

High-Frequency Continuous-Wave Optical Modulation via Stimulated Raman Scattering

by

David Claiborne Gold

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2020

Date of final oral examination: May 28, 2020

The dissertation is approved by the following members of the Final Oral Committee:

Deniz D. Yavuz, Professor, Physics

James E. Lawler, Professor, Physics

Thad G. Walker, Professor, Physics

Daniel van der Weide, Professor, Electrical and Computer Engineering

© Copyright by David Claiborne Gold, 2020

All Rights Reserved

For all the 1064 nm photons who dream of growing up to be 1555 nm

Acknowledgments

I've been very fortunate to have had the support and friendship of many people during my time in Madison, which contributed to both my work and happiness while here. I stumbled into my lab group somewhat by happenstance, but over the years I have realized how lucky an occurrence that was. My advisor, Deniz Yavuz, was everything I could have hoped for. He always made time to meet when I had questions, where his deep knowledge of physics showed through in his ease of discussing a difficult concept from many viewpoints and at any level of complexity. Despite his extensive knowledge, he always took my own and other group members' opinions very seriously and made us feel very valued. Furthermore, Deniz provided an excellent balance of guidance in my lab work, allowing me to be as independent as I wanted, while always being readily accessible when I needed help. While many students might be motivated in their research due to external pressure from their advisor, among our group members, motivation came solely from our passion for the material and from an earnest desire to be good students for Deniz. When funding was scarce in our lab, Deniz always did more than could be expected to insulate us from the effects of that, not only so that we could continue our research effectively, but also for our own financial security and comfort. The only complaint I could make is that when many of my peers might vent about their advisors in the security of the grad lounge, I had nothing to contribute and felt left out.

The research environment cultivated by Deniz prevailed among his students as well. Josh Weber was an excellent and patient teacher to me during the months he trained me to take over the modulation experiment, and became a good friend in the following years. He gave me my first introduction to many topics in optics and helped instill in me at least a

little of his attention to detail and meticulous nature in the lab. When Josh graduated, Nick Brewer continued to mentor me. In contrast to Josh's by-the-book teachings, Nick's unique approach to experiment, dubbed "street optics", taught me to how to make do with what equipment we had on hand, and helped prepare me for the future lean times in the lab. I'd like to thank Jared Miles as well for teaching me more than I wished to know about EIT and differential screws. In the past few years I've also had the pleasure of collaborating with Atul Bhadkamkar, who has been extremely generous with his time and was integral to my recent progress in the lab.

While far too many friends to list have impacted me and provided support and encouragement in one way or another, I am very grateful to them all. I would however like to extend additional thanks to those who made my time specifically in Chamberlin more enjoyable:

- The core Lab Lunch™ group consisting of Nick Brewer, Zach Buckholtz, Ben Lemberger, Josh Karpel, Dipto Das, and Eli Mueller. Because of them, I think I am one of the most cultured diners of the greater Chamberlin metropolitan area, or at least one of the best fed.
- Officemates in various combinations over the years in room 4340, Nick Brewer, Ben Lemberger, Eli Mueller, Megan Tabbutt, and the Candy Bandit, for making our office a fun place to take reprieve from the toils of the lab.
- Zach Williams for putting in many hours of work with me on Friday afternoons in the grad lounge.
- Brett Unks for frequently loaning me equipment from the optics lab without asking too many questions.

I would also like to give special thanks to:

- My longtime friends, William Garrett, David Hatcher, and Kyle Rothman. Despite extensive tomfoolery in our high school physics class, the many experiments, projects, and adventures with them over the years of their persistent friendship have helped cultivate my interest in physics.
- My parents James and Elizabeth for all their support as well as painstaking assistance with childhood science projects.

I am also especially grateful to Megan Mrozek and our young cat Scooter for all the love and support they've given me, particularly in these rather stressful last few months.

CONTENTS

List of Figures	viii
Abstract	xi
1 Introduction	1
1.1 Pulse Generation	3
1.1.1 Methods of Ultrashort Pulse Generation	10
1.1.2 Arbitrary Waveform Generation	11
1.2 An Optical Modulator	13
1.2.1 Molecular Modulation	14
2 Molecular Modulation	20
2.1 Formalism	23
2.2 Mixing Efficiency	31
3 Experimental Design	35
3.1 Original Experimental Design Overview	37
3.2 The ECDL	38
3.3 The Optical Cavity	40
3.3.1 A Brief Aside	41
3.3.2 Cavity Mode-Shape	42
3.3.3 Mode-Matching	46
3.3.4 Cavity Cooling	48
3.4 Laser Locking	50
4 Past Work	57

4.1	Previous Modulation Research	57
4.2	Past Work in the Yavuz Group	59
5	A 90 THz Modulator	63
5.1	The Mixing Beam	63
5.2	Sideband Detection	66
5.2.1	Detection by Eye	69
5.3	Results	72
6	Cavity Redesign	74
6.1	Gas Pressure	75
6.2	Cavity Geometry	78
6.3	Cavity Finesse	82
6.3.1	Choosing the Finesse	86
6.4	Performance	88
6.4.1	Ramping Efficiency	89
6.5	Looking Forward	92
7	Broadband Modulation	94
7.1	The Ti:sapph Laser	94
7.1.1	Initial Lasing	96
7.1.2	Mode-Locking	98
7.1.3	Beam Characterization	99
7.2	Experimental Design	101
7.2.1	Detection Scheme	103
7.3	Results	108
8	Microresonators: A New Direction	115

8.1	Whispering Gallery Modes	117
8.2	Tapered Fibers	120
8.2.1	Bi-Tapered Fibers	121
8.2.2	Fabrication	124
8.3	Microspheres	127
8.3.1	Experimental Design	130
8.3.1.1	High-Power Coupling	136
9	Microdisks	139
9.1	Detecting Resonances	141
9.2	Searching for Raman Lasing	144
9.3	Future Directions	156
	Bibliography	161

LIST OF FIGURES

1.1	<i>Laser Intensity vs Time (Single Frequency)</i>	4
1.2	<i>Laser Intensity vs Time (3 Frequencies)</i>	5
1.3	<i>Laser Intensity vs Time (5 Frequencies)</i>	6
1.4	<i>Laser Intensity vs Time (201 Frequencies)</i>	7
1.5	<i>Laser Intensity vs Time with Random Phases</i>	8
1.6	<i>Laser Intensity vs Time with Varying Modulation Frequency</i>	9
1.7	<i>An Ideal Optical Modulator</i>	14
2.1	<i>Energy Level Diagram for our Raman System</i>	20
2.2	<i>Rayleigh Scattering</i>	22
2.3	<i>Raman Scattering</i>	22
3.1	<i>The Original Experimental Design</i>	36
3.2	<i>Cavity Transmission vs Frequency</i>	52
3.3	<i>The Derivative of the Cavity Transmission with Respect to Frequency</i>	53
3.4	<i>High-Level Locking Setup Schematic</i>	54
4.1	<i>A Broadband Ro-Vibrational Spectrum</i>	62
5.1	<i>The Mixing Beam Experimental Design</i>	65
5.2	<i>The Mixing Beam Alignment Setup</i>	68
5.3	<i>The Mixing Beam Detection Setup</i>	70
5.4	<i>Conversion Efficiency vs Incident Pump Power</i>	73
6.1	<i>Cavity Output Pump and Stokes Power vs Gas Pressure</i>	76
6.2	<i>Conversion Efficiency vs Gas Pressure</i>	77
6.3	<i>Intensity Product for Various Cavity Geometries</i>	80

6.4	<i>Predicted Mixing Efficiency for Various Cavity Geometries</i>	81
6.5	<i>Original and Redesigned Cavity Mixing Efficiencies</i>	82
6.6	<i>Mixing Efficiency vs Cavity Finesse</i>	83
6.7	<i>Modulation Efficiency vs Cavity Finesse and Input Power</i>	84
6.8	<i>Optimal Cavity Finesse vs Pump Power</i>	85
6.9	<i>Modulation Efficiency vs Pump and Stokes Finesse</i>	87
6.10	<i>Cavity Output Time Dynamics</i>	91
6.11	<i>A Proud Moment</i>	93
7.1	<i>The Ti:sapph Laser</i>	96
7.2	<i>Ti:sapph CW Output Power</i>	97
7.3	<i>Beam Profiles</i>	98
7.4	<i>Ti:sapph Pulse Autocorrelation Signal</i>	101
7.5	<i>Ti:sapph Modulation Experimental Design</i>	102
7.6	<i>The Spectrometer</i>	104
7.7	<i>Spectrum From a 5 Picowatt HeNe Beam</i>	106
7.8	<i>Ti:sapph Power Balancing</i>	108
7.9	<i>Modulated Ti:sapph Spectrum</i>	110
7.10	<i>Ti:sapph Sideband Repetition Rate</i>	111
7.11	<i>Ti:sapph Mixing Efficiency vs Deuterium Pressure</i>	112
7.12	<i>Calculated Ti:sapph Pulse Shapes</i>	113
8.1	<i>Predicted Modulation Efficiency vs Pump Power and Fiber Separation</i>	117
8.2	<i>Coupling Efficiency vs Resonator Translation</i>	124
8.3	<i>A Bi-Tapered Fiber</i>	127
8.4	<i>The Tapered Fiber and Resonator Holder</i>	129
8.5	<i>A Silica Microsphere</i>	130

8.6	<i>Microresonator Experimental Design</i>	131
8.7	<i>Frequency Shift of 1064 nm ECDL vs Applied Voltage to Piezo</i>	133
8.8	<i>Microsphere Resonances</i>	135
8.9	<i>Microsphere Resonances and Fiber Transmitted Signal</i>	136
8.10	<i>Microsphere Coupling Efficiency</i>	138
9.1	<i>An Early Prototype of a Microdisk</i>	140
9.2	<i>A Wedge-Cut Microdisk</i>	142
9.3	<i>A Microdisk Resonance Scan</i>	143
9.4	<i>A Zoomed View of Microdisk Resonances Scan</i>	144
9.5	<i>A Degraded Microdisk</i>	147
9.6	<i>A Substantially Degraded Microdisk</i>	148
9.7	<i>The Repolished Crystal</i>	149
9.8	<i>Raman Lasing from the Microdisk</i>	150
9.9	<i>Raman Lasing from the Microdisk (Narrow Scan)</i>	151
9.10	<i>A Glowing Microdisk</i>	152
9.11	<i>Broad Spectral Scan of the Microdisk Transmitted Light</i>	153
9.12	<i>High-Resolution Spectral Scan of the Microdisk Transmitted Light</i>	155
9.13	<i>Pump to Stokes Conversion Efficiency</i>	156

Abstract

High-Frequency Continuous-Wave Optical Modulation via Stimulated Raman Scattering

David Claiborne Gold

Under the supervision of Professor Deniz D. Yavuz

at the University of Wisconsin-Madison

This thesis describes work towards realizing an efficient, broadband optical modulator at terahertz frequencies. Such a device would have immediate and important implications for atomic physics research and fiber-optic communications, and also serves as the basic building block of the so far elusive arbitrary optical waveform generator. Our modulator is based on the process of continuous-wave molecular modulation, in which stimulated Raman scattering is used to drive coherent oscillations of molecules. We use a high-finesse optical cavity to build up intense laser fields and establish molecular coherence. Using this coherence, we then modulate an independent mixing beam by 90 terahertz, which is the highest rate ever recorded in the continuous-wave regime. We apply this same technique to modulate a pulsed titanium-sapphire laser and show the broadband capability of our modulator and infer its ability to narrow the pulses of the laser. We also describe our recent work to extend our modulation techniques to microresonators, which have the potential for both miniaturization and high modulation efficiency. We demonstrate Raman lasing with one of these microresonators, producing 300 milliwatts of Stokes light, which is two orders of magnitude higher than previously observed in these types of systems.

1 Introduction

For the past few decades, the generation of very fast pulses of light has been a major area of interest in atomic physics. These "ultrashort" pulses, typically defined as shorter than the picosecond (10^{-12} seconds) level, with the shortest being at the femtosecond (10^{-15} seconds) level, have a wide variety of research and industry applications. A useful analogy to understand their utility is to think of the duration of the pulse as like the shutter speed of a camera. In order to sharply record a scene with a camera, there needs to be little movement during the time it takes for the shutter to open and close. The goal is to only capture light during a time period that is shorter than the timescale of any movement in the scene. The same is true when using light to probe physical phenomena—in order to clearly resolve a process, a pulse of light is needed with a duration that is on the same timescale as that process.

When generating ultrashort pulses, it is almost always desirable for the light source to be a laser, and so we will consider it as such hereon. There are a variety of ways one might generate a laser pulse. Perhaps the most obvious would be much like with a camera shutter, and to mechanically block and unblock an otherwise continuous laser beam. Most methods of this sort, though, would be limited to producing pulses down to the millisecond or microsecond range. A more sophisticated option would be to directly turn on or off the laser, for instance by electronically switching the power on and off. However, this kind of approach is limited to the maximum speed that electronic devices can send and respond to signals, which is typically at the nanosecond level.

While the implementation can vary greatly, these two families of mechanical and electronic techniques rely on the same basic principle of attempting to interrupt the propagation

of a laser beam and to only let it pass for an amount of time equal to the desired duration of the pulse. In understanding the basis for improving upon these techniques, it is important to note that as the duration of a pulse, τ , becomes shorter, Fourier synthesis dictates that the bandwidth of the pulse, $\Delta\nu$, must increase. That is, for an ideal bandwidth-limited pulse,

$$\Delta\nu = \frac{C}{\tau} \quad (1.1)$$

where C is a constant that depends on the exact type of pulse, but is close to 1.

