k_BT + U(\vec{r})\right) n(\vec{r}) \tag{9.34}$$

with the total energy $E = \int e(\vec{r})\mathrm{d}^3r = 3Nk_BT$. The phase space density is defined as $\rho = n_0\lambda_{dB}^3$ where $\lambda_{dB} = \frac{2\pi\hbar^2}{mk_BT}$ is the thermal de Broglie wavelength. The evaporation efficiency is defined as

$$\gamma_{eff} = -\frac{\ln(\rho_f/\rho_i)}{\ln(N_f/N_i)}, \tag{9.35}$$

with a good efficiency usually exceeding 1.

The evaportive cooling process works by preferentially kicking out particles with energy greater

than the trap depth and rethermalizing the rest via elastic scattering to lower the total energy. We need to consider the various cooling and heating mechanisms in this process. We model the number and temperature evolution as:

$$\dot{N} = \dot{N}_{ev} + \dot{N}_{ad} + \dot{N}_{1B} + \dot{N}_{2B} + \dot{N}_{3B} \tag{9.36}$$

$$\dot{E} = \dot{E}_{ev} + \dot{E}_{ad} + \dot{E}_{1B} + \dot{E}_{2B} + \dot{E}_{3B} \tag{9.37}$$

where the subcripts are: *ev* stands for evaporation, *ad* stands for adiabatic trap expansion, $1B(2B, 3B)$ iss the one-body (two-body, three-body) loss. Solving these coupled equations simultaneously will reveal the evaporation dynamics.

9.5.1 Simple model

First, we shall look at a simple model from [216] to gain intuition and to benchmark the simulation results. We model each term in eq. 9.36 as follows.

$\dot{N}_{ev} = -\Gamma_{ev} N$ where, for deep traps ($\eta > 6$), $\Gamma_{ev} = 2(\eta - 4)e^{-\eta}\Gamma_{el}$ and $\Gamma_{el} = n_0 \sigma \bar{v}/(2\sqrt{2})$ is the elastic collision rate, σ is the elastic collision cross section, and $\bar{v} = 4\sqrt{k_B T/\pi m}$ is the average relative velocity.

For deep traps $\dot{N}_{ad} \approx 0$ and we shall not include this in the calculations.

The other loss terms can be expressed as

$$\dot{N}_{1B} + \dot{N}_{2B} + \dot{N}_{3B} = -\Gamma_{1B} N - \Gamma_{2B} N - \Gamma_{3B} N, \tag{9.38}$$

where $\Gamma_{1B} = 1/\tau$ where τ is the combined one-body lifetime of the molecules, $\Gamma_{2B} = \beta_{2B} n_0/(2\sqrt{2})$ where β_{2B} is the two-body loss rate coefficient, and $\Gamma_{3B} = \beta_{3B} n_0^2/(3\sqrt{3})$ where β_{3B} is the three-body loss rate coefficient.

$\dot{E}_{ev} = -\Gamma_{ev}(\eta + \kappa)E$ where $\kappa = (\eta - 5)/(\eta - 4)$ and $\kappa k_B T$ is the average energy removed per evaporated molecule. The adiabatic expansion term is $\dot{E}_{ad} = \nu E \dot{T}/T$. The rest of the terms can be expressed as:

$$\dot{E}_{1B} + \dot{E}_{2B} + \dot{E}_{3B} = -\Gamma_{1B} E - \Gamma_{2B} \frac{3}{4} E - \Gamma_{3B} \frac{2}{3} E \tag{9.39}$$

We can now recast the master equation with these terms, and also convert E to T to find

$$\dot{N} = -(\Gamma_{ev} + \Gamma_{1B} + \Gamma_{2B} N + \Gamma_{3B}) N \tag{9.40}$$

$$\dot{T} = \frac{1}{3(1-\nu)} \left(-\Gamma_{ev}(\eta + \kappa - 3) + \Gamma_{2B} \frac{3}{4} + \Gamma_{3B} \right) T \tag{9.41}$$

where we now see that the two-body and three-body inelastic losses actually lead to heating. This is because these losses are density dependent and happen most often at the trap center, where the coldest particles reside. The hotter particles can travel to the outer edges of the trap, where these losses are not as likely. Since these losses happen most likely for the coldest molecules, they carry away less energy than the average energy, and this leads to heating. In our simulations, we will solve these for different η and ν to find optimum evaporation parameters.

We need to estimate these parameters for SrF molecules in our ODT. Some of these parameters are kept constant throughout all the simulations (unless stated otherwise), and they are summarized in Table 9.1. We borrow these values from the recently demonstrated shielding BEC

experiment [44, 202]. They demonstrated elastic cross sections as large as 10^{-8} cm^2, two body loss as low as $\beta_{2B} = 3 \times 10^{-13}$ cm^3 s^{-1} and almost no three-body loss. The one-body loss rate is set to the blackbody loss rate at room temperature [75].

Parameter	Value
σ	1×10^{-9} cm^2
Γ_{1B}	$1/3$ s^{-1}
β_{2B}	3×10^{-13} cm^3 s^{-1}
β_{3B}	0

Table 9.1: SrF parameters that are kept constant for evaporative cooling simulations.

9.5.1.1 Worst case for SrF

We have already demonstrated that we can load ODTs with $N_i = 4 \times 10^3$ at $T = 40\,\mu$K. Let us look at a scenario where we cannot increase/decrease the initial number/temperature in the ODT, and we are also stuck with the slow trap frequency along the axial direction. These parameters are listed in Table 9.2.

Parameter	Value
N_i	4×10^3
T_i	$40\,\mu$K
$\bar{\omega}_i$	$2\pi\times$ 500 Hz

Table 9.2: SrF parameters used for worst case evaporative cooling simulation using the simple model.

We can then iterate over η and ν to find the optimal parameters for forced evaporation, as shown in Figure 9.28. In this simulation, we let the integration run for a maximum of 10 s of evaporation (or until BEC is reached), and set the number to 0 if a PSD of 1 is not reached (or the value of N when BEC is reached). We see that for $\eta = 8$ and $\nu = 0$, we are left with a condensate of around 50 molecules at the end of the evaporation, which should be easily detectable in the experiment. The evaporation sequence takes a total of 2.5 s. This is already quite encouraging since this suggests that a BEC is within reach without improving any other parameters in the experiment. This is, of course, contingent on the shielding working as we expect.

9.5.1.2 Optimistic cases for SrF

We can, of course, be optimistic that the new slowing and trapping techniques will bear fruit and we can increase the ODT number by a factor of 10-100. Note that at some point, we will be limited by the Λ-light induced losses that we saw in Chapter 8, but we can get a factor of 10 gain if we implement shielding fast enough. We can also be optimistic that we will be able to implement a crossed beam ODT which would serve to greatly increase the trap frequency, thereby evaporating faster. Furthermore, we can also look at what would happen if the shielding is not as effective. We outline these results in Table 9.3, where we have explored a number of different combinations of starting number, starting temperature, the trap frequency and the two-body loss rate coefficient. This is, by no means, an exhaustive list, but is an indication that there are many configurations the experiment can try to produce a BEC. The most promising among them is the configuration

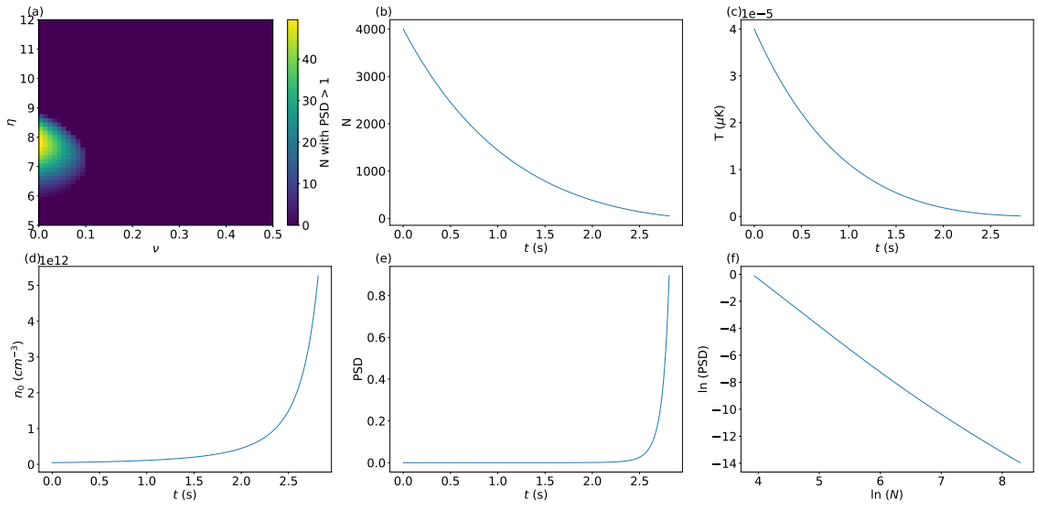

Figure 9.28: Results of the simulation using the simple model for worst case evaporation of SrF. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 2.5 s of evaporation.

where only the initial number is increased by a factor of 10, which would lead to large observable BECs.

Conditions	η	ν	N_{BEC}
$N_i = 4 \times 10^4$, $T_i = 40\,\mu$K, $\bar{\omega}_i = 2\pi \times 500$ Hz, $\beta_{2B} = 3 \times 10^{-13}$ cm^3 s^{-1}	10	0	4000
$N_i = 4 \times 10^3$, $T_i = 20\,\mu$K, $\bar{\omega}_i = 2\pi \times 500$ Hz, $\beta_{2B} = 3 \times 10^{-13}$ cm^3 s^{-1}	8	0	170
$N_i = 4 \times 10^4$, $T_i = 40\,\mu$K, $\bar{\omega}_i = 2\pi \times 500$ Hz, $\beta_{2B} = 10^{-12}$ cm^3 s^{-1}	9	0	2000
$N_i = 4 \times 10^4$, $T_i = 20\,\mu$K, $\bar{\omega}_i = 2\pi \times 500$ Hz, $\beta_{2B} = 5 \times 10^{-12}$ cm^3 s^{-1}	7	0	250
$N_i = 4 \times 10^3$, $T_i = 20\,\mu$K, $\bar{\omega}_i = 2\pi \times 1100$ Hz, $\beta_{2B} = 10^{-12}$ cm^3 s^{-1}	9	0	350
$N_i = 4 \times 10^4$, $T_i = 40\,\mu$K, $\bar{\omega}_i = 2\pi \times 1100$ Hz, $\beta_{2B} = 10^{-12}$ cm^3 s^{-1}	9	0	5000

Table 9.3: Optimistic evaporation cases for SrF. Other combinations of parameters, are, of course possible, and we list a few that would produce easily detectable condensates.

9.5.2 More complete evaporation models

The model from § 9.5.1 was a good starting point to develop intuition and it describes alkali gas evaporation very well, where the scattering lengths are low, leading to slow evaporation rates. However, in the case of microwave-dressed dipolar molecules, the scattering length can be quite large, and molecules can enter the so-called hydrodyniamic regime [43, 134, 202, 217] in which the elastic collision rate is so large that the mean free path becomes less than the size of the molecular cloud. The cloud is said to be collisionally thick, and in this case, the rethermalization rate is reduced from Γ_{el} to the geometric mean trap frequency $\bar{\omega}/(2\pi)$, causing the value of the elastic rate coefficient to saturate. This also makes the evaporative cooling efficiency worse, because the inelastic collision rate is not affected by the hydrodynamic regime, but the elastic collision rate is capped. Formally, the hydrodynamic regime occurs when $\Gamma_{el} = \frac{n_0 \sigma \bar{v}}{2\sqrt{2}} \gtrsim \frac{\bar{\omega}}{2\pi}$. Here, we look at two models that incorporate this effect.

9.5.2.1 Simple model with hydrodynamic limit

First, we implement a simple change to our model from § 9.5.1, where, during the simulation, we monitor the elastic collision rate Γ_{el} and $\bar{\omega}/(2\pi)$ and cap $\Gamma_{el} \leq \bar{\omega}/(2\pi)$. This is a simple handling

of entering the hydrodynamic regime, and the results for the worst case from § 9.5.1.1 are shown in Figure 9.29.

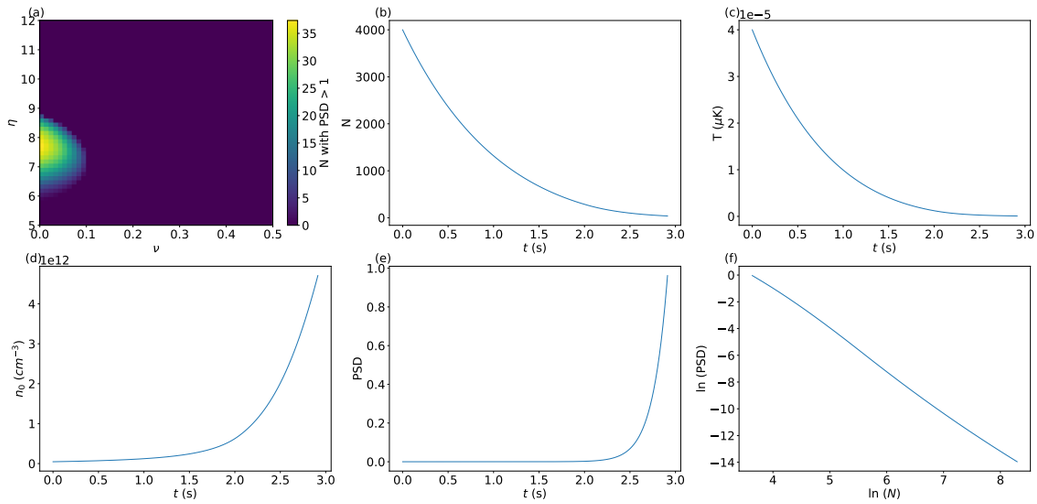

Figure 9.29: Results of the simulation using the simple model including the hydrodynamic regime (as described in the text) for worst case evaporation of SrF. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 3 s of evaporation, and one can see a small change in the efficiency in (f) as the PSD gets higher.

We can preliminarily see if we will enter the hydrodynamic regime by calculating the elastic collision rate, and for our initial conditions, we find that $\Gamma_{el} \approx 100$ Hz $< \bar{\omega}/(2\pi) = 500$ Hz. Thus, initially, we would not enter the hydrodynamic regime, and evaporation proceeds normally. However, as seen in Figure 9.29(f), for the ideal evaporation conditions, there is a change in slope of the ln (PSD) vs ln (N) curve at large PSD. This is the signature of entering the hydrodynamic regime, where the evaporation efficiency is decreased due to the capping of the elastic collision rate. We also see that it takes slightly longer to make a BEC in Figure 9.29 vs that in Figure 9.28, again due to the reduction of the evaporation efficiency. Even with this capping though, we are still able to achieve a BEC of around 37 molecules, which should be detectable in the experiment.

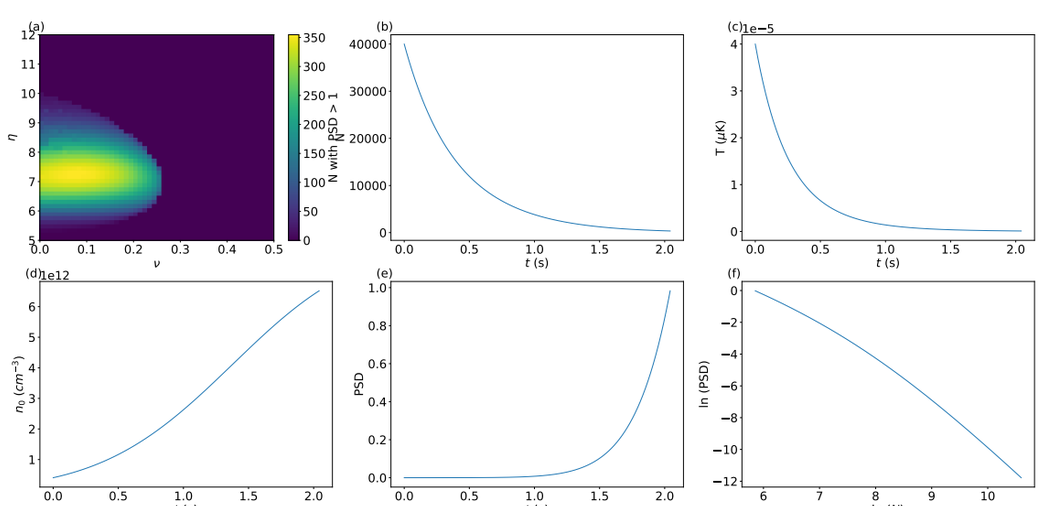

Figure 9.30: Results of the simulation using the simple model including the hydrodynamic regime (as described in the text) for a good case evaporation of SrF. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 2 s of evaporation, and one can see a substantial reduction in the evaporation efficiency in (f) as the PSD gets higher.