So even for the techniques described above that may not seem as if they would affect a laser's bandwidth, they must in fact be broadening it. While this effect will be small for modest pulse lengths, as pulses reach the picosecond level, the bandwidth required becomes an appreciable fraction of the center frequency of the pulse when working in the optical region of the spectrum. (For example, a picosecond pulse will have a bandwidth of $\approx \frac{1}{10^{-12} \text{ sec}} = 1 \text{ THz}$, or about 0.25% of the frequency of the deepest red that humans can see.) This relationship between pulse length and bandwidth provides a clue as to the types of techniques one might use to generate shorter pulses: Rather than focusing on essentially disrupting the propagation of a laser, one might try to more directly affect the bandwidth of the laser itself. And indeed, the shortest pulses are produced through optical techniques that serve to broaden the bandwidth of a laser, such as high harmonic generation, propagation through non-linear optical fibers, mode-locking, or molecular modulation. These techniques have produced pulses at the femtosecond level in the optical region of the spectrum. This work explores the use of molecular modulation in ultrashort pulse generation.

1.1 Pulse Generation

In order to broaden the bandwidth of a laser of frequency ω_0 , one might simply add a second laser at frequency ω'_0 . While the bandwidth of these combined beams is now $|\omega_0 - \omega'_0|$, in general this method will not produce a pulsed laser, or at least not one with pulses anywhere close to the limit implied by Equation 1.1. The issues lie with the fact that these two lasers will not typically be coherent with each other (having a constant phase relationship), and that additional evenly spaced spectral components are needed between the two frequencies.

To examine this further, let us consider the electric field of a laser with frequency ω_0 , which we can describe as

$$E(t) = E_0 \cdot \cos(\omega_0 t + \phi_0). \quad (1.2)$$

We assume no disruptions to the phase over time and choose $t = 0$ such that ϕ_0 is 0, so we can simply write the electric field as

$$E(t) = E_0 \cdot \cos(\omega_0 t). \quad (1.3)$$

The intensity of the beam, which is the relevant quantity when describing pulse width, is proportional to $E(t)^2$ and is plotted below. The full-width-half-max of a single cycle of such a laser is $\pi/2\omega_0$, which we call τ . We plot the time-axis in units of τ for ease of comparison when adding additional frequency components. Note that a single-frequency laser such as this is called "continuous-wave" and is not pulsed, or rather, has an infinite pulse width. To be truly pulsed, a laser needs appreciable "off" time before and after a rise in intensity with a duration that is much longer than the timescale of the pulse. This allows a system to only respond to a single pulse rather than a series of them.

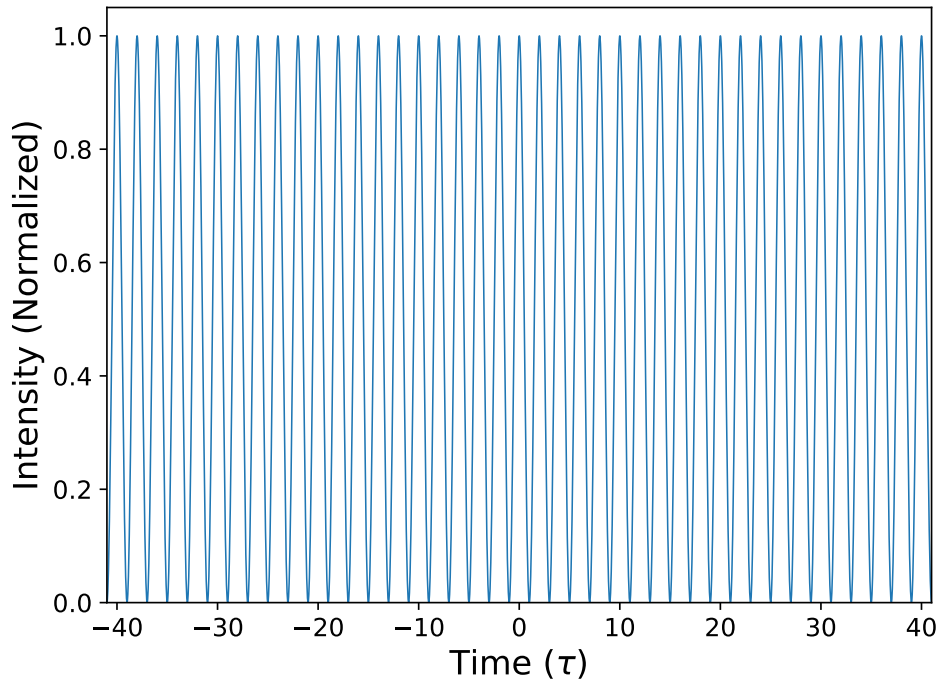


Figure 1.1: *Laser Intensity vs Time (Single Frequency)*

The intensity over time from a single-frequency laser. This is plotted in arbitrary intensity units, since we are only concerned with the relative changes in pulse width, and in time units of τ , the full-width-half-max of a single cycle, since we wish to describe a laser of arbitrary frequency. Note that for this single frequency, there is no "off" time between cycles, and so it is perfectly continuous-wave with no pulsed nature.

We now add two components to the electric field that correspond to two new lasers: one with a frequency greater than the original beam by ω_m and another with a frequency lower by ω_m . We will call ω_m the "modulation frequency", and these new lasers are often referred to as "sidebands"—terms that will be discussed more later. We assume these beams are coherent with the original beam, and so we set $\phi_0 = 0$ for these as well, and assume no phase jumps as before. We set their amplitudes also equal to E_0 for simplicity. The electric field is then

$$E(t) = E_0 [\cos(\omega_0 t) + \cos((\omega_0 + \omega_m)t) + \cos((\omega_0 - \omega_m)t)]. \quad (1.4)$$

We keep ω_m arbitrary in this expression, but for plotting choose it to be $\omega_0/10$. This is a large, but, as we will see later, not at all unrealistic fraction of the original beam frequency. The three beams interfere with each other and create an intensity vs time profile as shown below.




Figure 1.2: *Laser Intensity vs Time (3 Frequencies)*

The intensity over time from three lasers with frequencies spaced by $\omega_m = \frac{\omega_0}{10}$ and centered at ω_0 . The x-axis is in units of $\tau = \frac{\pi}{2\omega_0}$. The full-width-half-max for a given cycle does not change considerably from the single laser case, but the cycles come in clusters now, and begin to take on a pulsed nature.

In general, we can describe the electric field from adding N pairs of lasers to the original laser as

$$E = E_0 \cdot \cos(\omega_0 t) + E_0 \cdot \sum_{n=1}^N \left[\cos((\omega_0 + n\omega_m)t) + \cos((\omega_0 - n\omega_m)t) \right], \quad (1.5)$$

where the value of n is referred to as the "order" of the sideband or of the modulation. Plots for $N = 2$ (5 total components) and $N = 100$ (201 total components) are shown below.




Figure 1.3: *Laser Intensity vs Time (5 Frequencies)*

At 5 frequency components, the laser appears much more pulsed. Note that while the intensity is normalized in the plot, the peak intensity is proportional to the square of the number of frequency components. So while the average power is 5 times that of the original single frequency laser, the peak power is 25 times higher: Essentially all the power from the additional components gets compressed into short periods of time.




Figure 1.4: *Laser Intensity vs Time (201 Frequencies)*

At 201 frequency components, the laser is now almost perfectly pulsed with long periods of nearly 0 intensity between short bursts of high intensity. The spacing of these bursts is $\frac{2\pi}{\omega_m}$. Because the full bandwidth of the laser is now much larger than the center frequency, the width of the pulses is now less than 0.2τ —a "sub-cycle" pulse.

These have of course been idealized examples, but demonstrate the principle of pulse generation through introducing additional spectral components.

We next depart from these idealized cases to show the importance of generating equally spaced, coherent sidebands. We first allow the phases of the sidebands to vary randomly, again for the $N = 100$ case. This models a combination of lasers with no phase jumps but no set phase relationship between one another, or, alternatively, approximately mimics a system where the individual lasers may begin in phase but then experience frequent phase jumps.




Figure 1.5: *Laser Intensity vs Time with Random Phases*

For the $N = 100$ case, we let the phase of each added sideband vary randomly between 0 and 2π . This completely removes the pulsed nature of the waveform and it now contains many random intensity fluctuations.

The laser intensity now has large fluctuations and does not appear pulsed. The relatively low number of frequency components means the system is subject to considerable intensity noise from the random phases. Adding more components would reduce the relative variation and the intensity would approach a constant value. Although the intensity is plotted normalized so that the peak is still 1, the peak power is much lower than in the coherent case and the high-power of the pulses are now essentially spread out continuously over time. Clearly, using multiple lasers with no constant phase relationship, even with consistent frequency spacing, will not work for generating pulses. This demonstrates why common broadband sources, such as light bulbs, do not generate pulses of light through the interference of their spectral components.

We next reset the initial phases to 0 to make the sidebands once again coherent, but now allow ω_m to vary randomly by up to $\pm\frac{1}{2}\%$. The laser is relatively unperturbed close to $t = 0$, but after just a few pulse cycles (each of length $T = \frac{2\pi}{\omega_m}$), the phases of the frequency components are effectively scrambled relative to one another and the intensity looks again like the system in Figure 1.5. Several cycles are shown in the plot below.




Figure 1.6: *Laser Intensity vs Time with Varying Modulation Frequency*

At $t = 0$ when the pulse with perturbed sideband spacing is generated, the maximums of the components still line up in time. However, as time goes on, the components become misaligned relative to one another. With no common frequency multiple between them, it only takes a few cycles before the intensity turns to noise, much like when the phases were randomly varied.

So in order to generate consistent pulses of light by combining lasers, it seems that it is necessary for the lasers to have both a constant phase relationship (to be "in phase" with one another) and a precise frequency spacing.

1.1.1 Methods of Ultrashort Pulse Generation

There are a variety of ways of generating ultrashort pulses. Mode-locking lasers are a common approach, in which some broadband gain medium is made to lase at many different frequencies spaced by the laser cavity resonances. Several different active or passive schemes can ensure that these modes lase in phase with one another and produce short pulses in the manner described in the previous section. Ti:sapph lasers in particular have become highly developed in recent years and commercial systems that produce 4 femtosecond pulses are now available [1]. Another approach is high harmonic generation, which has produced even shorter pulses, well into the attosecond range [2]. This technique uses an already pulsed laser as a pump (often a Ti:sapph laser) and focuses it onto a target. Many of the atoms in the target are nearly ionized, and the high kinetic energy of the electrons when they return to the atoms can produce very high harmonics of the pulse. It is these high harmonics which give the required bandwidth to produce attosecond pulses, but also means that much of the light is very far into the UV or even x-ray regions of the spectrum, and so is limited to only certain applications. Free electron lasers, on the other hand, generate light through the interactions of electrons with magnetic fields, and so are not limited to atomic resonances, and can operate anywhere over a very wide wavelength range. The minimum pulse lengths are at the femtosecond level [3], but the lasers themselves tend to be very large and complex, and so are not well suited for many research labs. Techniques involving non-linear fibers have also been promising. Generally, a pulsed laser source is used as a pump and is further spectrally broadened through non-linear interactions in these fibers. Near single-cycle pulses have been generated using these techniques in the NIR region of the spectrum [4]. Another promising technique is molecular modulation, which is the main technique utilized in our research. Details of this method and our specific implementation of it are discussed in Chapter 2 and Chapter 3.

1.1.2 Arbitrary Waveform Generation

The previous sections described how combining different lasers with a fixed frequency spacing can be used to generate ultrashort pulses, and described some of the techniques used to produce these different laser frequencies. While pulse generation is the main goal that we focus on in our research, a pulse is really just one specific type of intensity vs time function that can be produced when coherent spectral components are combined. More generally, pulse generation falls under the domain of an "optical arbitrary waveform generator"—a theoretical device which aims to produce intensity waveforms of any shape. Such a device could have any number of uses. For instance, a square-wave that goes from 0 intensity to its maximum intensity nearly instantaneously might be more useful for examining the response of atoms to high intensity light than the sinusoidal shape of typical pulses. More exotic waveforms might be used to put an atomic system into a very specific quantum state (i.e. coherent and quantum control). There is also significant interest in arbitrary waveform generator for applications to precision and time-resolved spectroscopy. And there are surely many industry applications that would rapidly develop—for instance, perhaps pulses used for data transmission could be produced with shapes that pre-compensate for dispersion that would otherwise accumulate as they propagate through a fiber.

Arbitrary waveform generators do already exist for electrical signals. But as is typical for electrical devices, they are usually limited to operating at the GHz level. So while this allows them to be useful when interfacing with other electronic devices, they are less useful for engaging with the timescales of atoms and light, or for pushing data transmission rates to new levels. But although the signals for electronic and optical arbitrary waveform generators might be produced in different ways, the general operating principle in either case is the same direct application of Fourier synthesis.

The foundational principle of Fourier synthesis is that any periodic function $x(t)$, can be represented as an infinite sum of sines and cosines of harmonic frequencies with weighted amplitudes. (Well, *almost* any function. Functions can be defined that would not have a Fourier representation, but these would not exist outside the realm of mathematics. For all practical purposes, Fourier synthesis can be used to represent any type of electrical or optical signal we might want.) Mathematically, this is described as

$$x(t) = a_0 + \sum_{k=1}^{\infty} \left[a_k \cos(k\omega t) + b_k \sin(k\omega t) \right]. \quad (1.6)$$

So how would we generate an optical waveform described by $x(t)$? Well, we need to be able to produce a laser that corresponds to every term in the sum. To begin, first recall that we previously described the electric field from a laser in Equation 1.2 as

$$E(t) = E_0 \cdot \cos(\omega_0 t + \phi_0).$$

Starting with this laser and with an appropriate optical modulator, we set ω in Equation 1.6 as the frequency of our modulator, ω_m . We then would use our modulator to generate a given frequency term, $k\omega_m$, in $x(t)$ by an appropriate choice of modulation-order n , such that

$$k\omega_m = \omega_0 + n\omega_m. \quad (1.7)$$

Since all the frequency components in $x(t)$ differ by just a multiple of ω_m , choosing a laser with frequency ω_0 that satisfies Equation 1.7 for one value of k will work for all other values. Or alternatively, if the frequency of the modulator is adjustable, we may fix the frequency of our laser, ω_0 . Then, picking ω_m that satisfies Equation 1.7 for one value of k , will work for all others. Additionally, because sine and cosine differ from each other just by a phase factor of $\pi/2$, each term in Equation 1.6 can be generated from a single laser

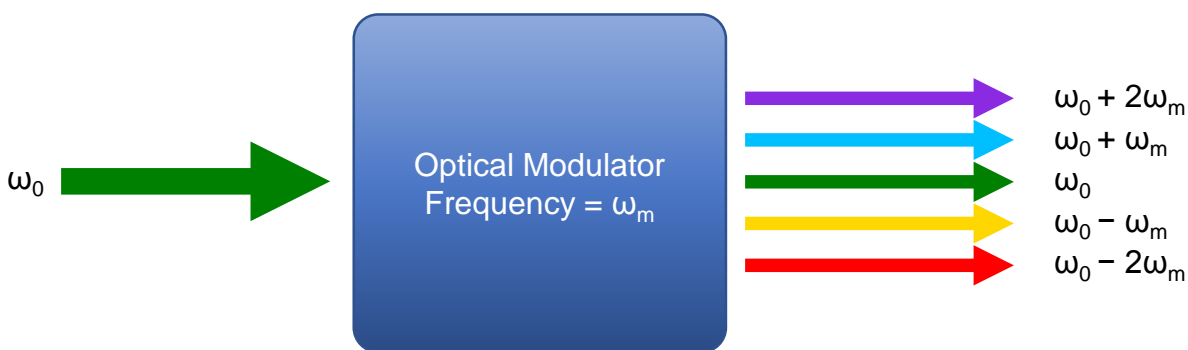
described by Equation 1.2 with appropriate phase and amplitude control. Note that if $x(t)$ is either an even or odd function, we can represent it with a combination of only sines or cosines, and so only amplitude control of the frequency components would be needed.

In practice, the resolution and accuracy of an optical arbitrary waveform generator would be limited by the frequency and stability of ω_0 and ω_m , the number of sidebands that can be produced, and overall level of control over the amplitude and phase of components. Producing many orders of coherent sidebands at a high frequency ω_m is one of the key challenges in creating an optical arbitrary waveform generator.

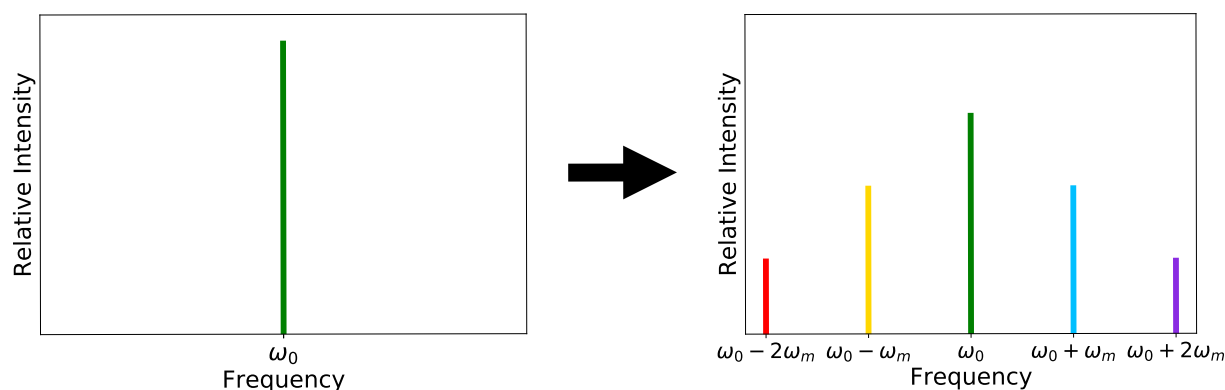
1.2 An Optical Modulator

The basic building block of an optical arbitrary waveform generator would be an optical modulator: A device that takes input light of a given frequency and outputs light of the same frequency along with one or more orders of sidebands shifted by multiples of the modulation frequency of the device. One of the long-term goals of this experiment, and of the work described in this thesis, is to produce a high-frequency, high-efficiency optical modulator that is able to operate anywhere within the optical region of the spectrum. A diagram of such a device is shown in Figure 1.7.

Optical modulators like in this diagram already exist and are in fact commonplace in atomic physics experiments, most often coming in the form of electro-optic modulators (EOMs) or acousto-optic modulators (AOMs). However, these devices have a maximum modulation rate of about 100 GHz, with rates typically being closer to 1 GHz, with the main limitation being the driving electronics [5, 6]. So while these have a wide variety of experimental applications, they are not fast enough to be of direct use in ultrashort pulse generation or arbitrary waveform generation.



(a) The Device



(b) In Frequency Space

Figure 1.7: *An Ideal Optical Modulator*

An ideal optical modulator takes an input laser at some frequency ω_0 and shifts a large fraction of the laser power to new frequencies, spaced from ω_0 by integer multiples of the modulation frequency ω_m (first and second-order modulation is shown here). In frequency space, this modulation puts sidebands on the original beam.

1.2.1 Molecular Modulation

Our work has been devoted to creating a high-frequency (10s of THz) optical modulator based on molecular modulation. Molecular modulation relies on Raman scattering, in which a photon exchanges energy with a molecule to shift to either a higher or lower frequency. This is in contrast to Rayleigh scattering, in which only the direction of propagation of a photon is affected. In Raman scattering, the molecules typically exchange

vibrational and/or rotational energy with photons and move back and forth between their ground state and an excited state. An effective molecular modulator, then, is one that has a large collection of molecules all vibrating or rotating in unison. When a laser beam passes through this region of oscillation, it will have its frequency modulated both up and down, imparting sidebands on the original beam with frequencies $\omega_0 \pm \omega_m$ where, as before, ω_0 is the frequency of the original laser and ω_m is the modulation frequency.

To understand how these molecules are set vibrating and how this creates modulation, we briefly examine two-level quantum mechanical systems. Recall that quantum mechanics dictates that when a particle is in some linear combination of the two eigenstates of a system, the probability of finding it in a new state (or staying in the original state) will have a sinusoidally oscillating component with a frequency proportional to the energy difference between the two eigenstates. (A particle that is entirely in an eigenstate of a system is in a stationary state and will not exhibit time dynamics.) Mathematically, we can write the initial state of the particle as $\psi_0 = \psi(t=0)$, where $\psi(t) = c_{a,0} |a\rangle e^{-i\omega_a t} + c_{b,0} |b\rangle e^{-i\omega_b t}$. Here $|a\rangle$ and $|b\rangle$ are the eigenstates of the system, and $\omega_a = E_a/\hbar$ and $\omega_b = E_b/\hbar$.

The probability of the particle being found in this same initial state at a later time t is then

$$\begin{aligned}
 P_{|\psi_0\rangle}(t) &= \left| \langle \psi_0 | \psi(t) \rangle \right|^2 \\
 &= \left| (c_{a,0} \langle a| + c_{b,0} \langle b|) (c_{a,0} |a\rangle e^{-i\omega_a t} + c_{b,0} |b\rangle e^{-i\omega_b t}) \right|^2 \\
 &= \left| c_{a,0}^2 e^{-i\omega_a t} + c_{b,0}^2 e^{-i\omega_b t} \right|^2 \\
 &= c_{a,0}^4 + c_{b,0}^4 + 2c_{a,0}^2 c_{b,0}^2 \cos[(\omega_b - \omega_a)t].
 \end{aligned} \tag{1.8}$$

Using the half angle identity for $\sin(\frac{x}{2})$, along with the fact that normalization requires $c_a^2 + c_b^2 = 1$, lets us rewrite this more cleanly as

$$P_{|\psi_0\rangle}(t) = 1 - 4c_{a,0}^2 c_{b,0}^2 \sin^2 \left[\frac{(\omega_b - \omega_a)t}{2} \right]. \tag{1.9}$$

Oscillations in $P_{|\psi_0\rangle}(t)$ are now readily apparent. At $t = 0$, the probability of the particle being found in its initial state is 1 (as should be expected). The probability then oscillates in time and will reach 0 if $c_{a,0} = c_{b,0} = \sqrt{2}/2$ (i.e. when the initial state is equally probable to collapse into either of the two eigenstates). The oscillation present here, though, is simply in the relative phase between the two terms in the $\psi(t)$, as $c_a(t)$ and $c_b(t)$ are static. So the probability of being found in $|a\rangle$ or $|b\rangle$ does not change over time. This means that just having a collection of particles with two energy levels is not going to cause them to physically oscillate or to form an optical modulator, which is perhaps an unsurprising fact.

If, however, a particle is in the presence of a driving oscillating potential (such as the electric field of a laser), the probability of finding the particle in $|a\rangle$ or $|b\rangle$ oscillates through the process of Rabi flopping. This is more complicated than in the previous example without a driving field, and will be explored more in Chapter 2. The main result, though, is that probability of a particle in an initial state $|\psi_0\rangle$ later being found in one of the two eigenstates, say, $|a\rangle$, oscillates in time at the generalized Rabi frequency. Mathematically we can write this as

$$P_{|a\rangle}(t) = \left| \langle a | \psi(t) \rangle \right|^2 = c_a(t)^2 \quad (1.10)$$

and

$$c_a(t)^2 \propto \left(\sin^2(\tilde{\Omega}t/2 + \phi) \right), \quad (1.11)$$

with $\tilde{\Omega} = \sqrt{\Omega^2 + \Delta^2}$ as the generalized Rabi frequency. Ω depends on the dipole moment between the two states and the driving potential amplitude; Δ is the detuning between the driving potential frequency, ω_0 , and the transition frequency of the states, $\omega_b - \omega_a$. In the types of systems we will work with, the detuning is very very large, in which case $\tilde{\Omega} \approx \Delta$. For our particle starting in $|\psi_0\rangle$ then, the probabilities for being found in $|a\rangle$ or in $|b\rangle$ will each have oscillatory terms at frequencies $\Delta_a = \omega_0 - \omega_a$ and $\Delta_b = \omega_0 - \omega_b$. The difference in these frequency terms is $\omega_b - \omega_a$, or the transition frequency between the two states.

This can be interpreted then as a particle that is in initial state $|\psi_0\rangle$ having an oscillatory probability in time of being found in either $|a\rangle$ or $|b\rangle$, and in the case that the particle interacts frequently with the environment, actually collapsing back and forth between $|a\rangle$ and $|b\rangle$. This oscillatory movement between $|a\rangle$ and $|b\rangle$ is significant for applications to molecular modulation because, for most molecules, these two states will have slightly different indices of refraction. Semi-classically then, this creates a time-varying index of refraction with an oscillation frequency of $\omega_b - \omega_a$, which we will call ω_m .

Now, we want to see how this will affect a laser of arbitrary frequency ω_0 , so we will write its electric field as

$$E(t) = E_0 \cdot e^{i(\omega_0 t + \phi)}, \quad (1.12)$$

where this time we write the field as a complex exponential for mathematical convenience. Next, we describe the index of refraction in the modulator as a static part (the average refractive index of the molecules) and a much smaller oscillating component as

$$n(t) = n_0 + \delta_n \cdot \sin(\omega_m t), \quad (1.13)$$

so that $\delta_n \ll n_0$. Because the speed of light in a medium is inversely proportional to the index of refraction, and because the phase delay that light accumulates within a medium is proportional to its velocity, the phase delay will take the same form as the index of refraction:

$$\phi(t) = \phi_0 + \delta_\phi \cdot \sin(\omega_m t). \quad (1.14)$$

We then substitute this expression for the phase in for the original phase term ϕ in Equation 1.12 to get

$$E(t) = E_0 \cdot e^{i(\omega_0 t + \phi_0 + \delta_\phi \cdot \sin(\omega_m t))}. \quad (1.15)$$

We can simplify the expression by absorbing the constant component of the phase, ϕ_0 , into

E_0 and separating the time-varying terms into frequency and phase components to get

$$E(t) = E'_0 \cdot e^{i\omega_0 t} \cdot e^{i\delta_\phi \cdot \sin(\omega_m t)}. \quad (1.16)$$

To see how this time-varying phase affects the laser, we can perform a Taylor expansion (given that δ_ϕ is very small). So this term becomes

$$e^{i\delta_\phi \cdot \sin(\omega_m t)} \approx 1 + i\delta_\phi \cdot \sin(\omega_m t). \quad (1.17)$$

Rewriting the sine term as complex exponentials will allow us to more easily combine it with the laser frequency term:

$$1 + i\delta_\phi \cdot \sin(\omega_m t) = 1 + \frac{\delta_\phi}{2}(e^{i\omega_m t} - e^{-i\omega_m t}). \quad (1.18)$$

Now, substituting this back in for the phase term into Equation 1.16 gives

$$E(t) \approx E'_0 e^{i\omega_0 t} \left(1 + \frac{\delta_\phi}{2}(e^{i\omega_m t} - e^{-i\omega_m t}) \right), \quad (1.19)$$

which we can write as

$$E(t) \approx E'_0 \left(e^{i\omega_0 t} + \frac{\delta_\phi}{2} e^{i(\omega_0 + \omega_m)t} - \frac{\delta_\phi}{2} e^{-i(\omega_0 - \omega_m)t} \right). \quad (1.20)$$

It is now clear in Equation 1.20 that in addition to the original electric field oscillating at the laser frequency of ω_0 , there are also frequencies of $\omega_0 \pm \omega_m$, with an amplitude reduced by a factor of $\frac{\delta_\phi}{2}$. The frequency downshifted sideband is called the Stokes beam and the frequency upshifted sideband is the anti-Stokes beam. Including the next term in the Taylor expansion would yield second-order sidebands with frequencies $\omega_0 \pm 2\omega_m$ and

with amplitudes reduced by $\frac{\delta\phi^2}{4}$, and so on for additional terms. In practice, there will be many more factors that affect the amplitudes between Stokes and anti-Stokes beams and between different orders, but this derivation gives the basics of how modulation works.