The effects of entering the hydrodynamic regime are most apparent in Figure 9.30 where we have now increased the starting number by a factor of 10, and kept all the other conditions exactly

the same, corresponding to one of the optimal cases we looked at in § 9.5.1.2. Here, we enter the hydrodynamic regime right from the beginning ($\Gamma_{el} \propto n_0 \propto N$) and the elastic collision rate is capped at the trap frequency during the entire evaporation sequence. We can see that for the optimal case in Figure 9.30, the evaporation efficiency keeps reducing as the PSD increases, and ultimately we are left with a BEC of only 350 molecules, as compared to the 4000 molecule prediction in § 9.5.1.2. Furthermore, the optimal values are $\eta = 7, \nu = 0.1$ here, vs the predicted $\eta = 1\nu = 0$ from § 9.5.1.2. This is still easily observable, giving us further optimism that a BEC is within reach.

9.5.2.2 General model with hydrodynamic limit

Finally, we look at a general model that was developed by [202] to describe dipolar collisions, which is relevant to us in the case of microwave shielding. The microwave dressing induces a dipole moment $d_{eff} = d_0\sqrt{12(1 + (\Delta/\Omega)^2)}$ where d_0 is the dipole moment in the molecule frame, Δ and Ω are the detuning and Rabi frequency of the microwave dressing. As the shielding parameter Δ/Ω is varied, the nature of the dipolar scattering changes - for large Δ/Ω, the collisions are semi-classical, but for small Δ/Ω, the collisional properties become strongly dependent on the dipole moment. Furthermore, the dipole-dipole interaction is anisotropic, which also plays a role in the rethermalization process. It changes the number of collisions needed to rethermalize, and the effective rethermalization rate can be written as $\Gamma_{ev} = \Gamma_{el}/N_{col}$ [218]. For far detuned microwaves, N_{col} converges to the bosonic s-channel value of 2.5, but as the microwaves are tuned to resonance, forward collisions that do not deflect the molecules' trajectories by a large angle get favored, and N_{col} sharply rises to as large as 15.

Furthermore, we now do not make any approximations about how large/small η is. We change the evaporation rate to $\Gamma_{ev} = \Gamma_{el}/N_{col}$. The number evolution is affected as per [202, 219, 220] $\dot{N}_{ev} = -N\nu_e(\eta)\Gamma_{ev}$ where $\nu_e(\eta) = (2 + 2\eta + \eta^2)/(2e^\eta)$. The energy evolution is affected as $\dot{E}_{ev} = -(1/3)E\alpha_e(\eta)\Gamma_{ev}$ where $\alpha(\eta) = (6 + 6\eta + 3\eta^2 + \eta^3)/(2e^\eta)$. In this case, the equations change to:

$$\dot{N} = -(\Gamma_{ev}\nu_t(\eta) + \Gamma_{1B} + \Gamma_{2B}N + \Gamma_{3B})N \tag{9.42}$$

$$\dot{T} = \frac{1}{3(1-\nu)}\left(-\Gamma_{ev}(\alpha_e(\eta) - 3\nu_e(\eta)) + \Gamma_{2B}\frac{3}{4} + \Gamma_{3B}\right)T \tag{9.43}$$

We also incorporate the hydrodynamic regime by capping $\Gamma_{el} \leq \bar{\omega}/(2\pi)$.

For our simulations, we need to be careful of the values we choose for each of these. For far detuned microwaves, while N_{col} is low, the inelastic collision rate is also higher, whereas for microwaves close to resonance, N_{col} is high, but the inelastic collision rate is lower. However, as [44] demonstrated, by adding a π-polarized component, one can effectively cancel out part of the large dipole-dipole interaction, making N_{col} small again, while also eliminating three-body losses. Here, let us assume that we can get this to work in a similar manner, and thus, the parameters stated in Table 9.1 are still valid. Nevertheless, we will look at two different $N_{col} = 2.5, 10$ to account for different experimental conditions.

9.5.2.2.1 Worst case with $N_{col} = 2.5$ We first look at the solution to the general model for the worst starting conditions from § 9.5.1.1, and with $N_{col} = 2.5$. The results are shown in

Figure 9.31, and the results are somewhat surprising. We see that we end up with around 120 molecules in the BEC, a lot more than the corresponding calculation in § 9.5.2.1.

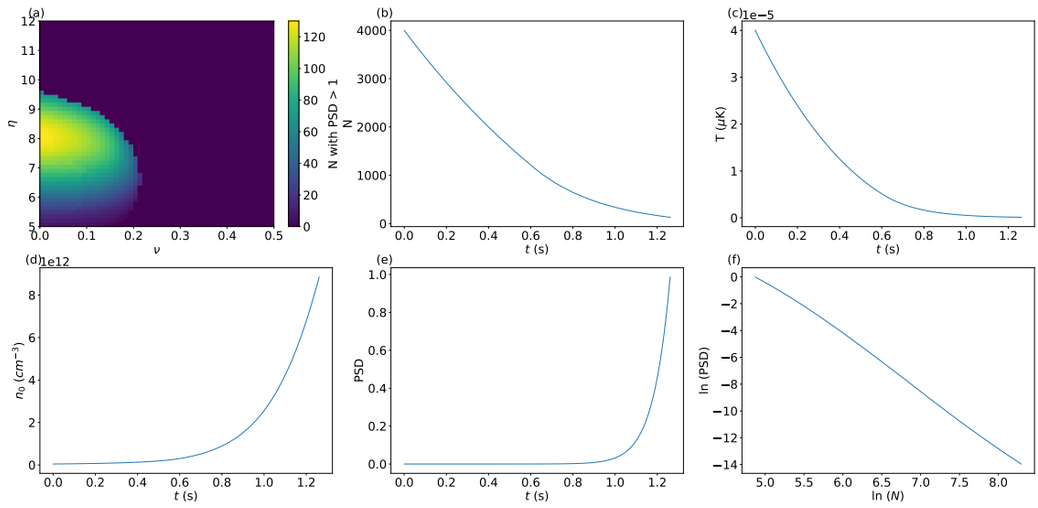

Figure 9.31: Results of the simulation using the general model including the hydrodynamic regime (as described in the text) for the worst case evaporation of SrF with $N_{col} = 2.5$. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 1.3 s of evaporation, and one can see a small change in the evaporation efficiency in (f) as the PSD gets higher.

This is because the effective $\Gamma_{ev} = \Gamma_{el}\frac{\nu(\eta)}{N_{col}}$ in the general model is larger than the $\Gamma_{ev} = \Gamma_{el}(\eta - 4)e^{-\eta}$ for low N_{col}. This means the effective elastic rate is higher and thus evaporation is more efficient, at least until we enter the hydrodynamic limit.

9.5.2.2.2 Worst case with $N_{col} = 10$ Next, we look at the results for the worst case initial conditions, but with $N_{col} = 10$, signifying that we do need to dress the molecules close to resonance to reduce inelastic losses. This means that the evaporation efficiency will be a lot worse, and indeed we find that under such a case, BEC is achievable for a small parameter space, as shown in Figure 9.32. Here, we will only get a 25 molecule BEC, maybe just about observable in the experiment. Thus it is clear that we would prefer to keep N_{col} as low as possible, and prefer to dress the molecules with somewhat far detuned microwaves.

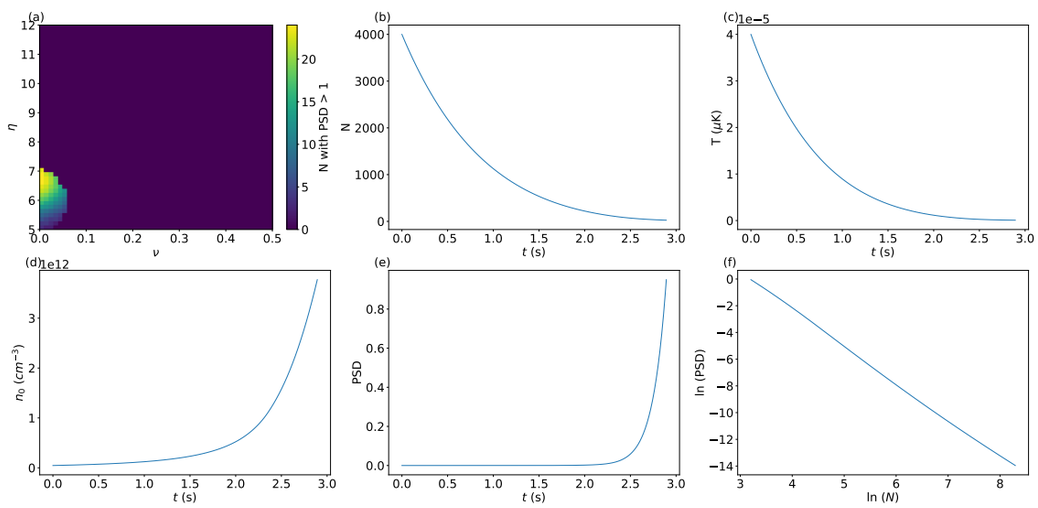

Figure 9.32: Results of the simulation using the general model including the hydrodynamic regime (as described in the text) for the worst case evaporation of SrF with $N_{col} = 10$. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 3 s of evaporation, and one can see a small change in the evaporation efficiency in (f) as the PSD gets higher.

9.5.2.2.3 Optimistic case with $N_{col} = 10$ Finally, let us look at what increasing the starting number by a factor of 10 gives us with a high N_{col}. This is shown in Figure 9.33 and here, we see that we can get a decently sized BEC of around 300 molecules. Furthermore, there is a drastic shift in the optimum $\eta = 6$, away from the large ηs from previous results. This is because, at large N_{col} it is preferred to keep the trap depth not as deep so as to reduce the inelastic loss rate. At larger trap depths, one would end up with a saturated elastic collision rate, but the inelastic rate just increases with density.

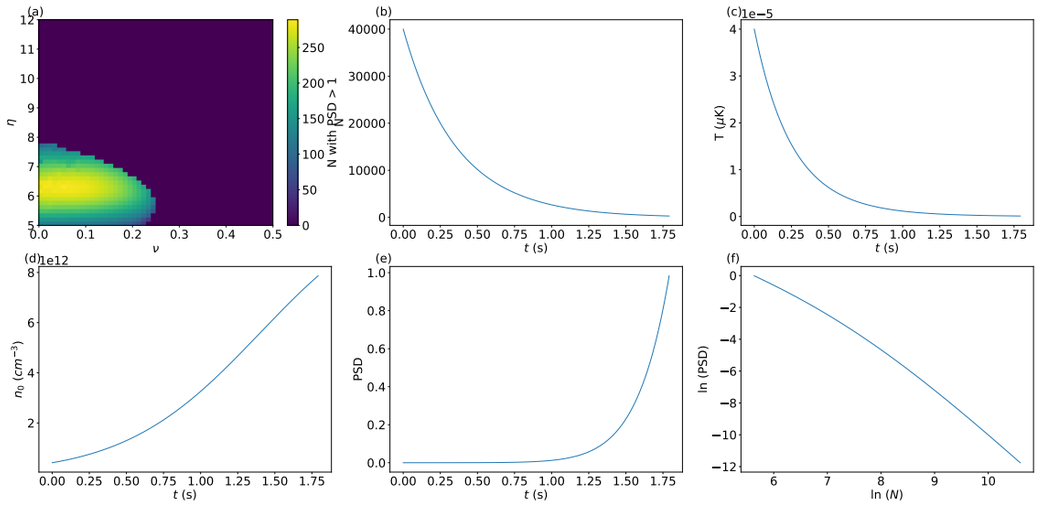

Figure 9.33: Results of the simulation using the general model including the hydrodynamic regime (as described in the text) for the worst case evaporation of SrF with $N_{col} = 10$. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 2 s of evaporation, and one can see a substantial reduction in the evaporation efficiency in (f) as the PSD gets higher.

These simulation results indicate that for a wide range of starting conditions, achieving a BEC is within reach if realistic elastic and inelastic collision rates can be achieved.

9.5.3 Evaporation and sympathetic cooling with Rb

In this section, we try to make some estimates of how well sympathetic cooling will work if we can co-trap SrF and Rb in the ODT. We follow a similar approach as from the previous section where we consider the number and temperature evolution of both Rb and SrF. To incorporate the cross-thermalization between the two species, we follow the approach outlined in [221]. The number and energy evolution of each species is governed by the following sets of equations:

$$\dot{N} = \dot{N}_{ev} + \dot{N}_{1B} + \dot{N}_{2B} + \dot{N}_{3B} \tag{9.44}$$

$$\dot{E} = \dot{E}_{ev} + \dot{E}_{ad} + \dot{E}_{1B} + \dot{E}_{2B} + \dot{E}_{3B} + \dot{E}_{th} \tag{9.45}$$

where $\dot{E}_{th}^{Rb} = -\dot{E}_{th}^{SrF} = \Gamma_{coll}\zeta k_B(T_{SrF} - T_{Rb})$, where $\zeta = 4m_{SrF}m_{Rb}/(m_{SrF} + m_{Rb})^2$ and the collision rate is given by

$$\Gamma_{coll} = 4\pi a_{SrF,Rb}^2 \bar{v}_{SrF,Rb} N_{Rb} N_{SrF} \left[\frac{2\pi k_B T_{SrF}}{m_{SrF}\bar{\omega}_{SrF}^2} + \frac{2\pi k_B T_{Rb}}{m_{Rb}\bar{\omega}_{Rb}^2}\right]^{-3/2} \tag{9.46}$$

where $a_{SrF,Rb}$ is the interspecies scattering length, and the relative thermal interspecies velocity is given by

$$\bar{v} = \sqrt{\frac{8k_B}{\pi}\left[\frac{T_{Rb}}{m_{Rb}} + \frac{T_{SrF}}{m_{SrF}}\right]} \tag{9.47}$$

We shall be making some assumptions about some of these quantities since we do not have any experimental data yet. The values we use for the different parameters are listed in Table 9.4. We assume that we can achieve the same conditions for SrF as used for the simulations in § 9.5.1.1. Furthermore, we assume that we can achieve good loading of Rb atoms in the trap, and at the same temperature as SrF. We use the known s-wave scattering lengths for Rb, and assume that by preparing Rb in the absolute ground state, we can eliminate two-body losses. We assume the SrF-Rb intraspecies scattering length to be $30a_0$ (a_0 is the bohr radius) based on some of the background scattering lengths in Ref.[211], and similar length scales outlined in [215]. We further assume that the two-body and three-body losses between Rb and SrF will be negligible. This is justified by Ref.[73], where, for molecules prepared in the rotational ground state, no two-body losses were observed. Lastly, we assume that we again have independent control over the two parameters η and ν, but we fix $\eta_{SrF} = 2\eta_{Rb}$ so that evaporation primarily occurs in Rb. We also fix $\nu_{SrF} = \nu_{Rb} = \nu$ for experimental simplicity.

SrF Parameter	Value	Rb Parameter	Value
σ_{SrF}	1×10^{-9} cm^2	σ_{Rb}	$8\pi(98a_0)^2$ [216]
$\Gamma_{1B,SrF}$	1/3 s^{-1}	$\Gamma_{1B,Rb}$	1/60 s^{-1}
$\beta_{2B,SrF}$	3×10^{-13} cm^3 s^{-1}	$\beta_{2B,Rb}$	0 [216]
$\beta_{3B,SrF}$	0	$\beta_{3B,Rb}$	4.3×10^{-29} cm^6 s^{-1} [216]
$N_{i,SrF}$	4×10^3	$N_{i,Rb}$	4×10^6
$T_{i,SrF}$	40 μK	$T_{i,Rb}$	40 μK
$\omega_{i,SrF}$	$2\pi \times 500$ Hz	$\omega_{i,Rb}$	$2\pi \times 500$ Hz
	SrF - Rb Parameter	Value	
	$a_{SrF,Rb}$	$30a_0$ [211]	
	$\Gamma_{2B,SrF,Rb}$ [73]	0	
	$\Gamma_{3B,SrF,Rb}$ [73]	0	
	η_{SrF}	$2\eta_{Rb}$	

Table 9.4: SrF and Rb parameters used for sympathetic cooling simulations.