We have briefly covered a lot of topics in this chapter, so to summarize, there are a variety of scientific and industry applications that make the ability to generate ultrashort pulses desirable. Generating these pulses by adding spectral components to a laser seems to be a fruitful method, and having control over spectral components in this way implies the additional ability to produce arbitrary waveforms and wavelengths of light, not just pulses. A high-frequency optical modulator would be the perfect tool to enable this, and this work explores the development of such a device based on the principle of molecular modulation. Molecular modulation works through generating a time-varying index of refraction in a material, and so much of the experimental work described here involves the development of techniques to generate such a property.

2 Molecular Modulation

In this chapter, we attempt a more rigorous analysis of the physical processes involved in molecular modulation. Most of our experiments used molecular deuterium gas as the modulation medium housed inside of an optical cavity, although we try to keep most of this analysis fairly general. We begin by looking at an energy level diagram for a single molecule in this system, shown in Figure 2.1 below.

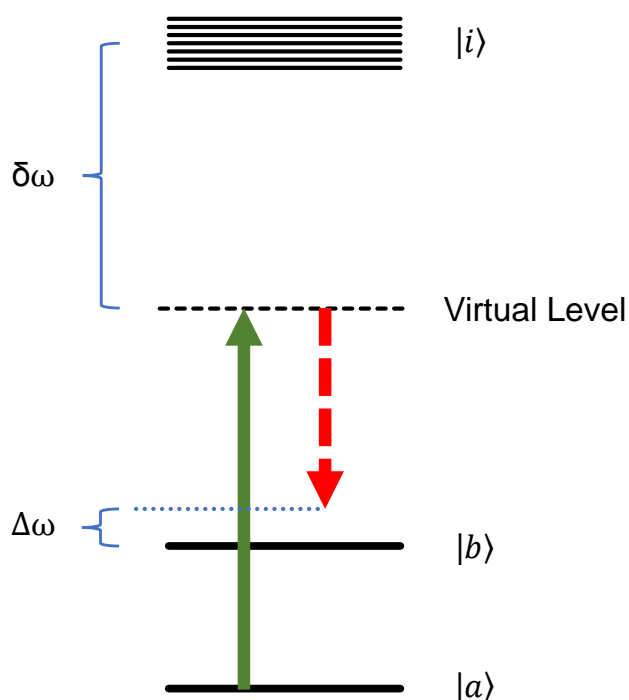
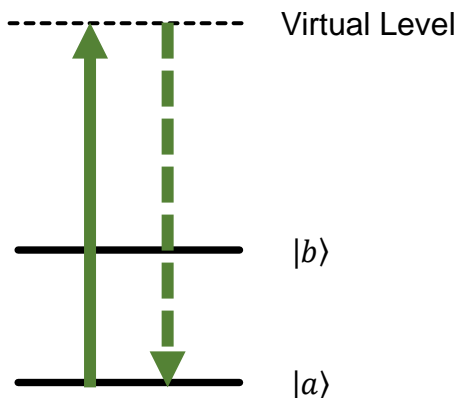


Figure 2.1: *Energy Level Diagram for our Raman System*

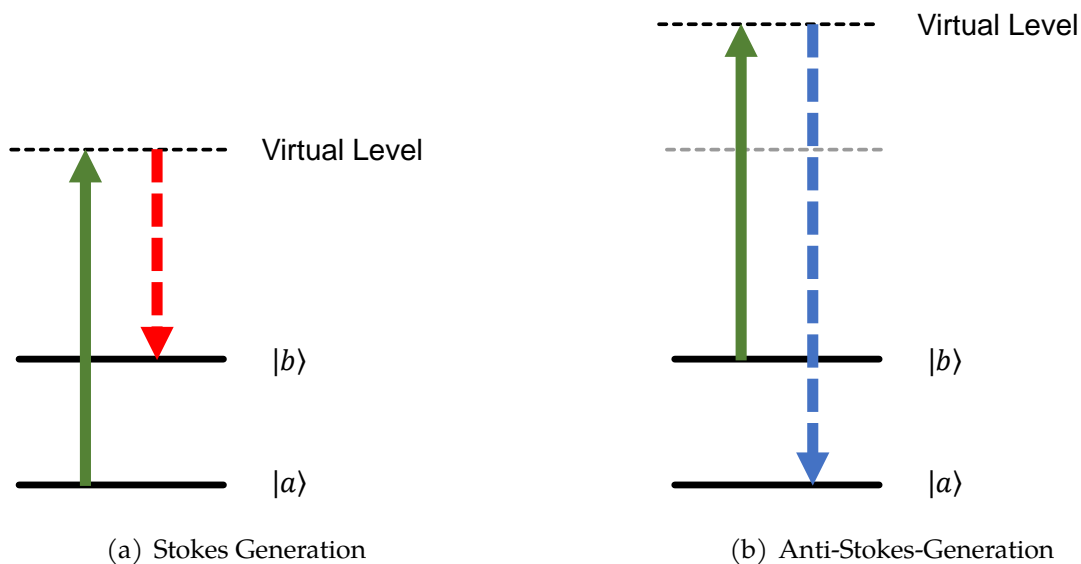
$|a\rangle$ and $|b\rangle$ are the ground and first excited vibrational level within the ground electronic state of the molecule. $|i\rangle$ refers to any member of a group of states in the first excited electronic state. When a pump laser (green) hits the molecule, it excites it to a "virtual" level, which is very far detuned from any of the $|i\rangle$ states (the diagram is not to scale). As $\omega_i \gg \omega_a, \omega_b$, we call the detuning between the pump and any of the $|i\rangle$ states $\delta\omega$. $\Delta\omega$ is the two-photon detuning.

Here $|a\rangle$ and $|b\rangle$ are the ground and first excited vibrational level within the ground electronic state of the molecule. There are generally rotational levels within each vibrational level as well, but it is usually possible to suppress transitions involving changes in rotational levels, so we assume both $|a\rangle$ and $|b\rangle$ are in the ground rotational state. In deuterium, $|a\rangle$ and $|b\rangle$ are spaced by 89.6 THz, which is a large fraction of the frequency of visible light. $|i\rangle$ refers to any member of a group of states in the first excited electronic state. These states are much higher energy than $|a\rangle$ or $|b\rangle$ and the pump laser.

Transitions directly between $|a\rangle$ and $|b\rangle$ are dipole forbidden in symmetric molecules such as deuterium. Therefore, when a pump laser of frequency ω_0 (shown in green) hits a molecule in the ground state $|a\rangle$, the dominant process is for the molecule to excite to a "virtual" state, which is just very far detuned from any of the $|i\rangle$ states. As $\omega_i \gg \omega_{a(b)}$, we call the detuning between the pump and any of the $|i\rangle$ states $\delta\omega$. Often, the molecule decays from the virtual level back to $|a\rangle$ almost immediately with no change in photon energy, which is the process of Rayleigh scattering that was mentioned in Chapter 1 and is now shown below in Figure 2.2. It is possible, though somewhat less probable, for the molecule to decay to $|b\rangle$ instead of $|a\rangle$. In this case, the emitted photon will have frequency ω_{-1} , which is reduced from the pump frequency by $\approx \omega_b - \omega_a$, and is called a Stokes photon. This is the process of Raman scattering and is shown in Figure 2.3(a). The difference between the frequency of the emitted photon and the energy level spacing is called the two-photon detuning and is given by $\Delta\omega = (\omega_b - \omega_a) - (\omega_0 - \omega_{-1})$. In the absence of interactions with other photons, the molecule will decay back to $|a\rangle$ from $|b\rangle$ non-radiatively. Another Raman scattering process, shown in Figure 2.3(b) can occur in which the molecule is initially in $|b\rangle$. In this case, when hit by a pump photon, the molecule excites to a (different) virtual level and then decays to $|a\rangle$. The emitted photon is now called an anti-Stokes photon and has a higher frequency than the pump frequency by $\omega_b - \omega_a$.

Figure 2.2: *Rayleigh Scattering*

In Rayleigh scattering, a pump photon excites a molecule to a virtual level. The molecule quickly decays back to its original state and a photon is emitted with no change in energy.

Figure 2.3: *Raman Scattering*

Raman scattering occurs when a molecule is excited to a virtual level by absorbing a photon, and then decays to a different energy level. If the molecule begins in $|a\rangle$, the emitted photon is called a Stokes photon and has a lower frequency by $\omega_b - \omega_a$. Alternatively, if the molecule begins in $|b\rangle$ and decays to $|a\rangle$, an anti-Stokes photon is produced with a frequency increased by $\omega_b - \omega_a$.

2.1 Formalism

With an overview of the energy levels of the system and the basic processes that can occur in mind, we next attempt to develop some formalism for the system. Much of this derivation is based on the original work of Harris and Sokolov [7], as well as the theses of previous students in our research group Josh Weber and Tyler Green [8, 9].

We begin with the time-dependent Schrödinger equation,

$$i\hbar \frac{d}{dt} |\psi\rangle = \hat{\mathcal{H}} |\psi\rangle. \quad (2.1)$$

As our modulation medium is a gas, we assume that it is diffuse enough that the molecules do not significantly interact with each other. We can then consider the wave function of a single isolated molecule, which will take the form

$$|\psi(t)\rangle = c_a(t)e^{-i\omega_a t} |a\rangle + c_b(t)e^{-i\omega_b t} |b\rangle + \sum_n c_{i_n}(t)e^{-i\omega_{i_n} t} |i_n\rangle. \quad (2.2)$$

The index n is over all the $|i\rangle$ states. However, as these states are all so far detuned from $|a\rangle$ and $|b\rangle$, have such a small population (almost certainly 0, in fact, excluding virtual levels) and we will never work with any specific $|i\rangle$ state, we just call any of their frequencies ω_i and drop the n subscript on those terms. Instead, in somewhat loose notation, we will write the summation index as over i . Next, we wish to find the Hamiltonian for one of these molecules, and we work in the interaction picture so that we can write the Hamiltonian as the sum of an unperturbed component, $\hat{\mathcal{H}}_0$, and an interaction component, $\hat{\mathcal{H}}_I$, so that

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_0 + \hat{\mathcal{H}}_I. \quad (2.3)$$

Given the wave function of Equation 2.2, the unperturbed Hamiltonian will be

$$\hat{\mathcal{H}}_0 = \hbar\omega_a |a\rangle \langle a| + \hbar\omega_b |b\rangle \langle b| + \sum_i \hbar\omega_i |i\rangle \langle i|. \quad (2.4)$$

For the interaction component, we assume that the only significant contribution from the laser will be in the electric dipole interaction. Recall though that in the systems we use in our experiment, transitions between $|a\rangle$ and $|b\rangle$ are dipole forbidden. We also neglect coupling terms between only the upper states $|i\rangle$, since, as mentioned, these are completely un-populated (see subsection 3.3.4 for a discussion of the relative population levels of our system). The only contributions, then, will be from coupling between $|a\rangle, |i\rangle$ and $|b\rangle, |i\rangle$, and we can write the interaction component as

$$\begin{aligned} \hat{\mathcal{H}}_I = & -E(t) \sum_i \mu_{a,i} |a\rangle \langle i| - E(t) \sum_i \mu_{b,i} |b\rangle \langle i| \\ & - E^*(t) \sum_i \mu_{a,i}^* |i\rangle \langle a| - E^*(t) \sum_i \mu_{b,i}^* |i\rangle \langle b|. \end{aligned} \quad (2.5)$$

In this equation, μ is the matrix element between two states from the electric dipole moment operator and $E(t)$ is the total electric field from any lasers present. Initially there is only a pump laser, but since we know the system will generate sidebands, we write the electric field as

$$E(t) = \sum_n E_n \cdot e^{i(\omega_n t - k_n z)}. \quad (2.6)$$

We also write the electric field in complex form for mathematical convenience. Recall that here n is the modulation-order and may be any integer. Additionally, we have included spatial dependence with the $k_n z$ term, where k_n is the wavenumber for a given modulation-order. For now, we assume homogeneity in the x and y directions so the electric field is a

plane wave. The full Hamiltonian can then be written as

$$\hat{\mathcal{H}} = -\frac{\hbar}{2} \begin{bmatrix} A & B \\ C & D - 2\Delta\omega \end{bmatrix}, \quad (2.7)$$

where

$$\begin{aligned} A &= \sum_n a_n |E_n|^2 \\ B &= \sum_n b_n E_n E_{n-1}^* \\ C &= B^* \\ D &= \sum_n d_n |E_n|^2. \end{aligned} \quad (2.8)$$

For convenience, we define the following constants, where a_n and d_n are generally called the dispersion constants, and b_n is called the coupling constant:

$$\begin{aligned} a_n &= \frac{1}{2\hbar^2} \sum_i |\mu_{a,i}|^2 \left[\frac{1}{(\omega_i - \omega_a) - \omega_n} + \frac{1}{(\omega_i - \omega_a) + \omega_n} \right] \\ d_n &= \frac{1}{2\hbar^2} \sum_i |\mu_{b,i}|^2 \left[\frac{1}{(\omega_i - \omega_b) - \omega_n} + \frac{1}{(\omega_i - \omega_b) + \omega_n} \right] \\ b_n &= \frac{1}{2\hbar^2} \sum_i \mu_{a,i} \mu_{b,i}^* \left[\frac{1}{(\omega_i - \omega_a) - \omega_n} + \frac{1}{(\omega_i - \omega_b) + \omega_n} \right]. \end{aligned} \quad (2.9)$$

In arriving at this expression for the Hamiltonian, we have used the rotating-wave approximation and assumed negligible transition linewidths.

In the basis of the unperturbed system (without the electric fields from lasers), the eigenstates of the full Hamiltonian, $|+\rangle$ and $|-\rangle$, come out to be

$$|\pm\rangle = 2B|a\rangle - \left[(A - D + 2\Delta\omega) \mp \sqrt{(A - D + 2\Delta\omega)^2 + 4B^2} \right] |b\rangle, \quad (2.10)$$

and have corresponding eigenvalues

$$E_{\pm} = -\frac{\hbar}{4} \left[(A + D - 2\Delta\omega) \mp \sqrt{(A - D + 2\Delta\omega + 4B^2)} \right]. \quad (2.11)$$

So far we have focused on a single molecule in our system, but as we wish to see how an ensemble of particles interacts with the pump laser and sidebands, we now move to the density matrix formulation. To this end, it is convenient to first express these eigenstates in a simpler form. For a two-level system like this, we only need to describe one of the eigenstates, since the probability of occupation of the other is defined by normalization. We focus on the $|-\rangle$ state and write it as

$$|-\rangle = \cos\left(\frac{\theta}{2}\right) e^{i\phi/2} |a\rangle + \sin\left(\frac{\theta}{2}\right) e^{-i\phi/2} |b\rangle, \quad (2.12)$$

where

$$\theta = \tan^{-1}\left(\frac{2|B|}{A - D + 2\Delta\omega}\right) \quad (2.13)$$

and ϕ is a phase factor so that $B = |B|e^{i\phi}$.

Now, with the density matrix as

$$\rho = \begin{bmatrix} \rho_{a,a} & \rho_{a,b} \\ \rho_{a,b}^* & \rho_{b,b} \end{bmatrix} = \begin{bmatrix} |c_a|^2 & c_a c_b^* \\ c_a^* c_b & |c_b|^2 \end{bmatrix}, \quad (2.14)$$

and using the wave function as defined in Equation 2.12, we have:

$$\rho_{a,a} = \cos^2\left(\frac{\theta}{2}\right) \quad (2.15)$$

$$\rho_{b,b} = \sin^2\left(\frac{\theta}{2}\right) \quad (2.16)$$

$$\text{and } \rho_{a,b} = \frac{1}{2} \sin(\theta) e^{i\phi}. \quad (2.17)$$

We assume the population of $|b\rangle$ stays low (i.e. there is weak coupling), in which case θ is small so that $\theta \approx \frac{2|B|}{A-D+2\Delta\omega}$ and $\sin(\theta) \approx \theta$. Then we can write the density matrix elements in a somewhat simplified manner using the original matrix elements of the full Hamiltonian. We get

$$\rho_{a,a} = 1 - |\rho_{a,b}|^2, \quad (2.18)$$

$$\rho_{b,b} = |\rho_{a,b}|^2, \quad (2.19)$$

$$\text{and } \rho_{a,b} = \frac{B}{A - D + 2\Delta\omega}. \quad (2.20)$$

$\rho_{a,a}$ and $\rho_{b,b}$ are the fraction of the molecular population in $|a\rangle$ and $|b\rangle$ respectively. $\rho_{a,b}$ can be understood as the "coherence" of the system, and describes the amount of coupling between the different states of the system through the pump and any generated electric fields. The coherence can have a maximum value of 0.5, which would occur when $\rho_{a,a} = \rho_{b,b} = 0.5$. In the language of our semi-classical description in subsection 1.2.1, this would correspond to the entire system oscillating between $|a\rangle$ and $|b\rangle$, giving a maximally varying index of refraction.

With expressions for the density matrix elements in hand, we can now consider how the pump laser and any generated sidebands will propagate through the system. The lasers can be well described by the Helmholtz equation in a non-homogenous medium (i.e., a

medium with a driven field that modifies the polarizability), so we have

$$\nabla^2 E_n = 2ik_n \frac{\partial E_n}{\partial z} - \mu_0 \omega_n^2 P_n, \quad (2.21)$$

with P_n as the polarizability of the medium at frequency ω_n that is associated with the n th sideband of the electric field. We still assume electric fields of the form given in Equation 2.6, so that $\frac{\partial^2 E}{\partial x^2} = \frac{\partial^2 E}{\partial y^2} = 0$. Additionally, we take the paraxial approximation to the Helmholtz equation so that $\frac{\partial^2 E}{\partial z^2} \approx 0$. Then we have

$$2ik_n \frac{\partial E_n}{\partial z} = -\mu_0 \omega_n^2 P_n, \quad (2.22)$$

which, by defining $\eta = \sqrt{\frac{\mu_0}{\epsilon_0}}$, we can write as

$$\frac{\partial E_n}{\partial z} = \frac{-i\eta}{2} \omega_n P_n. \quad (2.23)$$

Essentially, we are now looking at the propagation of the pump and sideband electric fields over small distances of our system where they can be fully treated as planar. The effects of beam size and shape are explored in Section 6.2. Additionally, recall that the Helmholtz equation gives time-independent solutions, or rather separable solutions, to the wave equation. So Equation 2.23 has no explicit time-dependence and E_n and P_n represent the steady-state amplitude of the electric field and polarizability at position z in the system. However, the net quantities will still oscillate in time with frequency components at ω_n , and it is in fact this oscillation in the polarizability that leads to coupling between different E_n components in the electric field and allows for the creation of sidebands.

Moving on from Equation 2.23, our analysis unfortunately becomes a rather tedious and very extensive exercise in algebraic manipulation, so we will just loosely sketch the

process. The curious reader is invited to reference Josh Weber and Tyler Green's theses for very detailed derivations [8, 9]. For an arguably unnecessarily detailed derivation, see Josh Weber's Appendix A. To continue with our approach though, the next step in our analysis is to evaluate the total polarizability, P , as a function of time, and express it as a sum of oscillating terms with frequencies ω_n and amplitudes P_n . By plugging in each P_n into Equation 2.23, we can get an equation describing the propagation of the n th electric field component.

In the time domain, the polarizability of a medium is given by

$$P(t) = \langle \psi(t) | \hat{P} | \psi(t) \rangle, \quad (2.24)$$

where $|\psi(t)\rangle$ is as defined by Equation 2.2 and \hat{P} is

$$\hat{P} = \sum_i \mu_{a,i} |a\rangle \langle i| + \sum_i \mu_{b,i} |b\rangle \langle i| + \sum_i \mu_{a,i}^* |i\rangle \langle a| + \sum_i \mu_{b,i}^* |i\rangle \langle b|. \quad (2.25)$$

To evaluate this, we must find expressions for $c_a(t)$, $c_b(t)$, and $c_i(t)$ that are present in $|\psi(t)\rangle$, which we had managed to avoid up until now by moving to the density matrix formulation. (It is these terms that will introduce expressions involving sums of E_n into the equation.) We can then arrange the sums and indices and, after a great deal of algebra, write $P(t)$ as

$$P(t) = \frac{1}{2} \sum_n P_n e^{i(\omega_n t - k_n z)} + \frac{1}{2} \sum_n P_n^* e^{-i(\omega_n t - k_n z)}. \quad (2.26)$$

Here P_n is given by

$$P_n = 2\hbar \left[|c_a|^2 E_n a_n + |c_b|^2 E_n d_n + c_a^* c_b E_{n+1} b_{n+1} e^{-i(\Delta\omega t + k_m z)} + c_a c_b^* E_{n-1} b_n^* e^{i(\Delta\omega t + k_m z)} \right], \quad (2.27)$$

with $k_m = \frac{\omega_b - \omega_a}{c}$, or the wavenumber associated with the spacing of the two energy levels that set the modulation frequency. We can now plug this expression for P_n back into Equation 2.23. Additionally, to simplify the expression, we define $\tilde{c}_a = c_a$ and $\tilde{c}_b = c_b e^{-i(\Delta\omega t + k_m z)}$. This gives us,

$$\frac{\partial E_n}{\partial z} = -i\eta\hbar\omega_n [|\tilde{c}_a|^2 E_n a_n + |\tilde{c}_b|^2 E_n d_n + \tilde{c}_a^* \tilde{c}_b E_{n+1} b_{n+1} + \tilde{c}_a \tilde{c}_b^* E_{n-1} b_n^*]. \quad (2.28)$$

Next, we multiply by the molecular density, N , so that we are now working with the macroscopic electric fields and polarizability. We also substitute in expressions for the density matrix elements from Equation 2.14 to further simplify the equation. Finally, we arrive at a useful expression for how the amplitude of the n th sideband will change in space.

$$\boxed{\frac{\partial E_n}{\partial z} = -i\eta\hbar\omega_n N [\rho_{a,a} E_n a_n + \rho_{b,b} E_n d_n + \rho_{a,b}^* E_{n+1} b_{n+1} + \rho_{a,b} E_{n-1} b_n^*]} \quad (2.29)$$

Notice that the propagation of the electric field of the n th sideband depends on the n th-1 and n th+1 fields, but also that the presence of an n th field allows for the creation of n th-1 and n th+1 fields that were not initially present. So starting with a pump laser, which gives the $n = 0$ field, one can then see how the Stokes ($n = -1$) and anti-Stokes ($n = 1$) fields would be generated; this is the process of Raman scattering. The dispersion constants are set by the transition dipole moment elements and detunings of the system, and so are largely fixed. The strength of this coupling effect between the electric fields, therefore, is mainly determined by the molecular density, electric field strengths, and coherence of the system.

Another important feature shown by this equation is that coherence is not tied to any one laser beam. So one laser or set of lasers can be used to establish a molecular coherence

and appropriate population levels (all the ρ terms), and then completely separate lasers can interact or "mix" with the molecules and coherence. This is our primary interest in this system, and is the basis for our optical modulator: After preparing the system using a pump laser, any other laser needs merely to pass through the region of molecular coherence and it will be modulated. We investigate this further below.

2.2 Mixing Efficiency

We next see how Equation 2.29 applies to the specific system we work with. For our system, the modulation medium is a molecular gas, usually deuterium, housed inside an optical cavity. The cavity has high-finesse at both the pump wavelength and the Stokes wavelength, which is 89.6 THz less than the pump and set by the vibrational energy level spacing of the gas. When the pump laser enters the cavity and builds up to high intensity, Stokes light is generated through the dynamics previously described. The Stokes light resonates and builds up to higher intensities, which in turn increases the coherence of the molecules and generates more Stokes light. The system quickly reaches steady-state powers. We now examine how a separate mixing beam will interact with this system.