We use the simple model developed in § 9.5.1 for the kinetics. Using the above assumptions, we find that the number and temperature evolution for Rb is given by

$$\dot{N}_{Rb} = -(\Gamma_{ev,Rb} + \Gamma_{1B,Rb} + \Gamma_{3B,Rb})N_{Rb} \tag{9.48}$$

$$\dot{T}_{Rb} = \frac{1}{3(1-\nu)}\left(-\Gamma_{ev,Rb}(\eta_{Rb} + \kappa_{Rb} - 3)T_{Rb} + \Gamma_{3B,Rb}T_{Rb} + \frac{\Gamma_{coll}}{N_{Rb}}(T_{SrF} - T_{Rb})\right) \tag{9.49}$$

and for SrF is given by:

$$\dot{N}_{SrF} = -(\Gamma_{ev,SrF} + \Gamma_{1B,SrF} + \Gamma_{2B,SrF})N_{SrF} \tag{9.50}$$

$$\dot{T}_{SrF} = \frac{1}{3(1-\nu)}\left(-\Gamma_{ev,SrF}(\eta_{SrF} + \kappa_{SrF} - 3)T_{SrF} + \frac{3}{4}\Gamma_{2B,SrF}T_{SrF} - \frac{\Gamma_{coll}}{N_{SrF}}(T_{SrF} - T_{Rb})\right) \tag{9.51}$$

We can once again numerically solve these and the results are shown in Figure 9.34. In these simulations, we cap the PSD above 1, as after that point, we only increase the fraction of atoms in the condensed core, but the simulation cannot handle this quantum behavior. Regardless of this, the results are optimistic. We see that Rb gets condensed much faster (within 0.5 s) than SrF, and up to the Rb condensation point, very few SrF molecules are lost. After this point, evaporation of Rb gets a little less efficient, and SrF is mainly evaporated until BEC is reached. We can also see the cross-thermalization scale, where the SrF temperature lags behind the Rb temperature almost for the entire evaporation sequence. Ultimately, within 1.2 s, SrF BEC can be achieved, with around 1400 molecules, indicating around 2/3 of the total molecules from the beginning are lost.

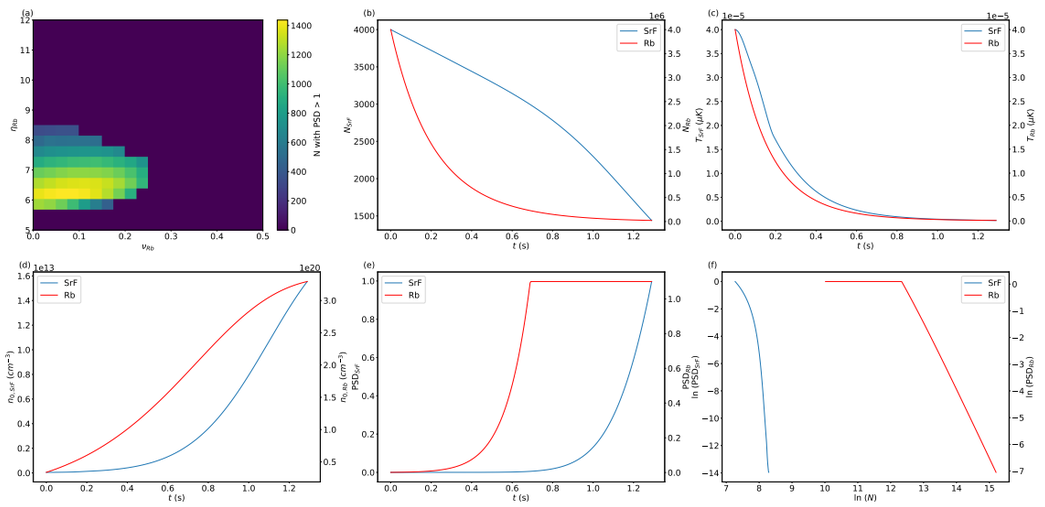

Figure 9.34: Results of the simulation using the simple model including the hydrodynamic regime for sympathetic cooling of SrF with Rb. Left axis shows SrF parameters, and right axis shows Rb parameters. PSD is capped at 1. (a) Numerical solution of the final number of atoms N_f to reach a final PSD of 1. (b) N vs t, (c) T vs t, (d) n_0 vs t, (e) PSD vs t, (f) ln(PSD) vs ln(N) for the optimum found in (a). BEC is achieved within 1.2 s of sympathetic cooling. Rb reaches degeneracy much faster (within 0.5 s) while SrF slowly thermalizes with Rb. SrF condensates with 1400 molecules can be achieved.

9.6 Conclusions

In this chapter, we have demonstrated that we can optically pump around 83% of the molecules in our trap into a single quantum state. Next, we have demonstrated preliminary coherent control by driving coherent Rabi oscillations between the rotational states in free space. We have also demonstrated high Ramsey coherence times of ms timescales, and current work is being pursued to try and increase the Rabi coherence times. Ultimately, we plan to transfer the molecules to the rovibrational ground state, where it is possible to shielding the molecules from inelastic loss using microwave shielding.

We have presented a detailed microwave generation and delivery scheme that should enable effective microwave shielding of SrF molecules. We have outlined ways to reduce the phase noise, as well as generate pure circular polarization, as needed. We have summarized the design of a next generation apparatus by Q. Wang, and current work is being done to characterize the components and assemble the machine. We have also briefly outlined the idea of implementing sympathetic cooling of SrF using Rb atoms, although this will require exploring a large parameter space, right from co-trapping SrF and Rb to characterizing collisions between atoms and molecules.

Finally, we have looked ahead towards direct evaporation of SrF and we have explored a variety of evaporation models. We have explored different starting conditions and different shielding criteria for these models, which indicate that direct evaporation of SrF can be efficient under certain trap conditions. We have also briefly explored sympathetic cooling of SrF with Rb, and find that this can be highly efficient under certain SrF-Rb collision and trap conditions. With these results, we believe that achieving a BEC of SrF molecules is within reach.

Chapter 10

Conclusions

This dissertation described progress on the SrF molecule cooling and trapping experiment, which ultimately aims to make a Bose-Einstein Condensate of SrF molecules.

In Chapter 4, we described the implementation of the Λ-enhanced gray molasses cooling protocol, and how this allowed us to discover an error in the branching ratios for the cooling transition, stemming from a sign error in the J-mixing coefficients. The resulting simulation results matched well with the experimental results, and we were able to cool the molecules down to the $\sim 10\,\mu$K regime.

In Chapter 5, we used the Λ-cooling method to actively load an ODT. We showed that due to the interplay of the AC Stark shifts caused by the ODT beam and the cooling beams, one can use the polarization of the ODT beam to enhance in-trap loading and cooling. Due to the peculiarities of our setup, we found that an intensity mismatch in the Λ-cooling beam intensities actually boosted the cooling - while this behavior was not replicable in OBE solvers, we were still able to achieve $\sim 14\,\mu$K temperatures in the ODT. This brought us to the brink of observing SrF molecule-molecule collisions for the first time, and we needed a factor of ~ 10 increase in density to be able to do so.

Chapter 6 discussed the design and implementation of a new molecule source based on chemical reactions between ablated Sr metal and SF_6 gas to make SrF molecules. This new source greatly enhanced the target lifetimes, while producing a similar trapped molecule number. We also discussed improvements to the cell geometry to get a colder and slower molecular beam.

In Chapter 7, we outlined a novel trapping technique called the blue-MOT, which produces a trapping force for blue-detuned light. By first capturing in the red-MOT, and then switching to the blue-MOT, we are able to achieve a factor of 100 increase in density, and compress the MOT to as low as $\sim 200\,\mu$m. The blue-MOT relies on the dual-frequency trapping mechanism, and is, in principle, a general technique that can be extended to other molecular species. We also detailed some of the peculiarities of the blue-MOT, such as it's sensitivity on the laser polarizations, and the deliberate misalignment of the MOT retro path to enhance trapping. The cause for this is still unclear, but a guess is that the MOT relies on sub-Doppler forces, which rely on robust dark states, the positions of which can potentially be tuned by changing the MOT beam alignment.

In Chapter 8, we detail the observation of collisions between SrF molecules in an ODT, the first such observation in a bulk gas of directly laser cooled molecules. This was enabled by the record high density loading ($\sim 30\%$ transfer efficiency) from the blue-MOT due to the small cloud size. We performed detailed statistical and systematic analysis of the loss date. We found that

the observed two-body loss rate is very close to the universal loss rate, akin to other experimental observations of molecular collisions.

Finally, in Chapter 9, we outline the ongoing work in the lab towards achieving quantum control, and shielding inelastic collisions. We describe an optical pumping protocol that allows us to achieve $\sim 80\%$ pumping efficiency in the $|N = 1, F = 0\rangle$ state in the ODT. We then describe the current progress and technical difficulties in transferring the molecules to the rovibrational ground state. We outline microwave shielding estimates, and describe the design of a low phase noise, high intensity microwave source. Lastly, we describe the pathways to degeneracy, including direct evaporation and sympathetic cooling with Rb.

With the estimates made in Chapter 9, we believe that achieving a BEC of SrF molecules in the next 2-3 years is within reach. Once a BEC is achieved, one may pursue a number of scientific milestones. Directly laser cooled molecules allow us to use non-destructive imaging in-situ, a crucial distinction from the bi-alkali approach, where, in each cycle, the molecule is dissociated into the individual atoms, and then detected. This is beneficial in many cases where one does not necessarily want to destroy a given run, such as, for example in quantum information applications with active error correction. In the future, one may imagine co-trapping Sr atoms (abundantly available in our new source, see Appendix B) and studying isotope substitution reactions, or even quantum interference in chemical reactions. One may also imagine the experiment going in a more precision measurement direction, where one may try to engineer spin-squeezing by using the strong dipole-dipole interaction, or even by constructing a cavity, and using cavity mediated interactions. Lastly, one may also build optical tweezers to harness these molecules for quantum information studies, as is being pursued in other groups. We believe that the SrF experiment has a bright future, with a lot of different interesting directions the experiment can take.

Appendix A

Sr target changing procedure

While direct exposure to Sr is not very harmful, Sr oxidizes very quickly in air to make SrO and SrOH, both of which can cause skin burns and respiratory injuries. Thus, there needs to be a robust target changing and storing protocol in place to avoid injuries, and also to avoid wasting useful Sr for the experiment.

Once the MOT signal has fallen by about a factor of 2, and cannot be recovered by any experimental optimization, it is time for a target replacement. The target change is best done on a Friday, so that there is enough time to pump down and start running on Monday.

A day before you plan to replace the target, you should prepare the new Sr target. This procedure is described below. Before beginning, make sure all the tools are cleaned with isopropanol or acetone.

1. Clean a new target holder piece by sonicating it in a solution of Citronox and DI water for 20 mins. After that, rinse it a few times with DI water, and sonicate again in DI water for 5 mins. Dry it out completely before proceeding.

2. Fill in 3 small beakers with hexane. These will be used to get rid of any oil on the Sr targets.

3. Wear P-95 masks. In a small metal weighing boat, weigh out the correct proportion of Stycast 2850ft and Cat 9 catalyst. Mix these thoroughly until a paste is formed.

4. Attach the Argon cylinder to the pump out station with the 1/4" plastic tube with Swagelok connections. The pump out station consists of a small turbo with a KF50 tee on top. One end of the KF50 tee has a Pirani gauge to monitor the pressure, and the other end has another tee. This second tee has a valve with the Argon connection on one end, and a window on the other end to view the target. Make sure the turbo is off, and the window flange is disconnected, so that the target can be placed inside.

5. Coat the center portion of the target holder piece with the epoxy paste. Stycast does not stick very well to copper until it is cured, so try to get as good a layer as you can.

6. Once everything is assembled, you are ready to open the Sr target. Things move pretty quickly after this. Start flushing the pump out chamber with Argon at a reasonably high flow rate. You should feel a positive pressure coming out of the open flange.

7. Open the can containing the Sr target. This has another plastic container with the Sr in it. Cut open the plastic and immediately put the Sr in the first hexane beaker. As long as the

Sr is under hexane, it won't oxidize.

8. Thoroughly swirl the Sr target in the first beaker. Then, use pliers to transfer the target to the second container. Again, swirl it and transfer to the third beaker. At this point, all of the oil should be removed from the target.

9. Remove the target from the hexane, and use kimwipes to immediately wipe it down (2-3 seconds). At this point, you may also have a second person blow nitrogen on the target.

10. Place the target on the copper piece, and try to smush it in the epoxy as hard as you can. Make sure to move it around so that the epoxy gets under and around the target. All this while, keep holding a kimwipe on top of the target so that oxygen does not get to the surface. Do this for around 10-15 seconds, until you feel confident that the epoxy has been well applied.

11. Immediately place the target inside the flange, valve off the Argon, and close the flange with the window. At the same time, turn on the turbo. The pressure should start falling quickly. You may also see some bubbles forming on the epoxy. This is normal, and is just the air escaping from under the epoxy.

12. The epoxy takes 12-24 hours to cure at room temperature. Keep the extra epoxy in the fume hood, so that you can make sure it is cured the next day. The target is now ready to be put in the cell.

Once the target is ready, you may proceed with the rest of the target change. Start with turning off the pulse tube, closing the gate valve, and setting the "Auto Pulse Tube On" button to the off position. Once the pulse tube is above 170 K, close the Edwards valve between the source and backing turbos, and turn off the source turbo. The source region is now fully isolated from anything else. Do this the night before so that the source region warms up, however, at this point, you may also flow in a small amount of Helium (1-10 Torr) to speed the warm up.

Much of the source opening/closing procedure is the same as that described in [93]. I will briefly summarize and add any changes below:

1. Once the source region is warmed up to 285 K, start flowing in nitrogen. This is to prevent the existing target from oxidizing. Once nitrogen is flowing, you can open the air valve slowly to bring the source up to atmosphere.

2. Remove the back and bottom plates on the 300 K shield. Remove the bottom plate on the 30 K shield. Remove the aluminum tape on the back plate of the 30 K shield, and then remove the back plate.

3. Remove the back plate on the 4 K shield.

4. Remove the target piece, and immediately dump it in some mineral oil. Stop flowing nitrogen. Remove the snorkel and remove the ablation window from the snorkel.

5. Remove the back plates and diffuser.

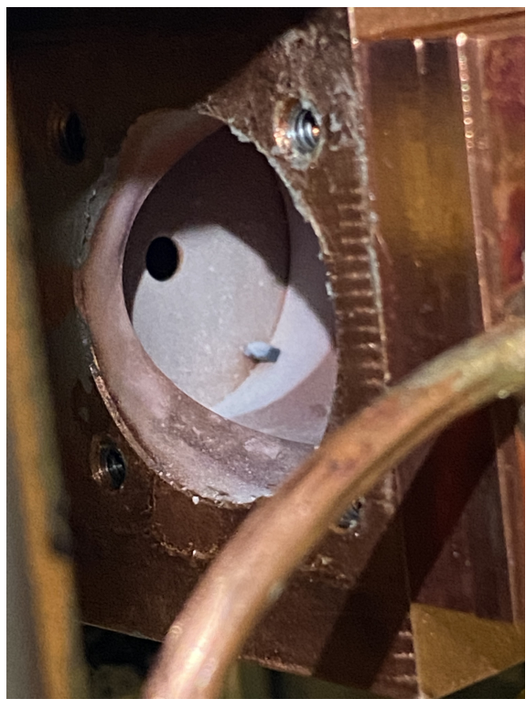

Figure A.1: Dusty cell after use. The inside of the cell looks gray, and is coated with dust on opening for a target change. You may even see small chunks inside the cell, presumably from ablation.

6. Clean any non-window parts with the snorkel by sonicating in Citronox and DI water. Clean any window parts with kimwipes and isopropanol.

7. Thoroughly clean the inside of the cell using kimwipes, DI water and isopropanol.

8. Clean all the windows with methanol using standard optics cleaning methods.

9. Use new screws which are cleaned with acetone. Clean the washers with acetone.

10. Scrape off as much Indium as possible on the back plate of the 4 K shield and replace it with new Indium strips. Use new Indium wire to seal the ablation window to the snorkel.