Consider a separate mixing beam with electric field E'_0 with modulation order denoted as E'_n . We assume this mixing beam is weak enough that its effects on the coherence and population levels are negligible (for most of our experiments, the mixing beam power in the cavity is three to four orders of magnitude lower than the pump beam). We also assume that the modulation efficiency will be low enough that the mixing beam is not appreciably depleted and can be taken as constant. For a given mixing beam sideband then, the only significant terms will be ones involving the carrier beam. So for the anti-Stokes sideband

(the Stokes sideband will work out similarly), Equation 2.29 would become

$$\frac{\partial E_1'}{\partial z} = -i\eta\hbar\omega_1'N[\rho_{a,a}E_1'a'd_1 + \rho_{b,b}E_1'd_1' + \rho_{a,b}^*E_2'b_2' + \rho_{a,b}E_0'b_1'^*] \quad (2.30)$$

$$\approx -i\eta\hbar\omega_1'N[\rho_{a,b}E_0'b_1'^*]. \quad (2.31)$$

Now, calling $\alpha = \eta\hbar\omega_1b_1'^*$ and imposing the boundary condition that $E_1'(z = 0) = 0$ (there is no sideband incident on the system), we get a solution for $E_1'(z)$ of

$$E_1'(z) \approx -i\alpha N\rho_{a,b}E_0'z. \quad (2.32)$$

We are generally interested in the modulation efficiency, or the ratio of incident mixing beam power to the generated power in a given order, so we square both sides of the equation and solve for the ratio of the electric field terms.

$$\epsilon = \frac{E_1'^2(z)}{E_0'^2} \approx \alpha^2 N^2 \rho_{a,b}^2 z^2 \quad (2.33)$$

We could in principle use Equation 2.33 to predict the modulation efficiency of our system. However, these are not very convenient variables to use when working in the lab. Additionally, the modulation efficiency tends to be low, and so we are really most interested in how it scales in the regime of low coherence to provide guidance as to how to modify the experiment. So we replace the molecular density N with the gas pressure p . We also rename the propagation length z as the cavity length l . Next, consider $\rho_{a,b}$ as given by Equation 2.20. Applying this to our system where everything but the $n = 0$ and $n = -1$ terms are negligible, we get

$$\rho_{a,b} \approx \frac{b_0 E_0 E_{-1}^*}{\alpha_0 |E_0|^2 + \alpha_{-1} |E_{-1}|^2 - d_0 |E_0|^2 - d_{-1} |E_{-1}|^2 + 2\Delta\omega} \quad (2.34)$$

Still considering the regime of low coherence, the $2\Delta\omega$ term will dominate in the denominator (the dispersion and coupling constants are all approximately proportional to the inverse of the single-photon detuning from the upper $|i\rangle$ states and so are very small). Then we can say

$$\rho_{a,b} \approx \frac{b_0 E_0 E_{-1}^*}{2\Delta\omega}. \quad (2.35)$$

Finally, making these changes, as well as assuming a constant two-photon detuning and dropping the proportionality constants, we arrive at

$$\epsilon \propto p^2 \times I_p \times I_S \times l^2 \quad (2.36)$$

Here, I_p is the intensity of the pump beam in the cavity and I_S is the intensity of the Stokes beam in the cavity.

For the benefit of the reader who has understandably skimmed this chapter, we now summarize the major assumptions we have made in arriving at Equation 2.36. First, some initial properties of the system: First we have assumed that our modulation medium is a diffuse enough gas that interactions between molecules are negligible. We then only look at the electric dipole interaction between the electric field and molecules and ignore interactions that are between only the $|i\rangle$ states, as these are essentially un-populated. The only interactions present then are between $|a\rangle, |i\rangle$ and $|b\rangle, |i\rangle$, and we assume the two-photon detuning is small compared to the transition and laser frequencies, and so the rotating wave approximation is valid. We also assume negligible transition linewidths.

Based on our available pump laser power, we have also assumed that the coherence, $\rho_{a,b}$, of the system is relatively low. We would need to adjust these calculations if we achieve a coherence above about 0.1 or so, but we would be thrilled to have to do so. A

related assumption is that the modulation efficiency is low enough that a separate mixing beam is not appreciably depleted in power when it is modulated. This also means that the only significant coupling terms for a mixing beam will be ones involving the carrier beam E'_0 . Additionally, this mixing beam is assumed to be weak enough that it has a negligible impact on the coherence of the system (note that its single-photon detuning will be the same order of magnitude as that of the pump for any optical wavelength, and its intensity will be three to four orders of magnitude lower than the pump in the cavity).

Finally, we have assumed the shape of the electric field to be planar, and have not accounted for the optical cavity beyond its effect on the magnitude of various electric fields due to its reflectivity. This means that our results are most valid over small propagation distances in the cavity, or for cavities in which the mirrors have large radii of curvature compared to the length of the cavity. These assumptions related to the beam shapes are relaxed in Chapter 6 and explored using a somewhat different formalism.

So having reached the guiding results of Equation 2.29 and Equation 2.36 that describe sideband generation, and with disclaimers to our formalism fully disclosed, we next examine how to create this system in the lab in Chapter 3.

3 Experimental Design

When I began working on this experiment and taking over for previous student Josh Weber in 2014, it was configured to use linearly polarized light from a continuous-wave 1064 nm fiber amplifier as the driving or "pump" beam. Deuterium gas was used as the modulation medium and was housed in an optical cavity. The main features are shown in Figure 3.1. This figure, as well as other optical layout figures in this document, do not typically show all the optical components present in the experimental implementation, as this would be impractical. Additionally, the specific arrangement of components is sometimes changed for clarity. Every effort has been made to ensure that the main operating principles are properly represented and can be determined from the diagrams.




Figure 3.1: *The Original Experimental Design*

The initial layout of the experiment when I began work on the project. Only the major components are shown. Light is produced at 1064 nm from an ECDL. It then passes through an isolator to block back-reflections from reaching the diode, followed by an EOM which puts 50 MHz sidebands on the laser so that it can be locked to a cavity resonance. The laser is then amplified to up to 20 watts and is shaped and aligned to the cavity. A photodiode detects reflected light from the cavity, which is used as feedback to the laser to lock it. A second isolator prevents the large amount of reflected light from reaching the fiber amplifier. The transmitted light from the cavity is picked-off, and bandpass filters separate the transmitted pump and Stokes beams to be observed on photodiodes. Some graphics in this and other optics layout figures in this thesis are adapted from *ComponentLibrary* by Alexander Frazen under a Creative Commons license [10].

3.1 Original Experimental Design Overview

We will first go through a high-level overview of the experiment and then examine some of the components in more details. The pump laser is produced by the 1064 nm diode shown in the upper left of the diagram. This provides tunable single-frequency light, which is next sent through an optical isolator to prevent back-reflections from perturbing the diode. The laser then passes through an EOM driven at 50 MHz, which creates frequency sidebands at about 5% efficiency. Several polarizing beam cubes and a Glan-Taylor polarizer (not shown in Figure 3.1) make for a high purity, linearly polarized beam. About 10 mW of light reaches the fiber amplifier to seed it. The spatial-mode and beam-shape of the amplifier output vary somewhat with the output power setting. To circumvent any issues this might cause, we typically set it to its maximum power of 30 W and then adjust the output power to our desired level with a polarizing beam cube and a half-wave plate, while sending the rest of the power to a beam dump (all not shown). As the beam approaches the cavity, various optics shape it to mode-match it to the cavity. Details of the cavity and of mode-matching are discussed in Section 3.3. When the laser is off-resonant from the cavity, nearly the entire beam will be back-reflected, so high damage threshold optics are in place after the fiber amplifier, and a second optical isolator is used to protect the amplifier.

During operation, the laser frequency is scanned using the piezo on the ECDL. When the laser frequency passes through a resonance of the optical cavity, the transmitted light is detected on a photodiode. Any Stokes light that is generated is separated using dielectric filters and detected as well. The laser can be locked to the cavity by using the reflected signal from the cavity when the EOM is running. In this case, additional processing is done to the reflected signal through a standard Pound-Drever-Hall technique to generate an error signal. High-frequency components are sent to the laser current controller and low-frequency components are sent to the laser piezo, which helps keep the laser on-resonance.

3.2 The ECDL

The pump beam is produced by an external cavity diode laser (ECDL) that is set in a Littrow configuration. This ECDL housing was designed in our lab, and serves to reduce the linewidth of the laser, make it lase in a single longitudinal-mode, and to provide a mechanism for tuning the wavelength of that mode. In a Littrow configuration, light from a collimated laser diode is directed onto a diffraction grating, such that the first-order diffracted light is sent back to the diode as optical feedback and the zeroth-order reflected light is used as the laser output.

The grating forms an optical cavity with the inner facet of the laser diode. This serves to add additional frequency restrictions on the laser diode, since output light from the whole ECDL now must be resonant with the diode-grating cavity. Additionally, since the grating spreads light out spatially into the diffracted orders as a function of frequency, careful adjustments to alignment can preferentially send back a narrow range of frequencies that are resonant with only a single mode of the cavity. As this light can cause stimulated emission within the diode, the overall effect is to enhance gain at a narrow band of frequencies, while suppressing it at others. This results in a narrow linewidth, single-mode laser. The grating is mounted on a piezoelectric transducer, which can slightly adjust the cavity length and resonant frequencies, thus tuning the output frequency of the laser. Broader tuning adjustments can be made by rotating the angle of the diffraction grating using a kinematic mount.

Lasers in the Littrow configuration often use a grating that is blazed to maximize that amount of light in the first diffracted order, and then use a geometry such that the blazed angle is equal to the incident angle (and so the diffraction angle too) to maximize optical feedback to the diode. The geometry of our ECDL housing requires the incident angle, or "Littrow angle", to be close to 45° , so in practice this means one would choose a grating

that has a line spacing necessary to reflect the first-order diffracted light back to the diode.

So starting with a grating with blaze angle $\theta_b = 45^\circ$, we then write the grating equation as

$$d(\sin \alpha + \sin \beta) = m\lambda, \quad (3.1)$$

with d as the grating line spacing, α as the incidence angle, β as the diffraction angle, and m as the diffracted order. We then set $\alpha = \beta = 45^\circ$ and $m = 1$ as per the Littrow configuration. Substituting in these values and solving for d gives

$$\begin{aligned} d(\sin 45^\circ + \sin 45^\circ) &= \lambda \\ d &= \frac{\sqrt{2}}{2}\lambda. \end{aligned} \quad (3.2)$$

For $\lambda = 1064$ nm, this gives a line spacing of $d = 752$ nm, or 1330 lines per mm.

In practice, it is not easy to obtain a grating with this exact line spacing, but this turns out not to matter. A variety of factors affect how much optical feedback a diode needs to perform well, but we find that for the 1064 nm diodes we use (QLD-1060-50S-AR from QPhotonics), only a small amount of optical feedback is needed. Directly sending the center of the diffracted order back to the diode would in fact give too much feedback. We instead use a grating with 1200 lines per mm, which has a reflected angle of 34.7° when light is incident at 45° . This results in much less light being reflected back on the diode. From here, small adjustments to the incidence angle are made to optimize the feedback and performance of the diode.

Our pump laser typically achieves a linewidth of ≈ 300 kHz, with a total tuning range of 10 nm and a continuous tuning range (i.e. without mode-hops) of 10 GHz. The maximum output power is 30 mW.

3.3 The Optical Cavity

The heart of the experiment is the optical cavity, which is where the modulation medium is housed (typically deuterium gas) and where the high laser intensities needed for modulation build up. The cavity mirrors are custom made by Precision Photonics with a radius of curvature of 100 cm, and with a high reflectivity at both the pump wavelength of 1064 nm and at the Stokes wavelength of 1555 nm (1064 nm downshifted by 90 THz). The pump reflectivity was measured by the manufacturer as $R_p = 0.999868$ and the Stokes reflectivity as $R_s = 0.999862$. Our own measurements showed good agreement with these values. The finesse of a cavity, \mathcal{F} , is defined as

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1 - R}, \quad (3.3)$$

which gives a finesse of approximately 23,000 at both wavelengths. The mirrors are relatively transparent in the rest of the optical region of the spectrum.

The cavity is 75 cm long, which gives it a free spectral range, ν_{FSR} , of

$$\nu_{\text{FSR}} = \frac{c}{2l} = 200 \text{ MHz}. \quad (3.4)$$

This is a fairly long cavity, which presents challenges for locking a laser to it. First, this makes the cavity more susceptible to mechanical perturbations that can disrupt the lock, and also increases the sensitivity to misalignment of the input beam. Second, this fundamentally narrows the inherent linewidth of the cavity resonances, which is given by

$$\nu_{\text{cavity}} = \frac{\nu_{\text{FSR}}}{\mathcal{F}}. \quad (3.5)$$

Equation 3.5 gives a linewidth of about 10 kHz at both the pump and the Stokes wavelengths, which is much narrower than the linewidth of the laser and means that significant feedback

is needed to achieve efficient coupling. The cavity was designed with this length because, despite these challenges, having a longer interaction length should increase the modulation efficiency as shown in Equation 2.36. This trade-off is examined more thoroughly in Section 6.2.

3.3.1 A Brief Aside

The reader may wonder why the cavity is not also reflective at the anti-Stokes wavelengths, since Equation 2.29 suggests that also having a strong anti-Stokes field could greatly increase the coherence of the system. There are several reasons for this. First, we found that it is very difficult and expensive to source mirrors that are highly reflective at three different wavelengths over such a wide range. Additionally, the mirrors we have been able to get are usually reflective over a fairly wide range of about 100 nm around each target wavelength. While this is not very problematic at the pump and the Stokes wavelengths, this would be for the anti-Stokes wavelength, which is at approximately 807 nm. Having the cavity be highly reflective there would block a large region of the spectrum that we would hope for our modulator to operate at. In the future, with more advanced mirrors, the experiment could be improved with a cavity that is highly reflective in a very narrow range around the pump, Stokes, and anti-Stokes wavelengths, and even second-order sidebands, while maintaining high transparency in all other regions of the spectrum. This would not only increase the modulation efficiency, but would allow us to use a lower gas pressure and improve the ease of operation of the experiment.