11. Once everything is clean, you can begin the closing up process. Apply a thin layer of Apiezon grease on each of the back plates, and secure them to the back of the cell.

12. Apply a think layer of Apiezon on the snorkel and screw it into the side of the cell. The top corner screw is tricky to get in, and I find it easiest to first get one of the bottom screws in, and then get this top one in. Be careful not to bump the SF6 tube in this process.

13. Apply Apiezon on the cell where the target piece goes. At this point, start flowing nitrogen again. Turn off the pump out turbo, quickly open the chamber, and bring the target over to the source region and screw it in. At this point, you should feel a positive pressure of nitrogen from the front hole of the cell.

14. Work quickly from here on out, and put in the back plate of the 4 K shield, the back plate of the 30 K shield, the aluminum tape around the corner of the 30 K shield, and the bottom plate of the 30 K shield.

15. During this process, close the Lesker valve between the UHV turbos and the backing region, and turn off all the turbos. Now the UHV is completely isolated from everything.

16. Screw in the back and bottom plates of the 300 K shields.

17. Once the backing turbo has spun down, open the Edwards valve between the source source and backing turbos slowly, so that the backing is exposed to atmosphere. Then you can turn on the source and backing turbos and start pumping down. Remember to stop flowing nitrogen, and close the air valve now.

18. The turbos take around 20-30 minutes to spin up and may fail 3-4 times during this process. Each time they fail, just restart them. The turbos hit a lull when the pressure gets down to 10 Torr in the backing region, and stay there for some time, slowly spinning up, until the source region gets below 10^{-3} Torr.

19. Once the turbos are spun up, the backing pressure should be around 2×10^{-2} Torr. Turn on the UHV turbos and once they are spun up, open the Lesker valve. At this point, only the gate valve should be closed.

Once the source ion gauge reads around 10^{-6} Torr, do a leak check, and if there are no leaks, the gate valve can be opened and the pulse tube turned on. The target may have developed a white layer of oxide on the surface during all this, but that will be blown away after a few ablation shots. As long as the target feels solid while replacing, it should be fine.

A few notes: if the turbos fail to spin up entirely, then there may be a few things wrong. First thing to check is whether all the screws on the 300 K shields are tight and there are no leaks. In a similar vein, you may check all the O-ring seals, and make sure they are not damaged. If that does not fix it, there may be something wrong with the diaphragm pump. The diaphragms need replacing every few years, so try replacing this next. This failure mode often shows up as the backing pressure falling too quickly, and then stalling. Finally, we have also seen the backing turbo itself fail. This is very rare, but can happen if there is oil in the system for some reason. It can be diagnosed by closing the Edwards, Lesker and gate valves, and trying to only pump on the backing region. If the turbo needs high power to do maintain vacuum, then it has failed. This is a more serious problem, and the turbo should be sent back to be refurbished.

Appendix B

Sr source characterization

B.1 Introduction

On switching to the Sr + SF_6 source in Chapter 6 we also performed some Sr characterization to detect the number of free Sr atoms produced per shot. This is important information if we want to simultaneously make a Sr MOT for instance, or even to see if the Sr coming out is thermalized with the Helium. The measurements presented here were performed mainly by M.Hu, in addition with G. Zheng, T. Langin and V. Jorapur.

B.2 Absorption measurements outside the cell

We measure the Sr atom number by looking at the absorption of a laser close to the $\left|^1S_0\right\rangle \rightarrow \left|^1P_1\right\rangle$ transition. The Beer-Lambert law relates the intensity transmitted through a optically thick cloud to the incident intensity as:

$$\frac{I}{I_0} = e^{-n\sigma L_s} \tag{B.1}$$

where I_0 is the incident beam intensity, σ is the absorption cross-section of the atoms, and L_s is the path length of the beam after the sample. The quantity $n\sigma L_s$ is also referred to as the optical depth (OD). This assumes that the beam is small enough to completely cover the transverse extent of the cloud. We then find the number density, and thus the number of atoms is:

$$N = nV = \frac{V}{\sigma L_s}\ln\frac{I_0}{I} \tag{B.2}$$

In our experimental setup, our "atomic cloud" is actually a pulse of Sr atoms that is transient in time. We place our absorption probe beam immediately after the cell exit aperture such that it propagates transverse to the atomic beam direction. We then measure the transmitted light power as the pulse of atoms flies by. The total number of atoms is simply the time integral of the absorption over the pulse.

The absorption cross-section is a function of detuning and since the molecules have a non-negligible transverse velocity, the Doppler broadened absorption is given by:

$$\sigma_\perp(v_\perp, \Delta) = \frac{\sigma_0}{1 + \frac{1}{\Gamma^2}(\Delta + kv_\perp)^2} \tag{B.3}$$

For our transition of interest, $\Gamma = 2\pi \times 32$ MHz, $\sigma_0 = 3\lambda^2/2\pi$. The transverse velocities follow a

gaussian distribution centered around 0 velocity, with standard deviation $\sigma_{v_\perp}$. To get the cross-section for all transverse velocities for a certain detuning, we integrate over the distribution as:

$$\sigma_\perp(\Delta) = \frac{1}{\sqrt{2\pi}\sigma_{v_\perp}} \int_{-\infty}^{\infty} e^{-\frac{v^2}{2\sigma_{v_\perp}^2}} \frac{\sigma_0}{1 + \frac{1}{\Gamma^2}(\Delta + kv_\perp)^2} dv_\perp \tag{B.4}$$

Finally, we can obtain N_{pulse} by integrating over the pulse duration:

$$N_{pulse} = \frac{Av_L}{\sigma_\perp(\Delta)L_s} \int_0^T \ln\frac{P_0}{P} dt \tag{B.5}$$

where A is the transverse area of the laser and v_L is the longitudinal velocity of the atom beam. For this characterization, we assume that the Sr atom velocities etc are similar to the SrF beam, and so we use the following parameters:

- $v_L \approx 140$ m/s
- $\sigma_{v_\perp} \approx 32$ m/s
- $L_s \approx 3$ mm
- $A = \pi L_s^2/4$
- $\lambda \approx 460.8623665$ nm
- $\Gamma = 2\pi \times 32$ MHz

B.3 Experimental setup and results

The transition of interest is at 461 nm, and so we decided to build a cats-eye ECDL using a commercial diode. M. Hu built this and found that we could lock it to a transfer cavity with a HeNe laser with a RMS jitter of around 1.5 - 2 MHz. Furthermore, on implementing a digital feedforward whereby the current was modulated as the piezo voltage was changed, we were able to obtain around 1.8 GHz (~ 3 FSR) mode-hop-free tuning range. This was found to be significantly more stable than a previous iteration that used a Littrow configuration. Using this, we were able to scan around 3 GHz total around the resonance of the Sr transition. The output of this laser is fiber coupled and sent along the same path as the SrF absorption beam. The beam passes in front of the cell, and is then detected on a photo-diode on the other side.

The data obtained is photo-diode voltage vs time. We first do some background subtraction by averaging the points where there are no atoms and subtracting the offset. Then, the total absorption may be determined by the integrated photo-diode voltage. The data run was repeated twice at each frequency and averaged.

We first look at the doppler-broadened absorption spectrum of the Sr beam. For this, we plot the peak OD and its standard deviation as a function of the detuning. This is shown in Figure B.1.

The data can be fit to a Gaussian with a full-width half-max of 295 MHz. This can then be translated to a temperature using the relation:

$$\Delta f_{FWHM} = \sqrt{\frac{8kT\ln 2}{mc^2}} f_0 \tag{B.6}$$

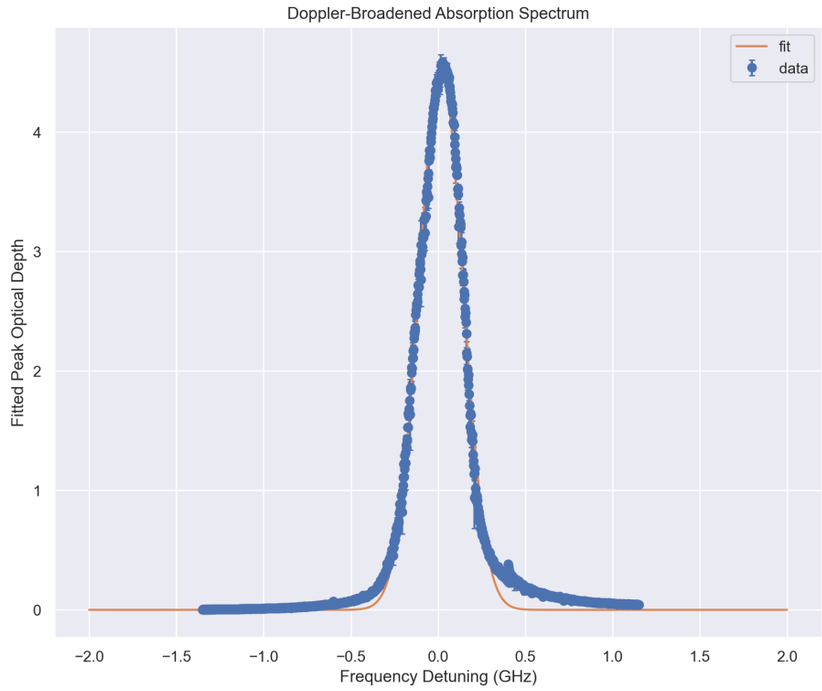

Figure B.1: Doppler-broadened absorption spectrum of Sr atoms on the Sr $^1S_0 \rightarrow ^1 P_1$ transition.

This yields a temperature $T \approx 35$ K, much hotter than the cell temperature of 4 K. However, it is also important to note that near resonance, there is basically 100% absorption, and it is very hard to determine the OD exactly. This can be seen from Figure B.1 where the peak OD near resonance is around 4.5, i.e. only around 0.01% of the light is transmitted.

This discrepancy can also be seen in the extracted raw number shown in Figure B.2. The Sr number exiting the cell should be independent of the detuning, however, we see a clear dependence on detuning, which is very pronounced near the resonance. This is again due to there being basically 100% absorption, and thus there arise significant errors in the determination of the true OD. There is also a number discrepancy at far detunings on either side of resonance, where the number is constant on either side, but not consistent with each other, also probably due to errors in OD determination.

Regardless of these uncertainties, we find that the average number of Sr atoms per pulse over all detunings is around $N_{Sr,pulse} \approx 6 \times 10^{12}$, with standard deviation $\sigma = 3 \times 10^{12}$ and standard error in the mean $\sigma/\sqrt{N} = 9 \times 10^{10}$.

If we ignore the data near resonance that has an obvious discrepancy, we find that the average number $N_{Sr,pulse,far-detuned} \approx 5.1 \times 10^{12}$ with $\sigma = 3 \times 10^{12}$ and $\sigma/\sqrt{N} = 1.1 \times 10^{11}$, which probably a better estimate of the actual number.

We also plot the OD vs time and detuning in Figure B.3. This gives us a transverse temperature $T \approx 35$ K.

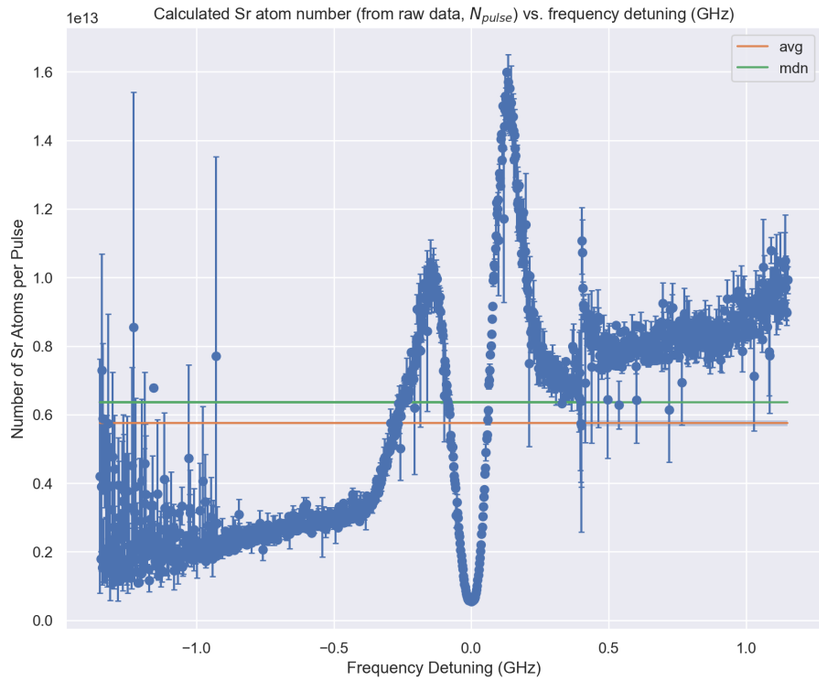

Figure B.2: Sr number extracted from the integrated OD. Blue points denote measured data. Orange (green) line denotes the average (median) number across all detunings.

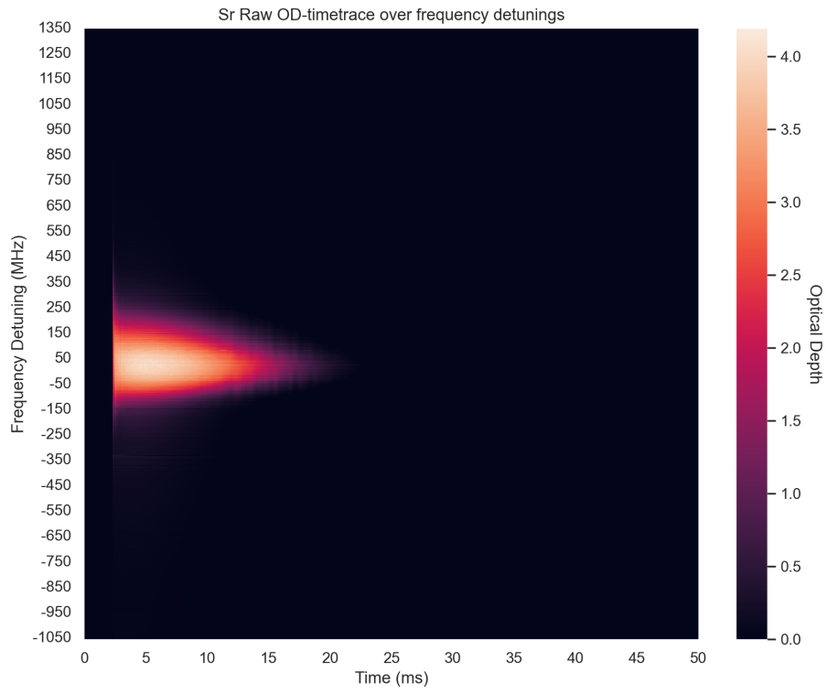

Figure B.3: Sr OD as a function of time on exiting the cell.

Appendix C

Experiment control and sequence

C.1 Experiment control

Before moving the experiment to UChicago, the experiment control was a hodgepodge of different systems working together. There were multiple delay generators talking to the same computer, and a master LabView program that controlled the experiment except that it did not fully control the experiment as we also had to manually change the timings on the delay generators. The LabView code's back panel could not be seen on the computer, thus it was impossible to edit the software. We will not describe this in detail as we have since upgraded to a better system. We refer the brave reader to [75, 76, 93] for further details about this.

The current experiment control is based on a python script written by Q. Wang and adapted from the Centrex control software. We have a number of devices that need to be configured and triggered each cycle - the analog waveforms for controlling the laser intensity, the field gradient, and the shim fields are controlled by NI DAQs. The timing control is accomplished by a SpinCore PulseBlaster module that triggers individual devices as well as any RF/MW switches at the appropriate time. The main program configures the waveform for each device, as well as the timing waveform for the PulseBlaster, and uploads it at the beginning of each cycle, see Figure C.1.

The experiment is triggered at the lowest point of the pulse tube temperature cycle by monitoring the temperature of a silicon temperature sensing diode. The output of the diode is purified and connected to a SR560 low noise amplifier, which under normal conditions is saturated. When the experiment start button is hit, the amplifier is blanked each cycle with a DAQ output causing the amplifier to send through the sinusoidal waveform of the temperature. At the lowest point in this cycle, a Rigol arbitrary function generator is triggered, which then produces a TTL pulse after a variable delay. This TTL pulse then triggers the PulseBlaster, and the experiment is set in motion.