Another natural question to arise at this point is "Why should we bother with an optical cavity at all?". Using a pulsed laser source as the pump could result in very high peak intensities of the pump and generated Stokes beam, which could lead to very high modulation efficiencies since that quantity scales as the product of the intensities. This

could eliminate the need for an optical cavity and greatly simplify the experimental setup. There are several reasons we choose a CW laser for the pump. First, although generating short pulses of light by modulating an already pulsed source is one of our end goals, using both a pulsed pump and a pulsed mixing beam would not work well—they would have to have very close repetition rates and pulse widths, or else quickly become out of sync with one another. More generally, a pulsed pump is not suitable for use in an arbitrary waveform generator, as it cannot produce waveforms with a longer duration than its own pulse width. Even its usefulness as a modulator to make individual frequency components is limited: Not only would any frequencies produced be limited to the duty cycle of the pump beam, they would also have a fundamental limit on how narrow they can be, as governed by Equation 1.1, which greatly limits their utility. Applications to precision spectroscopy, for instance, rely heavily on having a narrow linewidth source. This broadband nature inherent to pulsed sources also increases the uncertainty in the modulation frequency of the system. Modulation schemes using a pulsed pump source are most suitable for self-modulation for the purpose of generating narrow pulses, and our aim is for more general utility. There are few additional comments I'd like to make on the subject of pulsed lasers, but to keep this discussion civil, we will end the section here.

3.3.2 Cavity Mode-Shape

An ideal optical cavity uses lossless mirrors (i.e. all incident light is either transmitted or reflected), and in frequency space has narrow bands of 100% transmission spaced by the free spectral range, with near 0% transmission between these resonances. The circulating power in the cavity when a laser is on-resonance is given by $P_0(1 - R)^{-1}$, with P_0 as the incident power. In reality of course, many factors move a system away from this ideal case. As previously stated, the inherent linewidth of the cavity will limit the amount of

laser light coupled without electronic feedback, essentially acting as an extremely narrow bandpass filter. Another major factor is whether the shape of the laser beam matches (or is "mode-matched" to) the shape supported by the cavity. This eigenmode is the shape of laser beam that perfectly reflects back upon itself within the cavity and so will be unchanged as it propagates back and forth. While a cavity will generally support an infinite number of modes with intensity profiles described by Hermite polynomials, it is almost always desirable to couple to the lowest-order or "fundamental" Gaussian mode. This mode has the standard radially symmetric Gaussian intensity profile of the typical ideal laser beam, and has the highest peak intensity and smallest spatial extent of the modes. This is the mode we will discuss from here on. Additionally, this discussion applies primarily to linear cavities, which consist of two parallel mirrors.

In order to mode-match the pump laser to the cavity and couple light in, it is necessary to know the eigenmode of the cavity. To this end, we first examine the shape and propagation of laser beams. The propagation of most laser beams can be well described by Gaussian optics. For a beam propagating in the z direction, the intensity profile is given by

$$I(r, z) = I_0 \left(\frac{w_0}{w(z)} \right)^2 \exp \left(\frac{-2r^2}{w(z)^2} \right). \quad (3.6)$$

In this equation, I_0 is the peak intensity of the beam; w_0 is the beam waist, or the radius of the narrowest cross-section of the beam; $w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R} \right)^2}$ is the radius of the beam at a given point a distance z away from the waist; and $z_R = \pi w_0^2 \lambda^{-1}$ is the Rayleigh range and sets the length scale for the divergence of the beam. Another important quantity of the beam is the radius of curvature of the wavefront at position z , which is given by

$$R(z) = z \left[1 + \left(\frac{z_R}{z} \right)^2 \right]. \quad (3.7)$$

To find the mode-shape of a cavity, it can be helpful to use the complex beam parameter formulation, where the quantity $q(z)$ is defined as

$$q(z) = z + iz_R. \quad (3.8)$$

The real part of $q(z)$ is the distance z from the waist of the beam, and the imaginary part is the Rayleigh range, which contains the necessary information to calculate the spot size and radius of curvature of the beam at position z . Propagation through free space or through most optical elements can be described using a 2×2 transfer matrix, S , of the form

$$S = \begin{pmatrix} A & B \\ C & D \end{pmatrix}. \quad (3.9)$$

The transformation from the initial to final complex beam parameter, q_i to q_f , is described by

$$\begin{pmatrix} q_f \\ 1 \end{pmatrix} = k \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} q_i \\ 1 \end{pmatrix}, \quad (3.10)$$

with k as a normalization constant that is set so as to keep the second component of the vectors equal to 1. We can then calculate q_f as

$$q_f = \frac{Aq_i + B}{Cq_i + D}. \quad (3.11)$$

The effect of multiple optical elements in sequence can be found through simple matrix multiplication of their transfer matrices. This allows for the easy computation of their cumulative effect on a laser beam as it propagates through a system.

To find the cavity-mode, we need to first find the transfer matrix that corresponds to one round trip of our cavity. Then we find the complex beam parameter that is an eigenmode

of this matrix so that $q_f = q_i$ (meaning the beam is unchanged after a round trip). We will take the $z = 0$ reference point to be the center of the cavity.

If we consider a cavity of length l and with symmetric mirrors of radius of curvature R , one round trip will consist of free space travel of distance $l/2$, reflection by a mirror, free space travel of distance l , reflection by the other mirror, and free space travel of distance $l/2$. The cavity transfer matrix, M_C , is then

$$M_C = M_{l/2} M_R M_l M_R M_{l/2}. \quad (3.12)$$

Substituting in the appropriate transfer matrix for each element, we get

$$M_C = \begin{pmatrix} 1 & l/2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -2/R & 1 \end{pmatrix} \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -2/R & 1 \end{pmatrix} \begin{pmatrix} 1 & l/2 \\ 0 & 1 \end{pmatrix},$$

which evaluates to

$$M_C = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \quad (3.13)$$

with

$$\begin{aligned} A &= 1 + \frac{2l^2 - 4lR}{R^2} \\ B &= 2l + \frac{l^3 - 3l^2R}{R^2} \\ C &= \frac{4l - 2R}{R^2} \\ D &= A. \end{aligned}$$

(Note that A and D are not equal when the two mirrors do not have the same radius of curvature.) Additionally, with the condition that $q_f = q_i$, Equation 3.11 only has a solution

when $|A + D| < 2$, or, equivalently, when

$$-1 < \frac{l^2 - 2lR}{R^2} < 0. \quad (3.14)$$

This provides a stability condition for the geometry of a symmetric linear cavity. (For our cavity with $l = 0.75$ m and $R = 1.00$ m, the middle quantity of this equation is -0.9375 .) Solving Equation 3.11 with $q_f = q_i$ and using the ray transfer matrix of Equation 3.13 gives

$$q = \frac{il}{2},$$

which corresponds to a beam waist of

$$\omega_0 = \frac{l\lambda}{2\pi} \approx 405 \text{ microns}$$

located at $z = 0$, or the center of the cavity.

3.3.3 Mode-Matching

In order for the beam to couple well to the cavity, the beam shape must be approximately the eigenmode of the cavity, or "mode-matched" to the cavity. Note that for a laser beam of a given wavelength, the size and position of the waist fully define the shape of the beam. So in practice, mode-matching can be achieved by ensuring that the size and position of the beam's waist are the same as that of the cavity-mode. In our case, we need our laser to focus at the middle of our cavity to a diameter of 810 microns.

While a single lens should be enough to appropriately shape the beam to a cavity, in general, because lenses are not readily available in arbitrary focal lengths, it may require several lenses to do so. Past group members have used different techniques to determine

an appropriate combination of focal lengths and positions of lenses, to varying effect. The following description applies to the techniques I developed and used.

We first measure the spot size of the beam in several known positions using a beam profiler camera. This allows us to fit the beam shape to an intensity distribution described by Equation 3.6 and to determine the complex beam parameter of our pump beam, and then to calculate the beam shape inside the cavity (this shape in the cavity cannot be directly measured with the camera as the cavity is enclosed). The level of mode-matching can then be quantified by calculating the overlap integral between the inferred beam shape in the cavity and the cavity-mode. From here, a relatively simple program written in Python lets us input the focal lengths of lenses that we have available to use, and will output several possible configurations that should achieve good mode-matching. It determines this by simply iterating through all possible configurations of the available lenses in discrete position intervals along the beam path. At each configuration, it uses the ray matrix formulation to quickly calculate the resulting beam parameter in the cavity. It then ranks these configurations by the value of their overlap integral.

Typically, by starting with a reasonably collimated pump beam and allowing two lenses of commonly available focal lengths from 50 to 1000 mm to be placed anywhere over about 1 meter of path length, we will find a number of possible configurations with mode-matching greater than 95% (i.e. the calculations say that less than 5% of the pump beam should be rejected from the cavity due to spatial-mode mismatch). After choosing and setting up one of these configurations, small adjustments are typically needed to optimize the coupling, but we find that this procedure gives a very good starting point.

3.3.4 Cavity Cooling

For molecular modulation using either the vibrational or rotational states of deuterium, we need the molecules to begin in their ground state. The pump beam can then excite them to a far detuned upper state before they emit Stokes photons as they decay down to the first excited state. This process repeats and coherence is built between the ground and excited states. From the Boltzmann distribution, we can determine the ratio of molecules in the ground state to excited states. We first look at the vibrational states and start with the Boltzmann distribution for an arbitrary state i , which gives the probability of a molecule being in state i (compared to being in any other state):

$$p_i = \frac{e^{-\epsilon_i/kT}}{\sum_{j=0}^{\infty} e^{-\epsilon_j/kT}} \quad (3.15)$$

Setting i as the ground vibrational state and summing over the vibrational energy levels v (with frequency spacings close to 90 THz for the first few levels) and choosing T as a typical room temperature of 300 K, gives

$$p_{v=0} \approx 1 - (5.9 \cdot 10^{-8}),$$

so the system is almost completely in the ground vibrational state.

We repeat this calculation for the rotational levels, but note that the energy spacings are much smaller, from just a couple THz up to a few tens of THz. We now find that

$$p_{J=0} \approx 0.14,$$

so only about 14% of the molecules will be in the $|v = 0, J = 0\rangle$ state and available to build coherence.

To help address the population distribution issue, the optical cavity was designed with a reservoir around the main chamber that can hold several liters of liquid nitrogen to cool the gas. Once equilibrium is reached, the gas temperature is approximately 77 K. The effect on the vibrational level distribution is insignificant (note that while the ratio of molecules in the $|v = 0\rangle$ to excited levels changes substantially, only the actual number of molecules in the $|v = 0\rangle$ state is important here), but has a large effect on the rotational levels. At 77 K, we find that $p_{J=0} \approx 0.46$, so overall there are about 3.29 times more molecules in the ground state and available to form a coherently oscillating system.

Cooling with liquid nitrogen has additional benefits as well. The linewidth of the Raman transition used in the modulator is most strongly affected by Doppler broadening and pressure broadening. Doppler broadening occurs due to the increased width of the velocity distribution of particles as temperature is increased. This effectively broadens the transition's linewidth from the perspective of the laser, or broadens the laser's linewidth from the perspective of the ensemble of particles. The spread of velocities is described by the Maxwell distribution, which gives a \sqrt{T} dependence for the width.

Pressure broadening occurs due to collisions between molecules, which can shorten the natural decay time of a molecule in an excited state (thus increasing the uncertainty in the energy of an emitted photon, or the width of the emission line). Molecular collisions occur at a rate proportional to the velocity of the particles, which the Maxwell distribution also says will depend on \sqrt{T} . So Doppler broadening and pressure broadening both affect the linewidth of the Raman transition like

$$\Delta\nu_D, \Delta\nu_P \propto \nu_R \sqrt{T}, \quad (3.16)$$

with $\Delta\nu_D$ as the Doppler linewidth, $\Delta\nu_P$ as the pressure broadening linewidth, and ν_R as the frequency of the Raman transition. With a temperature change from 300 K to 77 K, we

get a linewidth decrease for both for both $\Delta\nu_D$ and $\Delta\nu_P$ by a factor of

$$\sqrt{\frac{300\text{K}}{77\text{K}}} \approx 1.97.$$

A literature review gives the Doppler and pressure linewidths for the first vibrational transition in deuterium at 300 K and 1 atm as approximately 556 MHz and 175 MHz respectively [11, 12]. Lowering the temperature to 77 K will then narrow these to 282 MHz and 89 MHz. The inhomogeneous Doppler profile is described by a Gaussian, and the homogeneous pressure profile is described by a Lorentzian. The overall effect is given by the convolution of these distributions, which is a Voigt profile.

A Voigt profile's linewidth, σ_V , can be well approximated by the equation $\sigma_V = 0.5346\sigma_L + \sqrt{0.2166\sigma_L^2 + \sigma_G^2}$, with σ_L as the width of the Lorentzian and σ_G as the width of the Gaussian. Applying this equation to our system, we can find the ratio of the overall linewidth of the Raman transition at the two temperatures as

$$\frac{\Delta\nu_{R,300\text{K}}}{\Delta\nu_{R,77\text{K}}} = \frac{0.5346\nu_{P,300\text{K}}^2 + \sqrt{(0.2166(\nu_{P,300\text{K}})^2 + (\nu_{D,300\text{K}})^2)}}{0.5346\nu_{P,77\text{K}}^2 + \sqrt{(0.2166(\nu_{P,77\text{K}})^2 + (\nu_{D,77\text{K}})^2)}} \approx \frac{655\text{ MHz}}{333\text{ MHz}}.$$

This gives an overall narrowing by a factor of $\frac{655}{333} \approx 1.97$, meaning that we probably could have just treated all the linewidths as Gaussian. Since the laser and cavity linewidths are much smaller than the Raman linewidth at either of these temperatures, this narrowing effectively corresponds to an increase in the Raman-gain by a factor of 1.97.