The PulseBlaster is the timing control of the experiment - it sends out triggers to all the devices and switches - thus everything is referenced to the same start time (modulo the differences in the internal clocks of each device, which is not a problem for most purposes) making the experiment repeatable. The DAQs are all triggered at the start of the experiment, and the DAQ outputs are controlled using the analog waveforms created in the beginning. Once the PulseBlaster waveform is complete, it enters a WAIT mode where it waits to receive the next trigger. This makes it very easy to run the experiment at any arbitrary multiple of the pulse tube cycle (1.4 Hz), since any waveform longer than this time will not be triggered until after an additional cycle has passed.

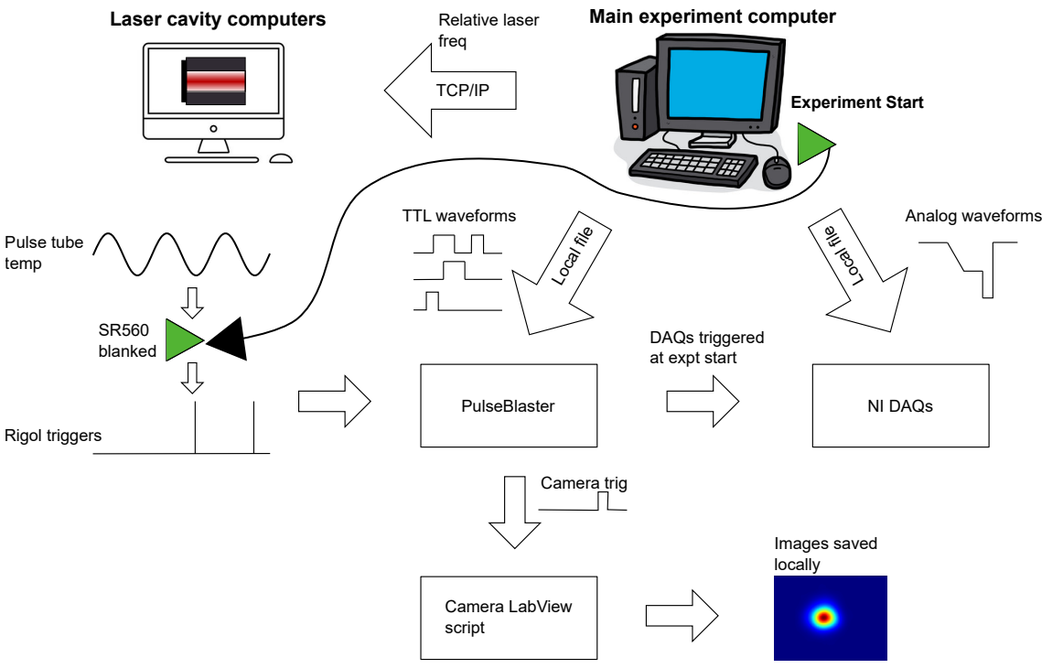

Figure C.1: Experiment control block diagram. The main experiment computer houses both the main python program as well as the camera LabView script. This computer communicates with the laser locking computers over TCP/IP. When the start button is hit on the main program, the SR560 is blanked each cycle and the Rigol starts sending triggers to the PulseBlaster.

Figure C.4 shows an example of a MOT waveform implemented from the PulseBlaster. Analog waveforms and the TTL timings are are configured in the PulseBlaster tab. When the start button is hit, the waveform is uploaded to the PulseBlaster and the devices, the SR560 starts blanking, and the experiment runs. At the end of each cycle, a new waveform is uploaded to the PulseBlaster and all the devices, and the experiment keeps running until either the stop button is hit, or it has finished the data run.

The main program allows for easy and arbitrary scanning of different parameters, and all the configuration data is saved locally to reference later for analysis. The program is modular - any new device can be easily added with all the functionalities of the other devices - making it a highly versatile program. Figure C.2 shows an example of the devices tab in the main program. Each device added to the experiment has its own tab where one can specify the device parameters such as the COM port, sampling rate, trigger inputs etc. or in cases of devices like the Agilent, these parameters could be the default frequency and power. The main program also has a tab for devices that are used for slow monitoring of parameters, see Figure C.3. Only those devices that are critical should be put here. For instance, the pressure in the source gauge needs to be in a certain range for the experiment to run - there is an interlock for this that prevents the experiment from running otherwise. Similarly, the MOT coil temperatures need to be below a certain value. For more details, see the group github [1].

The camera is still controlled by a LabView script - we were unable to find python packages for the camera. This makes it somewhat limited in what it can achieve - however, for all the results presented in this thesis, it was sufficient. In the future, a newer camera that has python packages will be used that will make the experiment control entirely in python and make the data acquisition fully arbitrary.

[1] `https://github.com/SrFDeMilleGroup/SrF-lab-control-v2`

C.2 Experimental sequence

We detail the experimental sequence up to the MOT, since that is the starting point of all the experiments. See Figure C.4 for a detail of all the devices that are used to make the MOT.

1. **t = 0 ms**: The PulseBlaster is triggered at the minimum of the pulse tube cycle. This sets the experiment in motion. All the DAQs are also triggered at this point and start outputting their analog waveforms.

2. **t = 2 ms**: The YAG is triggered. SrF molecules are produced by laser ablation of either a SrF_2 target or a Sr + SF_6 target.

3. **Laser slowing (0 - t_{slow})**:

   (a) **t = 0 - 35 ms**: $X \to A$ slowing is applied for a total of 35 ms due to the low scattering rate from the Λ-system. After 32 ms, the slowing shutter is turned off while the repump light is allowed to pass through to repump all molecules to $v = 1$. After 35 ms, the longitudinal shutter is turned off which blocks all the slowing light. Slowing coils are turned on for the entire duration of slowing.

   (b) **t = 0 - 17 ms**: $X \to B$ slowing is applied for a total of 17 ms, taking advantage of the higher scattering rate due to the breaking of the Λ-system. Slowing coils are turned on as well. After 14.5 ms, the slowing shutter is turned off while the repump light is allowed to pass through to repump all molecules to $v = 1$. After 17 ms, the longitudinal shutter is turned off which blocks all the slowing light. Slowing coils are turned on for the entire duration of slowing.

4. **MOT loading (t_{slow} - t_{slow} + 35 ms = t_{MOT})**: Molecules are accumulated in the MOT over 35 ms. The MOT coils and the MOT light are turned on at the beginning of the experiment and the MOT is accumulating molecules throughout - however, the most significant accumulation happens in these 35 ms. The DC MOT takes a similar amount of time to load as the RF MOT.

5. **MOT compression (t_{MOT} - t_{MOT} + 40 ms = t_{cMOT})**: After molecules are accumulated in the MOT, the MOT is compressed by linearly increasing the field gradient and lowering the laser intensity over 40 ms. This lowers both the temperature and the size of the molecules, thereby increasing density, see [55, 76, 93].

6. **Hold (t_{cMOT} - t_{cMOT} + 20 ms)**: The molecules are held in the compressed MOT for 20 ms to allow the dynamics to equilibrate.

The compressed MOT is the starting point of all our experiments, and we shall refer to this compressed MOT as the MOT in this thesis.

Figure C.2: Example devices tab in the main program.

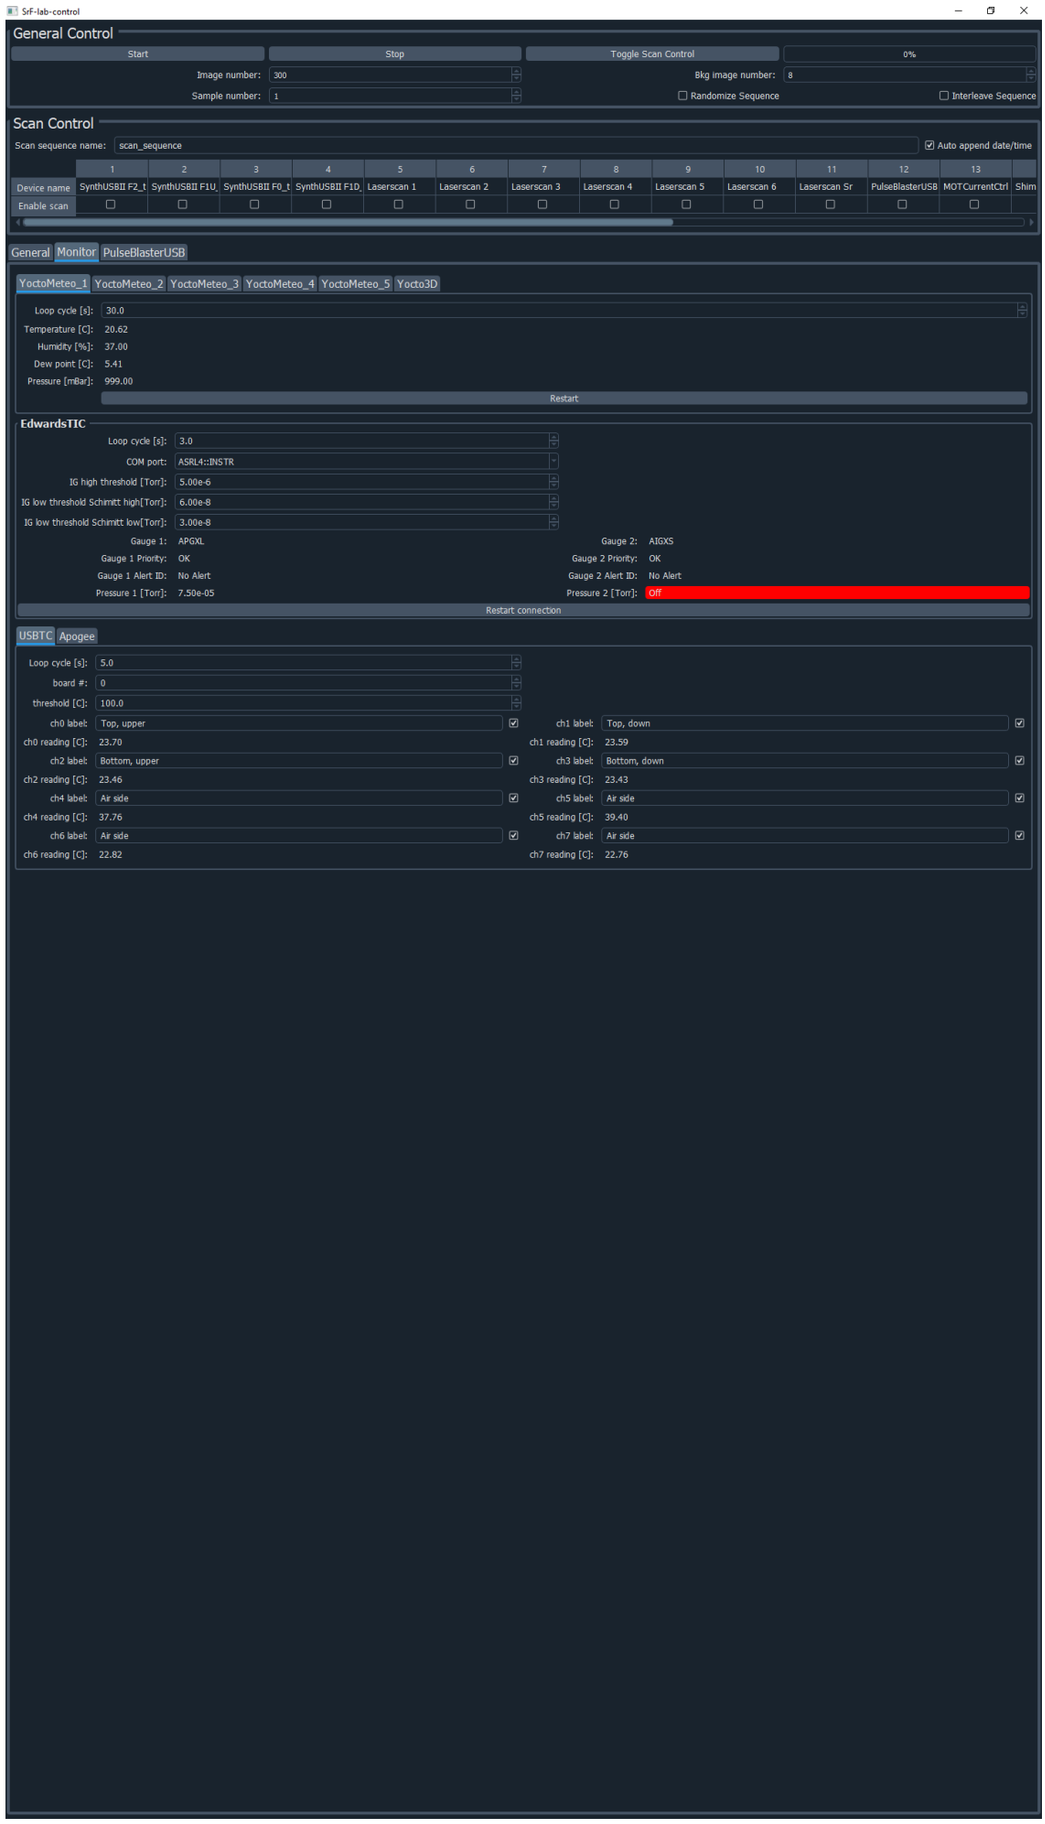

Figure C.3: Example monitor tab in the main program.

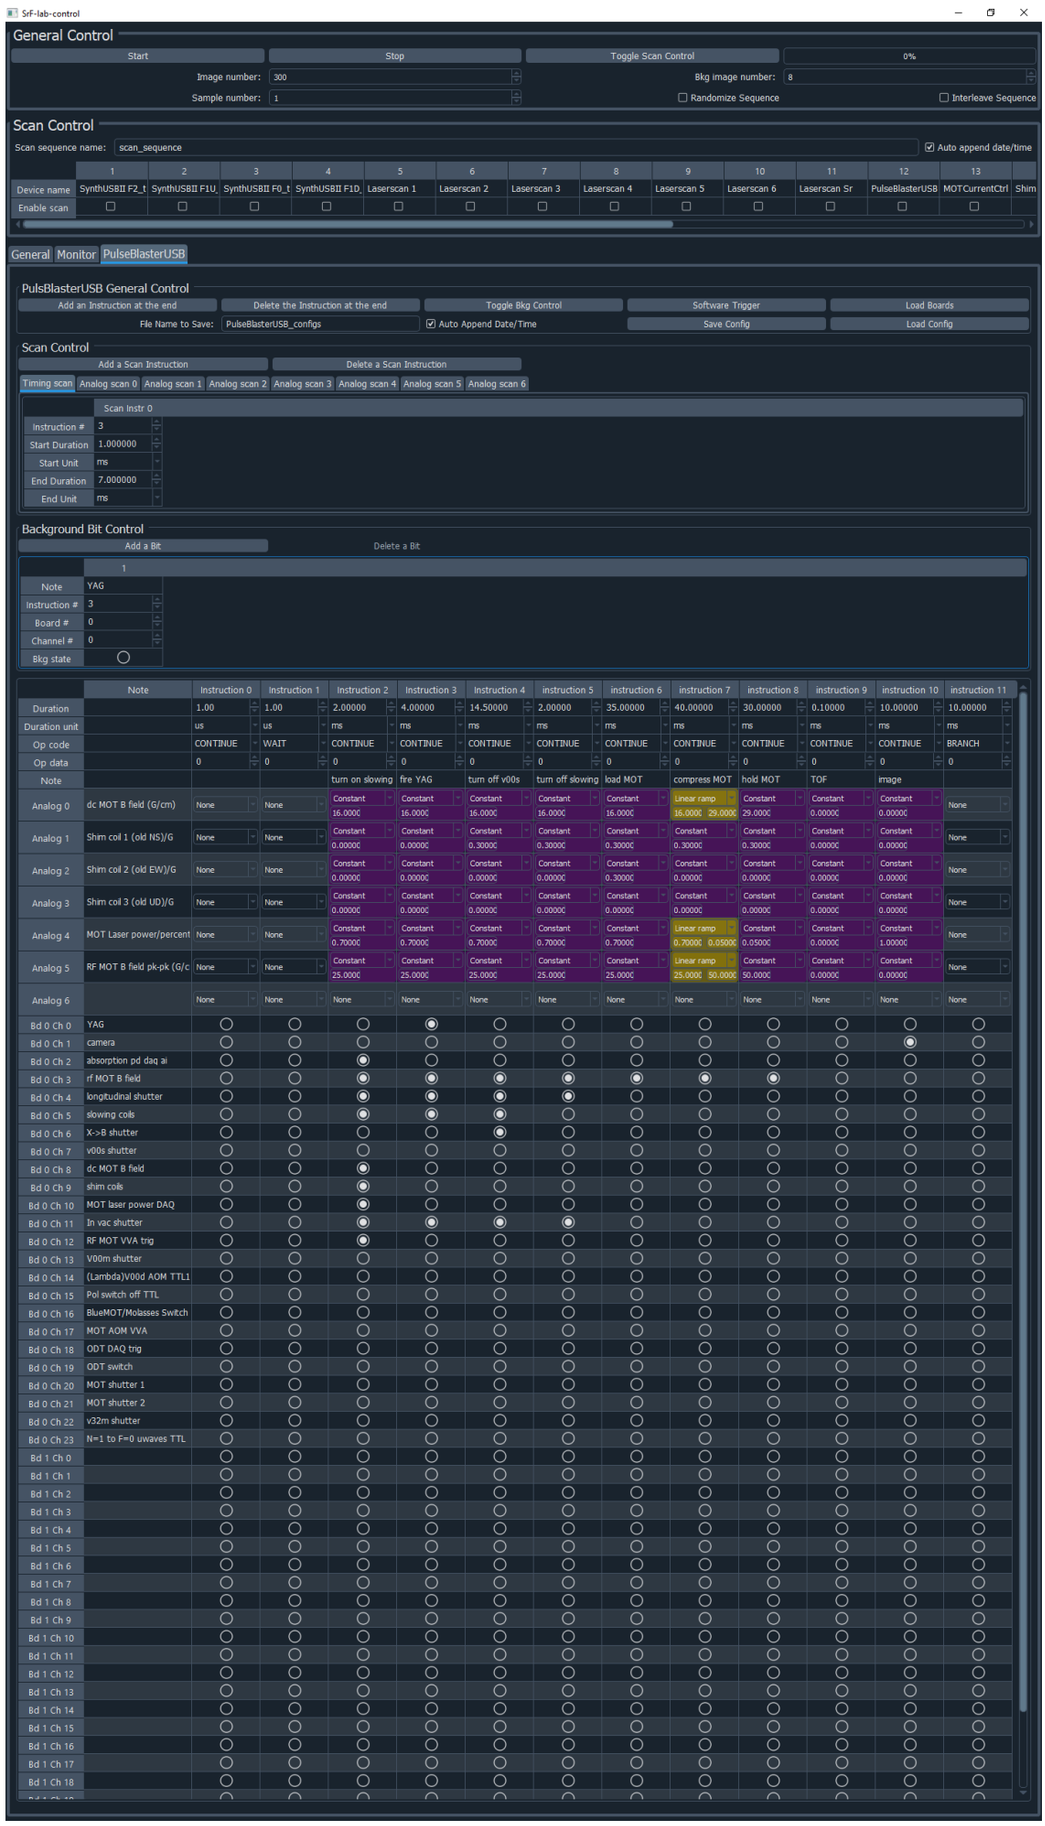

Figure C.4: Example MOT waveform in the main program.

References

1. Carr, L. D., DeMille, D., Krems, R. V. & Ye, J. Cold and ultracold molecules: Science, technology and applications. *New Journal of Physics* **11,** 055049 (2009).
2. Di Rosa, M. D. Laser-cooling molecules: Concept, candidates, and supporting hyperfine-resolved measurements of rotational lines in the A-X(0,0) band of CaH. *European Physical Journal D* **31,** 395–402 (2004).
3. Ospelkaus, S., Ni, K. K., Wang, D., *et al.* Quantum-state controlled chemical reactions of ultracold potassium-rubidium molecules. *Science* **327,** 853–857 (Feb. 2010).
4. McDonald, M., McGuyer, B. H., Apfelbeck, F., *et al.* Photodissociation of ultracold diatomic strontium molecules with quantum state control. *Nature* **535,** 122–126 (July 2016).
5. Hu, M. G., Liu, Y., Grimes, D. D., *et al.* Direct observation of bimolecular reactions of ultracold KRb molecules. *Science* **366,** 1111–1115 (Nov. 2019).
6. Krems, R. V. *Cold controlled chemistry* July 2008.
7. DeMille, D. Quantum Computation with Trapped Polar Molecules. *Physical Review Letters* **88,** 067901 (Jan. 2002).
8. Yu, P., Cheuk, L. W., Kozyryev, I. & Doyle, J. M. A scalable quantum computing platform using symmetric-Top molecules. *New Journal of Physics* **21,** 093049 (2019).
9. Albert, V. V., Covey, J. P. & Preskill, J. Robust encoding of a qubit in a molecule, arXiv:1911.00099v2 (2019).
10. Yelin, S. F., Kirby, K. & Côté, R. Schemes for robust quantum computation with polar molecules. *Physical Review A* **74,** 050301 (2006).
11. Holland, C. M., Lu, Y. & Cheuk, L. W. On-Demand Entanglement of Molecules in a Reconfigurable Optical Tweezer Array (2022).
12. Bao, Y., Yu, S. S., Anderegg, L., *et al.* Fast optical transport of ultracold molecules over long distances (2022).
13. Anderegg, L., Cheuk, L. W., Bao, Y., *et al.* An optical tweezer array of ultracold molecules. *Science* **365,** 1156–1158 (Sept. 2019).
14. Rabl, P., Demille, D., Doyle, J. M., Lukin, M. D., Schoelkopf, R. J. & Zoller, P. Hybrid quantum processors: Molecular ensembles as quantum memory for solid state circuits. *Physical Review Letters* **97,** 033003 (July 2006).
15. Micheli, A., Brennen, G. K. & Zoller, P. A toolbox for lattice-spin models with polar molecules. *Nature Physics* **2,** 341–347 (Apr. 2006).

16. André, A., DeMille, D., Doyle, J. M., *et al.* A coherent all-electrical interface between polar molecules and mesoscopic superconducting resonators. *Nature Physics* **2,** 636–642 (Aug. 2006).

17. Andreev, V., Ang, D. G., DeMille, D., *et al.* Improved limit on the electric dipole moment of the electron. *Nature* **562,** 355–360 (Oct. 2018).

18. Altuntaş, E., Ammon, J., Cahn, S. B. & Demille, D. Demonstration of a Sensitive Method to Measure Nuclear-Spin-Dependent Parity Violation. *Physical Review Letters* **120,** 142501 (Apr. 2018).

19. Cairncross, W. B., Gresh, D. N., Grau, M., *et al.* Precision Measurement of the Electron's Electric Dipole Moment Using Trapped Molecular Ions. *Physical Review Letters* **119,** 153001 (2017).

20. Kozyryev, I. & Hutzler, N. R. Precision Measurement of Time-Reversal Symmetry Violation with Laser-Cooled Polyatomic Molecules. *Physical Review Letters* **119,** 133002 (2017).

21. Roussy, T. S., Caldwell, L., Wright, T., *et al.* An improved bound on the electron's electric dipole moment. *Science (New York, N.Y.)* **381,** 46–50 (July 2023).

22. Baron, J., Campbell, W. C., DeMille, D., *et al.* Order of magnitude smaller limit on the electric dipole moment of the electron. *Science* **343,** 269–272 (Jan. 2014).

23. Yu, P. & Hutzler, N. R. Probing Fundamental Symmetries of Deformed Nuclei in Symmetric Top Molecules. *Physical Review Letters* **126,** 023003 (Jan. 2021).

24. Hudson, E. R., Lewandowski, H. J., Sawyer, B. C. & Ye, J. Cold molecule spectroscopy for constraining the evolution of the fine structure constant. *Physical Review Letters* **96,** 143004 (Apr. 2006).

25. Tarbutt, M. R., Sauer, B. E., Hudson, J. J. & Hinds, E. A. Design for a fountain of YbF molecules to measure the electron's electric dipole moment. *New Journal of Physics* **15,** 053034 (May 2013).

26. Fried, D. G., Killian, T. C., Willmann, L., *et al.* Bose-einstein condensation of atomic hydrogen. *Physical Review Letters* **81,** 3811–3814 (Nov. 1998).

27. Bradley, C. C., Sackett, C. A., Tollett, J. J. & Hulet, R. G. Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions. *Physical Review Letters* **75,** 1687–1690 (Aug. 1995).

28. Davis, K. B., Mewes, M. O., Andrews, M. R., *et al.* Bose-Einstein condensation in a gas of sodium atoms. *Physical Review Letters* **75,** 3969–3973 (Nov. 1995).

29. Modugno, G., Ferrari, G., Roati, G., Brecha, R. J., Simoni, A. & Inguscio, M. Bose-Einstein condensation of potassium atoms by sympathetic cooling. *Science* **294,** 1320–1322 (Nov. 2001).

30. Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. & Cornell, E. A. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science* **269,** 198–201 (July 1995).

31. Weber, T., Herbig, J., Mark, M., Nägerl, H. C. & Grimm, R. Bose-Einstein condensation of cesium. *Science* **299,** 232–235 (Jan. 2003).

32. Kraft, S., Vogt, F., Appel, O., Riehle, F. & Sterr, U. Bose-Einstein Condensation of Alkaline Earth Atoms: Ca40. *Physical Review Letters* **103,** 130401 (Sept. 2009).

33. Stellmer, S., Tey, M. K., Huang, B., Grimm, R. & Schreck, F. Bose-Einstein condensation of strontium. *Physical Review Letters* **103,** 200401 (Nov. 2009).

34. Mickelson, P. G., Martinez De Escobar, Y. N., Yan, M., Desalvo, B. J. & Killian, T. C. Bose-Einstein condensation of Sr88 through sympathetic cooling with Sr87. *Physical Review A - Atomic, Molecular, and Optical Physics* **81,** 051601 (May 2010).

35. Aikawa, K., Frisch, A., Mark, M., *et al.* Bose-Einstein condensation of erbium. *Physical Review Letters* **108,** 210401 (May 2012).

36. Takasu, Y., Maki, K., Komori, K., *et al.* Spin-Singlet Bose-Einstein Condensation of Two-Electron Atoms. *Physical Review Letters* **91,** 040404 (July 2003).

37. Fukuhara, T., Sugawa, S. & Takahashi, Y. Bose-Einstein condensation of an ytterbium isotope. *Physical Review A - Atomic, Molecular, and Optical Physics* **76,** 051604 (Nov. 2007).

38. Rom, T., Best, T., Mandel, O., *et al.* State selective production of molecules in optical lattices. *Physical Review Letters* **93,** 073002 (Aug. 2004).

39. Reinaudi, G., Osborn, C. B., McDonald, M., Kotochigova, S. & Zelevinsky, T. Optical production of stable ultracold Sr288 molecules. *Physical Review Letters* **109,** 115303 (Sept. 2012).

40. Ni, K. K., Ospelkaus, S., De Miranda, M. H., *et al.* A high phase-space-density gas of polar molecules. *Science* **322,** 231–235 (Oct. 2008).

41. Danzl, J. G., Haller, E., Gustavsson, M., *et al.* Quantum gas of deeply bound ground state molecules. *Science* **321,** 1062–1066 (Aug. 2008).

42. De Marco, L., Valtolina, G., Matsuda, K., Tobias, W. G., Covey, J. P. & Ye, J. A degenerate Fermi gas of polar molecules. *Science* **363,** 853–856 (Feb. 2019).

43. Schindewolf, A., Bause, R., Chen, X. Y., *et al.* Evaporation of microwave-shielded polar molecules to quantum degeneracy. *Nature* **607,** 677–681 (July 2022).

44. Bigagli, N., Yuan, W., Zhang, S., *et al.* Observation of Bose-Einstein Condensation of Dipolar Molecules (Dec. 2023).

45. Barry, J. F., Shuman, E. S., Norrgard, E. B. & Demille, D. Laser radiation pressure slowing of a molecular beam. *Physical Review Letters* **108,** 103002 (Mar. 2012).

46. Narevicius, E. & Raizen, M. G. Toward cold chemistry with magnetically decelerated supersonic beams. *Chemical Reviews* **112,** 4879–4889 (2012).

47. Van De Meerakker, S. Y., Bethlem, H. L., Vanhaecke, N. & Meijer, G. Manipulation and control of molecular beams. *Chemical Reviews* **112,** 4828–4878 (2012).

48. Lu, H. I., Kozyryev, I., Hemmerling, B., Piskorski, J. & Doyle, J. M. Magnetic trapping of molecules via optical loading and magnetic slowing. *Physical Review Letters* **112,** 113006 (Mar. 2014).

49. Chervenkov, S., Wu, X., Bayerl, J., *et al.* Continuous centrifuge decelerator for polar molecules. *Physical Review Letters* **112,** 013001 (Jan. 2014).

50. Shuman, E. S., Barry, J. F. & Demille, D. Laser cooling of a diatomic molecule. *Nature* **467,** 820–823 (Oct. 2010).

51. Prehn, A., Ibrügger, M., Glöckner, R., Rempe, G. & Zeppenfeld, M. Optoelectrical Cooling of Polar Molecules to Submillikelvin Temperatures. *Physical Review Letters* **116,** 063005 (Feb. 2016).

52. Son, H., Park, J. J., Ketterle, W. & Jamison, A. O. Collisional cooling of ultracold molecules. *Nature* **580,** 197–200 (Apr. 2020).

53. Cheuk, L. W., Anderegg, L., Augenbraun, B. L., *et al.* Λ-Enhanced Imaging of Molecules in an Optical Trap. *Physical Review Letters* **121,** 083201 (2018).

54. Caldwell, L., Devlin, J. A., Williams, H. J., *et al.* Deep Laser Cooling and Efficient Magnetic Compression of Molecules. *Physical Review Letters* **123,** 033202 (July 2019).

55. Mccarron, D. J., Steinecker, M. H., Zhu, Y. & Demille, D. Magnetic Trapping of an Ultracold Gas of Polar Molecules. *Physical Review Letters* **121,** 13202 (2018).

56. Truppe, S., Williams, H. J., Hambach, M., *et al.* Molecules cooled below the Doppler limit. *Nature Physics* **13,** 1173–1176 (Aug. 2017).

57. Ding, S., Wu, Y., Finneran, I. A., Burau, J. J. & Ye, J. Sub-Doppler Cooling and Compressed Trapping of YO Molecules at μ K Temperatures. *Physical Review X* **10,** arXiv:2002.00056v1 (June 2020).

58. Barry, J. F., Mccarron, D. J., Norrgard, E. B., Steinecker, M. H. & Demille, bibinitperiod D. Magneto-optical trapping of a diatomic molecule. *Nature* **512** (2014).

59. Norrgard, E. B., Mccarron, D. J., Steinecker, M. H., Tarbutt, M. R. & Demille, D. Submillikelvin Dipolar Molecules in a Radio-Frequency Magneto-Optical Trap. *Physical Review Letters* **116,** 063004 (Feb. 2016).

60. Steinecker, M. H., McCarron, D. J., Zhu, Y. & DeMille, D. Improved Radio-Frequency Magneto-Optical Trap of SrF Molecules. *ChemPhysChem* **17,** 3664–3669 (Nov. 2016).

61. Williams, H. J., Truppe, S., Hambach, M., *et al.* Characteristics of a magneto-optical trap of molecules. *New Journal of Physics* **19,** 113035 (Nov. 2017).

62. Anderegg, L., Augenbraun, B. L., Chae, E., *et al.* Radio Frequency Magneto-Optical Trapping of CaF with High Density. *Physical Review Letters* **119,** 103201 (Sept. 2017).

63. Hummon, M. T., Yeo, M., Stuhl, B. K., Collopy, A. L., Xia, Y. & Ye, J. 2D magneto-optical trapping of diatomic molecules. *Physical Review Letters* **110,** 143001 (Apr. 2013).

64. Collopy, A. L., Ding, S., Wu, Y., *et al.* 3D Magneto-Optical Trap of Yttrium Monoxide. *Physical Review Letters* **121,** 213201 (Nov. 2018).

65. Devlin, J. A. & Tarbutt, M. R. Three-dimensional Doppler, polarization-gradient, and magneto-optical forces for atoms and molecules with dark states. *New Journal of Physics* **18,** 123017 (2016).

66. Tarbutt, M. R. & Steimle, T. C. Modeling magneto-optical trapping of CaF molecules. *Physical Review A - Atomic, Molecular, and Optical Physics* **92,** 53401 (2015).

67. Devlin, J. A. & Tarbutt, M. R. Laser cooling and magneto-optical trapping of molecules analyzed using optical Bloch equations and the Fokker-Planck-Kramers equation. *Physical Review A* **98,** 063415 (Dec. 2018).

68. Wu, Y., Burau, J. J., Mehling, K., Ye, J. & Ding, S. High Phase-Space Density of Laser-Cooled Molecules in an Optical Lattice. *Physical Review Letters* **127,** 263201 (Dec. 2021).

69. Vilas, N. B., Hallas, C., Anderegg, L., *et al.* Magneto-optical trapping and sub-Doppler cooling of a polyatomic molecule. *Nature* **606,** 70–74 (June 2022).

70. Augenbraun, B. L., Frenett, A., Sawaoka, H., *et al.* Zeeman-Sisyphus Deceleration of Molecular Beams. *Physical Review Letters* **127,** 263002 (Dec. 2021).

71. Baum, L., Vilas, N. B., Hallas, C., *et al.* 1D Magneto-Optical Trap of Polyatomic Molecules. *Physical Review Letters* **124,** 133201 (Apr. 2020).

72. Cheuk, L. W., Anderegg, L. L. L., Bao, Y., *et al.* Observation of collisions between two ultracold ground-State CaF molecules. *Physical Review Letters* **125,** 043401 (July 2020).

73. Jurgilas, S., Chakraborty, A., Rich, C. J., *et al.* Collisions between Ultracold Molecules and Atoms in a Magnetic Trap. *Physical Review Letters* **126,** 153401 (Apr. 2021).

74. Segev, Y., Pitzer, M., Karpov, M., Akerman, N., Narevicius, J. & Narevicius, E. Collisions between cold molecules in a superconducting magnetic trap. *Nature* **572,** 189–193 (Aug. 2019).

75. Barry, J. F. *Laser cooling and slowing of a diatomic molecule* PhD thesis (2013).

76. Norrgard, E. B. *Magneto-optical trapping of diatomic molecules* PhD thesis (Yale University, 2016).

77. Brown, J. M. & Carrington, A. *Rotational Spectroscopy of Diatomic Molecules* (Cambridge University Press, Apr. 2003).

78. Hutzler, N. R. *A New Limit on the Electron Electric Dipole Moment: Beam Production, Data Interpretation, and Systematics* PhD thesis (Harvard University, June 2014), 1–322.

79. Hao, Y., Pašteka, L. F., Visscher, L., *et al.* High accuracy theoretical investigations of CaF, SrF, and BaF and implications for laser-cooling. *Journal of Chemical Physics* **151** (July 2019).

80. Karman, T. & Hutson, J. M. Microwave Shielding of Ultracold Polar Molecules. *Physical Review Letters* **121,** 163401 (Oct. 2018).

81. Lassablière, L. & Quéméner, G. Controlling the Scattering Length of Ultracold Dipolar Molecules. *Physical Review Letters* **121,** 163402 (Oct. 2018).

82. Quéméner, G. & Bohn, J. L. Shielding Σ 2 ultracold dipolar molecular collisions with electric fields. *Physical Review A* **93,** 012704 (Jan. 2016).

83. Li, J. R., Tobias, W. G., Matsuda, K., *et al.* Tuning of dipolar interactions and evaporative cooling in a three-dimensional molecular quantum gas. *Nature Physics* **17,** 1144–1148 (Sept. 2021).

84. Grimm, R., Weidemüller, M. & Ovchinnikov, Y. B. Optical Dipole Traps for Neutral Atoms. *Advances in Atomic, Molecular and Optical Physics* **42,** 95–170 (Feb. 2000).

85. Ludlow, A. D., Boyd, M. M., Ye, J., Peik, E. & Schmidt, P. O. Optical atomic clocks. *Reviews of Modern Physics* **87,** 637 (June 2015).

86. Baldwin, C. H., Bjork, B. J., Foss-Feig, M., *et al.* High-fidelity light-shift gate for clock-state qubits. *Physical Review A* **103,** 012603 (Jan. 2021).

87. Lisdat, C., Dörscher, S., Nosske, I. & Sterr, U. Blackbody radiation shift in strontium lattice clocks revisited. *Physical Review Research* **3,** L042036 (Dec. 2021).

88. Saffman, M. *Quantum computing with atomic qubits and Rydberg interactions: Progress and challenges* Oct. 2016.

89. Jonathan, D. & Plenio, M. B. Light-shift-induced quantum gates for ions in thermal motion. *Physical Review Letters* **87,** 127901/1–127901/4 (Sept. 2001).

90. Langhoff, S. R., Bauschlicher, C. W., Partridge, H. & Ahlrichs, R. Theoretical study of the dipole moments of selected alkaline-earth halides. *The Journal of Chemical Physics* **84,** 5025–5031 (May 1986).

91. Sheridan, P. M., Wang, J. G., Dick, M. J. & Bernath, P. F. Optical-optical double resonance spectroscopy of the C 2Π- A 2Π and D 2Σ +-A 2Π Transitions of SrF. *Journal of Physical Chemistry A* **113,** 13383–13389 (2009).

92. Zhu, Y. *Experiments with 87Rb: Towards Co-trapping 88Sr19F and 87Rb* PhD thesis (Yale University, Apr. 2021).

93. Steinecker, M. H. *Sub-Doppler Laser Cooling and Magnetic Trapping of SrF Molecules* PhD thesis (Yale University, 2019).

94. Metcalf, H. J. & van der Straten, P. *Laser Cooling and Trapping* (Springer New York, New York, NY, 1999).

95. Fleischhauer, M., Imamoglu, A. & Marangos, P. J. Electromagnetically induced transparency. *Reviews of Modern Physics* **77,** 633–673 (Apr. 2005).

96. Jones, M. P. A. Atom-Light Interactions.

97. Aspect, A., Arimondo, E., Kaiser, R., Vansteenkiste, N. & Cohen-Tannoudji, C. Laser cooling below the one-photon recoil energy by velocity-selective coherent population trapping. *Physical Review Letters* **61,** 826–829 (1988).

98. Grier, A. T., Ferrier-Barbut, I., Rem, B. S., *et al.* Λ-enhanced sub-Doppler cooling of lithium atoms in D1 gray molasses. *Physical Review A - Atomic, Molecular, and Optical Physics* **87,** 063411 (June 2013).

99. Rio Fernandes, D., Sievers, F., Kretzschmar, N., Wu, S., Salomon, C. & Chevy, F. Sub-Doppler laser cooling of fermionic 40K atoms in three-dimensional gray optical molasses. *EPL* **100,** 63001 (Dec. 2012).

100. Salomon, G., Fouché, L., Wang, P., Aspect, A., Bouyer, P. & Bourdel, T. Gray-molasses cooling of 39K to a high phase-space density. *EPL* **104,** 63002 (Jan. 2013).

101. Chen, H. Z., Yao, X. C., Wu, Y. P., *et al.* Production of large K 41 Bose-Einstein condensates using D1 gray molasses. *Physical Review A* **94,** 033408 (Sept. 2016).

102. Colzi, G., Durastante, G., Fava, E., Serafini, S., Lamporesi, G. & Ferrari, G. Sub-Doppler cooling of sodium atoms in gray molasses. *Physical Review A* **93,** 023421 (Feb. 2016).

103. Bouton, Q., Chang, R., Hoendervanger, A. L., *et al.* Fast production of Bose-Einstein condensates of metastable helium. *Physical Review A - Atomic, Molecular, and Optical Physics* **91,** 061402 (June 2015).

104. Burchianti, A., Valtolina, G., Seman, J. A., *et al.* Efficient all-optical production of large Li 6 quantum gases using D1 gray-molasses cooling. *Physical Review A - Atomic, Molecular, and Optical Physics* **90,** 043408 (Oct. 2014).

105. Nath, D., Easwaran, R. K., Rajalakshmi, G. & Unnikrishnan, C. S. Quantum-interference-enhanced deep sub-Doppler cooling of 39K atoms in gray molasses. *Physical Review A - Atomic, Molecular, and Optical Physics* **88,** 053407 (Nov. 2013).

106. Sievers, F., Kretzschmar, N., Fernandes, D. R., *et al.* Simultaneous sub-Doppler laser cooling of fermionic Li 6 and K 40 on the D1 line: Theory and experiment. *Physical Review A - Atomic, Molecular, and Optical Physics* **91,** 023426 (Feb. 2015).

107. Boiron, D., Michaud, A., Lemonde, P., *et al.* Laser cooling of cesium atoms in gray optical molasses down to 1.1 μK. *Physical Review A - Atomic, Molecular, and Optical Physics* **53,** R3734–R3737 (June 1996).

108. Triché, C., Verkerk, P. & Grynberg, G. Blue-sisyphus cooling in cesium gray molasses and antidot lattices. *European Physical Journal D* **5,** 225–228 (1999).

109. Esslinger, T., Ritsch, H., Weidemüller, M., Sander, F., Hemmerich, A. & Hänsch, T. W. Purely optical dark lattice. *Optics Letters* **21,** 991 (July 1996).

110. Langin, T. K. & DeMille, D. Toward improved loading, cooling, and trapping of molecules in magneto-optical traps. *New Journal of Physics* **25,** 043005 (Apr. 2023).

111. Langin, T. K., Jorapur, V., Zhu, Y., Wang, Q. & DeMille, D. Polarization Enhanced Deep Optical Dipole Trapping of -Cooled Polar Molecules. *Physical Review Letters* **127,** 163201 (Oct. 2021).

112. Tarbutt, M. R. Magneto-optical trapping forces for atoms and molecules with complex level structures. *New Journal of Physics* **17,** 015007 (Jan. 2015).

113. Overstreet, K. R., Zabawa, P., Tallant, J., Schwettmann, A. & Shaffer, J. P. Multiple scattering and the density distribution of a Cs MOT. *Optics Express* **13,** 9672 (Nov. 2005).

114. Migdall, A. L., Prodan, J. V., Phillips, W. D., Bergeman, T. H. & Metcalf, H. J. First observation of magnetically trapped neutral atoms. *Physical Review Letters* **54,** 2596–2599 (June 1985).

115. Petrich, W., Anderson, M. H., Ensher, J. R. & Cornell, E. A. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Physical Review Letters* **74,** 3352–3355 (Apr. 1995).

116. Moore, K. L., Purdy, T. P., Murch, K. W., *et al.* Bose-Einstein condensation in a mm-scale Ioffe-Pritchard trap. *Applied Physics B: Lasers and Optics* **82,** 533–538 (Mar. 2006).

117. Adams, C. S., Lee, H. J., Davidson, N., Kasevich, M. & Chu, S. Evaporative cooling in a crossed dipole trap. *Physical Review Letters* **74,** 3577–3580 (May 1995).

118. Weinstein, J. D., DeCarvalho, R., Guillet, T., Friedrich, B. & Doyle, J. M. Magnetic trapping of calcium monohydride molecules at millikelvin temperatures. *Nature* **395,** 148–150 (Sept. 1998).

119. Hess, H. F., Kochanski, G. P., Doyle, J. M., Masuhara, N., Kleppner, D. & Greytak, T. J. Magnetic trapping of spin-polarized atomic hydrogen. *Physical Review Letters* **59,** 672–675 (Aug. 1987).

120. Reens, D., Wu, H., Langen, T. & Ye, J. Controlling spin flips of molecules in an electromagnetic trap. *Physical Review A* **96,** 063420 (Dec. 2017).

121. Brink, D. M. & Sukumar, C. V. Majorana spin-flip transitions in a magnetic trap. *Physical Review A - Atomic, Molecular, and Optical Physics* **74,** 035401 (Sept. 2006).

122. Bause, R., Schindewolf, A., Tao, R., *et al.* Collisions of ultracold molecules in bright and dark optical dipole traps. *Physical Review Research* **3,** 033013 (Sept. 2021).

123. Heinz, A., Park, A. J., Šantić, N., *et al.* State-Dependent Optical Lattices for the Strontium Optical Qubit. *Physical Review Letters* **124,** 203201 (May 2020).

124. Lu, Y., Holland, C. M. & Cheuk, L. W. Molecular Laser Cooling in a Dynamically Tunable Repulsive Optical Trap. *Physical Review Letters* **128,** 213201 (May 2022).

125. Dagdigian, P. J., Cruse, H. W. & Zare, R. N. Radiative lifetimes of the alkaline earth monohalides. *Journal of Chemical Physics* **60,** 2330–2339 (Mar. 1974).

126. Berg, L. E., Ekvall, K., Hansson, T., *et al.* Time-resolved measurements of the B 2Σ state of SrF by laser spectroscopy. *Chemical Physics Letters* **248,** 283–288 (Jan. 1996).

127. Ernst, W. E. & Schröder, J. O. The B 2Σ+-X 2Σ+ system of SrF: Precise spectroscopic constants from a combined fit of microwave and sub-doppler optical spectra. *Chemical Physics* **78,** 363–368 (Aug. 1983).

128. De Melo, G. F. & Ornellas, F. R. The Λ + S and Ω electronic states of SrF below 40,000 cm1: A theoretical contribution. *Journal of Quantitative Spectroscopy and Radiative Transfer* **237,** 106632 (Nov. 2019).

129. Nitsch, C., Schröder, J. O. & Ernst, W. E. Optical-optical double-resonance spectroscopy of SrF: The F 2Σ+-B 2Σ+ and G 2Π-B 2Σ+ systems. *Chemical Physics Letters* **148,** 130–135 (July 1988).

130. Sandars, J. R. P. A. & H., P. G. The hyperfine structure Stark effect I. Theory. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **305,** 125–138 (May 1968).

131. Lundblad, N., Schlosser, M. & Porto, J. V. *Erratum: Experimental observation of magic-wavelength behavior of Rb87 atoms in an optical lattice (Physical Review A (2010) 81 (031611))* Mar. 2010.

132. Zheng, T. A., Yang, Y. A., Safronova, M. S., *et al.* Magic wavelengths of the Yb (6s2 S0 1 -6s6p P1 3) intercombination transition. *Physical Review A* **102,** 062805 (Dec. 2020).

133. Guo, K., Wang, G. & Ye, A. Dipole polarizabilities and magic wavelengths for a Sr and Yb atomic optical lattice clock. *Journal of Physics B: Atomic, Molecular and Optical Physics* **43,** 135004 (June 2010).

134. O'Hara, K. M., Granade, S. R., Gehm, M. E. & Thomas, J. E. Loading dynamics of CO2 laser traps. *Physical Review A - Atomic, Molecular, and Optical Physics* **63,** 1–5 (Mar. 2001).

135. Roy, R., Green, A., Bowler, R. & Gupta, S. Rapid cooling to quantum degeneracy in dynamically shaped atom traps. *Physical Review A* **93,** 043403 (Apr. 2016).

136. Rosenband, T., Grimes, D. D. & Ni, K.-K. Elliptical polarization for molecular Stark shift compensation in deep optical traps. *Optics Express* **26,** 19821 (Aug. 2018).

137. Rosi, S., Burchianti, A., Conclave, S., *et al.* Λ-enhanced grey molasses on the D 2 transition of Rubidium-87 atoms. *Scientific Reports* **8,** 1–9 (Jan. 2018).

138. Kulatunga, P., Blum, T. & Olek, D. Loading characteristics of a microscopic optical dipole trap (Sept. 2010).

139. Kuppens, S. J., Corwin, K. L., Miller, K. W., Chupp, T. E. & Wieman, C. E. Loading an optical dipole trap. *Physical Review A - Atomic, Molecular, and Optical Physics* **62,** 13 (June 2000).

140. Fuhrmanek, A., Bourgain, R., Sortais, Y. R. & Browaeys, A. Light-assisted collisions between a few cold atoms in a microscopic dipole trap. *Physical Review A - Atomic, Molecular, and Optical Physics* **85,** 062708 (June 2012).

141. Anderegg, L. *Ultracold molecules in optical arrays: from laser cooling to molecular collisions* PhD thesis (Harvard University, Jan. 2020).

142. Idziaszek, Z. & Julienne, P. S. Universal rate constants for reactive collisions of ultracold molecules. *Physical Review Letters* **104,** 113202 (Mar. 2010).

143. Gregory, P. D., Frye, M. D., Blackmore, J. A., *et al.* Sticky collisions of ultracold RbCs molecules. *Nature Communications* **10,** 1–7 (July 2019).

144. Frye, M. D., Julienne, P. S. & Hutson, J. M. Cold atomic and molecular collisions: Approaching the universal loss regime. *New Journal of Physics* **17,** 045019 (Apr. 2015).

145. Barry, J. F., Shuman, E. S. & Demille, D. A bright, slow cryogenic molecular beam source for free radicals. *Physical Chemistry Chemical Physics* **13,** 18936–18947 (Nov. 2011).

146. Truppe, S., Hambach, M., Skoff, S. M., *et al.* A buffer gas beam source for short, intense and slow molecular pulses. *Journal of Modern Optics* **65,** 648–656 (Mar. 2018).

147. Patterson, D., Rasmussen, J. & Doyle, J. M. Intense atomic and molecular beams via neon buffer-gas cooling. *New Journal of Physics* **11,** 055018 (May 2009).

148. Hutzler, N. R., Lu, H. I. & Doyle, J. M. *The buffer gas beam: An intense, cold, and slow source for atoms and molecules* Sept. 2012.

149. Bulleid, N. E., Skoff, S. M., Hendricks, R. J., Sauer, B. E., Hinds, E. A. & Tarbutt, M. R. Characterization of a cryogenic beam source for atoms and molecules. *Physical Chemistry Chemical Physics* **15,** 12299–12307 (Aug. 2013).

150. Egorov, D., Lahaye, T., Schöllkopf, W., Friedrich, B. & Doyle, J. M. Buffer-gas cooling of atomic and molecular beams. *Physical Review A - Atomic, Molecular, and Optical Physics* **66,** 8 (Oct. 2002).

151. Xiao, D., Lancaster, D. M., Allen, C. H., *et al.* Shaped nozzles for cryogenic buffer-gas beam sources. *Physical Review A* **99,** 013603 (Jan. 2019).

152. Skoff, S. M., Hendricks, R. J., Sinclair, C. D., *et al.* Diffusion, thermalization, and optical pumping of YbF molecules in a cold buffer-gas cell. *Physical Review A - Atomic, Molecular, and Optical Physics* **83,** 023418 (Feb. 2011).

153. Lu, H. I., Rasmussen, J., Wright, M. J., Patterson, D. & Doyle, J. M. A cold and slow molecular beam. *Physical Chemistry Chemical Physics* **13,** 18986–18990 (Oct. 2011).

154. Hutzler, N. R., Parsons, M. F., Gurevich, Y. V., *et al.* A cryogenic beam of refractory, chemically reactive molecules with expansion cooling. *Physical Chemistry Chemical Physics* **13,** 18976–18985 (Oct. 2011).

155. Patterson, D. & Doyle, J. M. Bright, guided molecular beam with hydrodynamic enhancement. *Journal of Chemical Physics* **126,** 154307 (Apr. 2007).

156. Santamaria, L., Sarno, V. D., Natale, P. D., *et al.* Comb-assisted cavity ring-down spectroscopy of a buffer-gas-cooled molecular beam. *Physical Chemistry Chemical Physics* **18,** 16715–16720 (June 2016).

157. Porterfield, J. P., Satterthwaite, L., Eibenberger, S., Patterson, D. & McCarthy, M. C. High sensitivity microwave spectroscopy in a cryogenic buffer gas cell. *Review of Scientific Instruments* **90,** 53104 (May 2019).

158. Spaun, B., Changala, P. B., Patterson, D., *et al.* Continuous probing of cold complex molecules with infrared frequency comb spectroscopy. *Nature* **533,** 517–520 (May 2016).

159. Maxwell, S. E., Brahms, N., Decarvalho, R., *et al.* High-flux beam source for cold, slow atoms or molecules. *Physical Review Letters* **95,** 173201 (Oct. 2005).

160. Wright, S. C., Doppelbauer, M., Hofsäss, S., *et al.* Cryogenic buffer gas beams of AlF, CaF, MgF, YbF, Al, Ca, Yb and NO–a comparison. *Molecular Physics* **121,** 17–18 (Sept. 2022).

161. Takahashi, Y., Shlivko, D., Woolls, G. & Hutzler, N. R. Simulation of cryogenic buffer gas beams. *Physical Review Research* **3,** 023018 (Apr. 2021).

162. Doppelbauer, M. J., Schullian, O., Loreau, J., *et al.* Using a direct simulation Monte Carlo approach to model collisions in a buffer gas cell. *Journal of Chemical Physics* **146,** 044302 (Jan. 2017).

163. Singh, V., Samanta, A. K., Roth, N., *et al.* Optimized cell geometry for buffer-gas-cooled molecular-beam sources. *Physical Review A* **97,** 032704 (Mar. 2018).

164. Jadbabaie, A., Pilgram, N. H., Kłos, J., Kotochigova, S. & Hutzler, N. R. Enhanced molecular yield from a cryogenic buffer gas beam source via excited state chemistry. *New Journal of Physics* **22,** 022002 (Feb. 2020).

165. Anderegg, L., Augenbraun, B. L., Bao, Y., *et al.* Laser cooling of optically trapped molecules. *Nature Physics* **14,** 890–893 (2018).

166. McCarron, D. J., Norrgard, E. B., Steinecker, M. H. & DeMille, D. Improved magneto-optical trapping of a diatomic molecule. *New Journal of Physics* **17,** 1–12 (2015).

167. Jarvis, K. N., Devlin, J. A., Wall, T. E., Sauer, B. E. & Tarbutt, M. R. Blue-Detuned Magneto-Optical Trap. *Physical Review Letters* **120,** 083201 (Feb. 2018).

168. Burau, J. J., Aggarwal, P., Mehling, K. & Ye, J. Blue-Detuned Magneto-optical Trap of Molecules. *Physical Review Letters* **130,** 193401 (May 2023).

169. Barry, J. F. *Laser cooling of a diatomic molecule* PhD thesis (Yale University, 2013).

170. Dalibard, J. & Cohen-Tannoudji, C. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *Journal of the Optical Society of America B* **6,** 2023 (Nov. 1989).

171. Hallas, C., Vilas, N. B., Anderegg, L., *et al.* Optical Trapping of a Polyatomic Molecule in an -Type Parity Doublet State. *Physical Review Letters* **130,** 153202 (Apr. 2023).

172. Rosi, S., Burchianti, A., Conclave, S., *et al.* Λ-enhanced grey molasses on the D 2 transition of Rubidium-87 atoms. *Scientific Reports* **8,** 1–9 (Dec. 2018).

173. Li, S. J., Holland, C. M., Lu, Y. & Cheuk, L. W. A Blue-Detuned Magneto-Optical Trap of CaF Molecules. *arXiv preprint* (Nov. 2023).

174. He, J., Ye, X., Lin, J., Guo, M., Quéméner, G. & Wang, D. Observation of resonant dipolar collisions in ultracold Na 23 Rb 87 rotational mixtures. *Physical Review Research* **3,** 013016 (Jan. 2021).

175. Voges, K. K., Gersema, P., Meyer Zum Alten Borgloh, M., *et al.* Ultracold Gas of Bosonic Na 23 K 39 Ground-State Molecules. *Physical Review Letters* **125,** 083401 (Aug. 2020).

176. Bause, R., Christianen, A., Schindewolf, A., Bloch, I. & Luo, X. Y. Ultracold Sticky Collisions: Theoretical and Experimental Status. *Journal of Physical Chemistry A* **127,** 729–741 (Jan. 2023).

177. Takekoshi, T., Reichsöllner, L., Schindewolf, A., *et al.* Ultracold dense samples of dipolar RbCs molecules in the rovibrational and hyperfine ground state. *Physical Review Letters* **113,** 205301 (Nov. 2014).

178. Park, J. W., Will, S. A. & Zwierlein, M. W. Ultracold Dipolar Gas of Fermionic Na23K40 Molecules in Their Absolute Ground State. *Physical Review Letters* **114,** 205302 (May 2015).

179. Drews, B., Deiß, M., Jachymski, K., Idziaszek, Z. & Denschlag, J. H. Inelastic collisions of ultracold triplet Rb2 molecules in the rovibrational ground state. *Nature Communications* **8,** 1–7 (Mar. 2017).

180. Gregory, P. D., Aldegunde, J., Hutson, J. M. & Cornish, S. L. Controlling the rotational and hyperfine state of ultracold Rb 87 Cs 133 molecules. *Physical Review A* **94,** 041403 (Oct. 2016).

181. Rvachov, T. M., Son, H., Sommer, A. T., *et al.* Long-Lived Ultracold Molecules with Electric and Magnetic Dipole Moments. *Physical Review Letters* **119,** 143001 (Oct. 2017).

182. Ye, X., Guo, M., González-Martínez, M. L., Quéméner, G. & Wang, D. Collisions of ultracold 23Na87Rb molecules with controlled chemical reactivities. *Science Advances* **4** (Jan. 2018).

183. Anderegg, L., Burchesky, S., Bao, Y., *et al.* Observation of microwave shielding of ultracold molecules. *Science* **373,** 779–782 (Aug. 2021).

184. Lin, J., Chen, G., Jin, M., *et al.* Microwave Shielding of Bosonic NaRb Molecules. *Physical Review X* **13,** 031032 (July 2023).

185. Bigagli, N., Warner, C., Yuan, W., *et al.* Collisionally stable gas of bosonic dipolar ground-state molecules. *Nature Physics,* 1–6 (Sept. 2023).

186. Ramachandran, P. S., Rajamanickam, N., Bagare, S. P. & Kumar, B. C. Franck-Condon factors and r-centroids for certain band systems of astrophysical molecules SrF and ScF. *Astrophysics and Space Science* **295,** 443–449 (Jan. 2005).

187. Holland, C. M., Lu, Y. & Cheuk, L. W. Synthesizing optical spectra using computer-generated holography techniques. *New Journal of Physics* **23,** 033028 (Mar. 2021).

188. Truppe, S., Williams, H. J., Fitch, N. J., *et al.* An intense, cold, velocity-controlled molecular beam by frequency-chirped laser slowing. *New Journal of Physics* **19,** 022001 (Feb. 2017).

189. Meyer, E. R. & Bohn, J. L. Chemical pathways in ultracold reactions of SrF molecules. *Physical Review A - Atomic, Molecular, and Optical Physics* **83,** 032714 (Mar. 2011).

190. Landau, L. D. & Lifshitz, E. M. *Course of Theoretical Physics - Vol 3 Quantum Mechanics - nonrelativistic theory 3rd ed* (Butterworth-Heinemann, 1981).

191. Lepers, M. & Dulieu, O. in *RSC Theoretical and Computational Chemistry Series* 11, 150–202 (Mar. 2018).

192. Quéméner, G., Bohn, J. L., Petrov, A. & Kotochigova, S. Universalities in ultracold reactions of alkali-metal polar molecules. *Physical Review A - Atomic, Molecular, and Optical Physics* **84,** 062703 (Dec. 2011).

193. Caldwell, L. *Deep Laser Cooling and Coherent Control of Molecules* PhD thesis (2020).

194. Hudson, J. J., Kara, D. M., Smallman, I. J., Sauer, B. E., Tarbutt, M. R. & Hinds, E. A. Improved measurement of the shape of the electron. *Nature* **473,** 493–496 (May 2011).

195. Quéméner, G. & Bohn, J. L. Ultracold molecular collisions in combined electric and magnetic fields. *Physical Review A - Atomic, Molecular, and Optical Physics* **88,** 012706 (July 2013).

196. Matsuda, K., De Marco, L., Li, J. R., *et al.* Resonant collisional shielding of reactive molecules using electric fields. *Science* **370,** 1324–1327 (Dec. 2020).

197. Rosenberg, J., Christakis, L., Raj, R., *et al.* Collisional shielding of a bosonic gas of polar molecules. *APS* **2023,** K07.003 (2023).

198. Mukherjee, B., Frye, M. D., Le Sueur, C. R., Tarbutt, M. R. & Hutson, J. M. Shielding collisions of ultracold CaF molecules with static electric fields. *Physical Review Research* **5,** 033097 (July 2023).

199. Karman, T. & Hutson, J. M. Microwave shielding of ultracold polar molecules with imperfectly circular polarization. *Physical Review A* **100,** 052704 (Nov. 2019).

200. Karman, T. Microwave shielding with far-from-circular polarization. *Physical Review A* **101,** 042702 (Apr. 2020).

201. Gorshkov, A. V., Rabl, P., Pupillo, G., *et al.* Suppression of inelastic collisions between polar molecules with a repulsive shield. *Physical Review Letters* **101,** 073201 (Aug. 2008).

202. Bigagli, N., Warner, C., Yuan, W., *et al.* Collisionally stable gas of bosonic dipolar ground-state molecules. *Nature Physics* **19,** 1579–1584 (Sept. 2023).

203. Chen, X.-Y., Duda, M., Bause, R., *et al.* Microwave dressed polar molecules in two-dimensional optical lattices. *APS* **2022,** M08.010 (2022).

204. Boudot, R. & Rubiola, E. Phase noise in RF and microwave amplifiers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **59,** 2613–2624 (2012).

205. Microwave, M. *How do amplifiers affect signal phase noise? – Marki Microwave RF & Microwave*

206. Goldsmith, P. F. *Quasioptical systems : Gaussian beam quasioptical propagation and applications* (IEEE Press, 1998).

207. Murphy, J. A. Aperture efficiencies of large axisymtnetric reflector antennas fed by conical horns. *IEEE Transactions on Antennas and Propagation* **36,** 570–575 (1988).

208. Chen, X. Y., Schindewolf, A., Eppelt, S., *et al.* Field-linked resonances of polar molecules. *Nature* **614,** 59–63 (Feb. 2023).

209. Alauze, X., Lim, J., Trigatzis, M. A., *et al.* An ultracold molecular beam for testing fundamental physics. *Quantum Science and Technology* **6,** 044005 (July 2021).

210. Lim, J., Frye, M. D., Hutson, J. M. & Tarbutt, M. R. Modeling sympathetic cooling of molecules by ultracold atoms. *Physical Review A - Atomic, Molecular, and Optical Physics* **92,** 053419 (Nov. 2015).

211. Bird, R. C., Tarbutt, M. R. & Hutson, J. M. Tunable Feshbach resonances in collisions of ultracold molecules in 2 states with alkali-metal atoms. *Physical Review Research* **5,** 023184 (Apr. 2023).

212. Park, J. J., Lu, Y. K., Jamison, A. O. & Ketterle, W. Magnetic trapping of ultracold molecules at high density. *Nature Physics* **19,** 1567–1572 (July 2023).

213. Park, J. J., Lu, Y. K., Jamison, A. O., Tscherbul, T. V. & Ketterle, W. A Feshbach resonance in collisions between triplet ground-state molecules. *Nature* **614,** 54–58 (Feb. 2023).

214. Son, H., Park, J. J., Ketterle, W. & Jamison, A. O. Collisional cooling of ultracold molecules. *Nature* **580,** 197–200 (Apr. 2020).

215. Morita, M., Kosicki, M. B., Zuchowski, P. S. & Tscherbul, T. V. Atom-molecule collisions, spin relaxation, and sympathetic cooling in an ultracold spin-polarized Rb(S 2)-SrF(Σ + 2) mixture. *Physical Review A* **98,** 042702 (Oct. 2018).

216. Olson, A. J., Niffenegger, R. J. & Chen, Y. P. Optimizing the efficiency of evaporative cooling in optical dipole traps. *Physical Review A - Atomic, Molecular, and Optical Physics* **87,** 053613 (May 2013).

217. Ma, Z. Y., Thomas, A. M., Foot, C. J. & Cornish, S. L. The evaporative cooling of a gas of caesium atoms in the hydrodynamic regime. *Journal of Physics B: Atomic, Molecular and Optical Physics* **36,** 3533–3540 (Aug. 2003).

218. Wang, R. R. & Bohn, J. L. Anisotropic thermalization of dilute dipolar gases. *Physical Review A* **103,** 063320 (June 2021).

219. Davis, K. B., Mewes, M. O. & Ketterle, W. An analytical model for evaporative cooling of atoms. *Applied Physics B Laser and Optics* **60,** 155–159 (Feb. 1995).

220. Hung, C. L., Zhang, X., Gemelke, N. & Chin, C. Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps. *Physical Review A - Atomic, Molecular, and Optical Physics* **78,** 011604 (July 2008).

221. Warner, C., Lam, A. Z., Bigagli, N., Liu, H. C., Stevenson, I. & Will, S. Overlapping Bose-Einstein condensates of Na 23 and Cs 133. *Physical Review A* **104,** 033302 (Sept. 2021).