3.4 Laser Locking

In order for high intensities of pump light to build up within the cavity, the laser frequency must be adjusted so that it is resonant with the cavity. As is often the case with high-finesse

optical cavities, our laser has a much larger linewidth than the cavity resonance. In our experiment, the laser has a linewidth of about 300 kHz on short timescales and the cavity has a linewidth of about 10 kHz. Additionally, thermal and mechanical fluctuations can cause the center frequencies of both of these resonances to drift by 10s of MHz or more over a period of minutes. This means that in order to couple significant amounts of light, the laser linewidth needs to be narrowed, as well as the center frequency continuously adjusted to keep it on-resonance.

The solution to these issues in any laser system is generally to lock the laser to a stable reference frequency or to the target resonator itself. We use a standard Pound-Drever-Hall (PDH) technique to lock our pump laser to the optical cavity. The original paper by Drever and Hall gives a very accessible explanation of the technique [13], and past students in the Yavuz Group, Josh Weber and Tyler Green, both give detailed descriptions of our specific setup in their dissertations [8, 9], so this discussion will be limited to a brief overview. A guide aimed towards someone setting up a locking system like this for the first time given by Fox, Oates, and Hollberg, is very useful as well [14].

To start, let's say that a laser is completely on-resonance with a cavity, and so transmission is at a maximum. The transmission as a function of frequency will be as shown in Figure 3.2, with narrow resonances spaced by the free spectral range of the cavity. At this point, any drift in frequency, either of the laser or of the cavity resonance, will cause the transmission to drop. Due to the symmetric nature of the cavity resonance shape, we cannot simply adjust the laser frequency to fix this since we do not know whether it is now too high or too low. To combat this with the PDH technique, frequency sidebands are put on the laser before it enters the cavity. In our case, this is done via an EOM operating at 50 MHz. This effectively allows the system to sample the resonance at both a higher and lower frequency. If the laser is not perfectly resonant, one of these sidebands will have a higher transmission than the other and indicates the direction that the laser frequency

needs to be shifted.

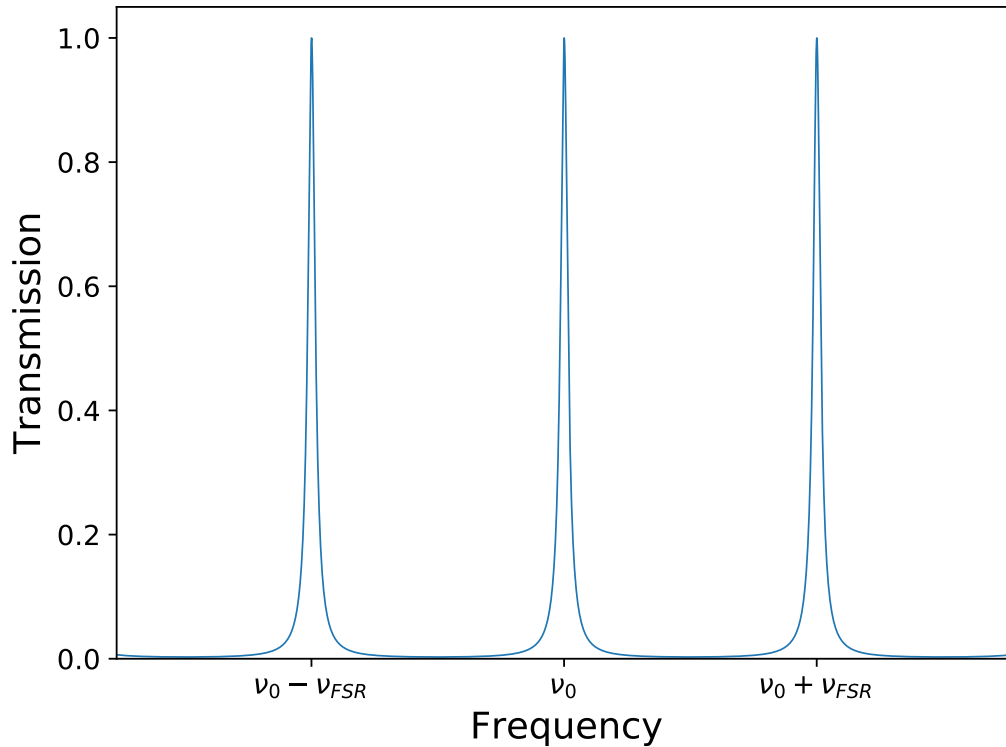


Figure 3.2: *Cavity Transmission vs Frequency*

The calculated transmission for an optical cavity aligned to the fundamental mode. The length is set at 75 cm for a FSR of 200 MHz. The mirror reflectivities are set at a rather low value of 0.9 so that the resonances are more visible.

In practice, it is usually the reflected signal that is monitored rather than the transmitted. This signal is demodulated using the same source that drives the EOM. Appropriate filtering is applied to the demodulated signal, yielding the "error signal", which is similar to the derivative of the reflected laser signal with respect to frequency, and is asymmetric about the resonance. This error signal is sent to a proportional-integral-derivative (PID) controller, which, with appropriate transfer-function tuning, can then be sent back to any number of mechanisms that can control the laser frequency. A plot of an idealized error signal as the laser frequency is scanned is shown in below.

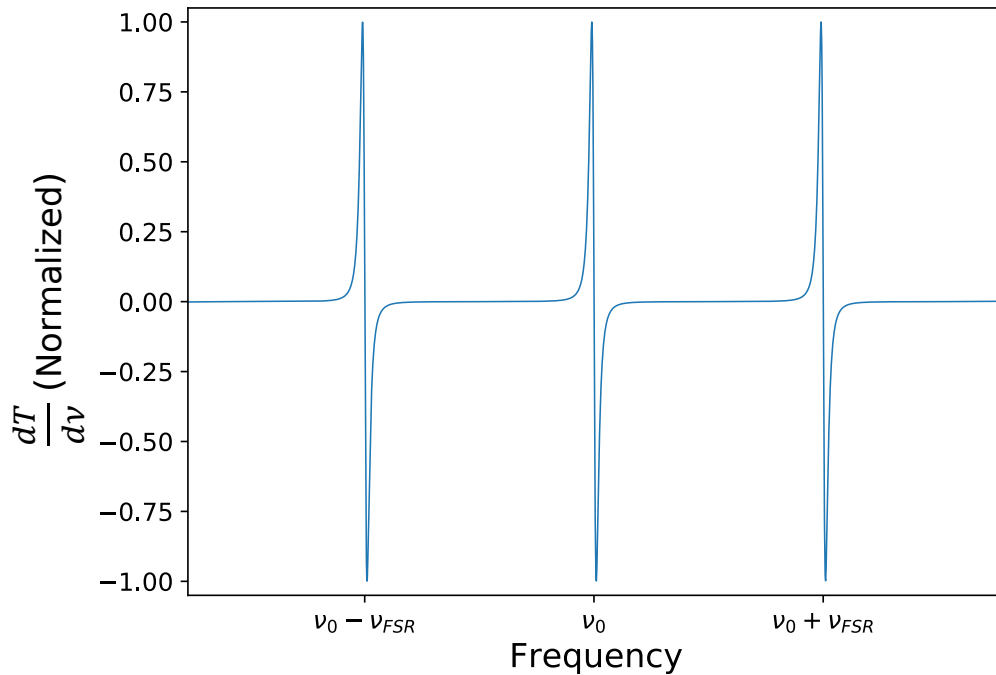


Figure 3.3: *The Derivative of the Cavity Transmission with Respect to Frequency*

The derivative of the transmission for the same cavity parameters as of Figure 3.2. This is closely related to the error signal and has the useful properties of being asymmetric around a resonance and being very linear when the frequency drift is much less than the cavity linewidth.

A diagram of our specific implementation of the PDH technique is shown in Figure 3.4. An RF generator produces a 50 MHz sine wave which is split into two signals. One branch is amplified and drives the EOM that was shown in Figure 3.1. The other branch is sent to a mixer where it is combined with the reflected photodiode signal that has been sent through a DC block and an amplifier. Note that light that has been modulated by the EOM before it reaches this photodiode, and so it contains the necessary information about the interference of the carrier and sidebands to generate the error signal and is centered at 50 MHz. A phase shifter consisting of a variable length of cable is in place before the RF signal reaches the mixer, to ensure it has the proper phase relationship with the reflected signal. The output of the mixer goes through a low-pass filter, which leaves only the difference

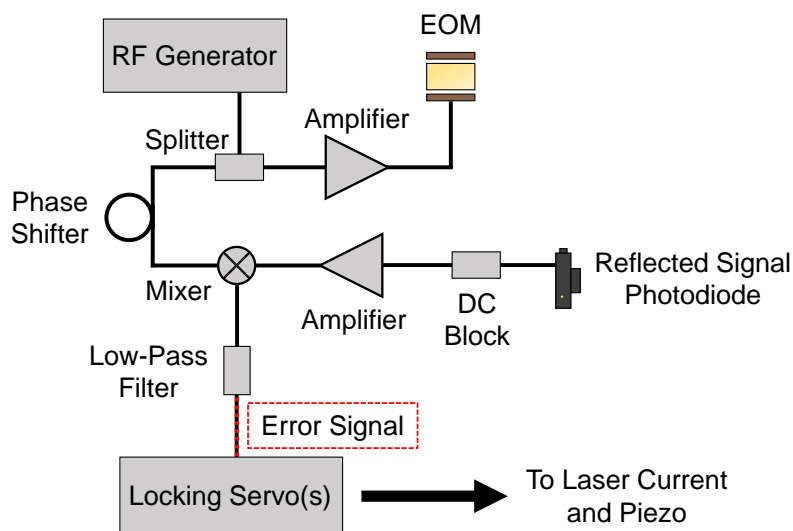


Figure 3.4: *High-Level Locking Setup Schematic*

An RF generator produces a 50 MHz sine wave which is split into two signals. One branch is amplified and drives the EOM that was shown in Figure 3.1. The other branch is sent to a mixer where it is combined with the reflected photodiode signal that has been sent through a DC block and an amplifier. A phase shifter consisting of a variable length of cable is in place before the RF signal reaches the mixer, to ensure it has the proper phase relationship with the reflected signal. The output of the mixer goes through a low-pass filter, which leaves only the difference of the mixer inputs and so produces a demodulated version of the reflected signal. This is the error signal and it is sent to a locking servo which provides feedback to the laser current and piezo.

of the mixer inputs, and so produces a demodulated version of the reflected signal. This is the error signal and it is sent to a locking servo which provides feedback to the laser current and piezo.

Our original setup split the error signal into low-frequency and high-frequency signals that were sent to two separate homemade PID controllers. The low-frequency signal (approximately DC to 100 Hz) was sent to the piezo on the diffraction grating of the ECDL. The mechanical response of the piezo is limited to these low-frequencies, but the piezo can make relatively large laser frequency adjustments, so is appropriate for long-term drift corrections. The high-frequency signal (approximately 100 Hz to 1 MHz) was sent to the

laser current controller for smaller, faster corrections.

When locked, the transmitted light from the cavity had a typical linewidth of about 20 kHz, indicating significant narrowing had occurred. The overall transmitted power was relatively low though. Many cavity systems regularly achieve >90% transmission when locked, however ours was typically closer to 10-20%. The stability of the lock was also very poor, often losing resonance after only a few seconds. In 2014, we replaced our two homemade controllers with a single commercial one (a Vescent D2-125), that could output two different signals for "fast" and "slow" response. This somewhat improved the consistency of the lock, but did not have a large impact on the transmitted power or narrowing of the laser linewidth.

There are a number of factors that contribute to the locking performance issues, not all of which were entirely anticipated during the initial design of this experiment. As previously stated, the 75 cm length of cavity results in narrower resonances, as well as increases its susceptibility to vibrations. The high-finesse of the cavity, while necessary for achieving high laser intensities and molecular coherence, also decreases the linewidth and makes locking more challenging. Additionally, more than 10 W of laser power is incident on the cavity, and the high intensities that build up in the cavity can rapidly heat up the cavity mirrors when on-resonance. This cannot only temporarily affect their material properties and cause fluctuations in circulating power, but also can affect the length of the cavity and cause large shifts in the cavity resonances. The liquid nitrogen cooling can have similar effects, causing the cavity to expand and contract with changing temperature, as well as creating mechanical perturbations as it boils on the exterior surface of the cavity. Finally, the deuterium gas itself creates issues. Although the gas pressure is low (usually about 0.1 - 1.0 atm), because a photon in the cavity makes many thousands of round trips on average before exiting, it effectively travels through tens of kilometers of gas. This creates many opportunities for undesirable Rayleigh scattering, which can result in losses and power

fluctuations. Additionally, the generation of Stokes light depletes power from the pump beam, which again creates intensity fluctuations.

All of these factors together resulted in the relatively low transmission, as well poor stability in the lock. Unsurprisingly, when testing the cavity under vacuum rather than filled with gas, and without liquid nitrogen cooling, locking performance was significantly better. We attempted to address some of these issues with experimental redesigns, as discussed in Chapter 6.

4 Past Work

In 2009, when Deniz and his student Tyler Green first began to work on the molecular modulation project, the foundations for the technique had been well established by other research groups. Josh Weber's thesis [8] gives a good overview of this work, as do several of our group's publications [15, 16], so this chapter will provide just a short summary of this work and how it set the stage for the specific direction of our group's research. We then describe the major results from our lab from before I began work on the project.

4.1 Previous Modulation Research

Work on molecular modulation began in the early 1990s, with one of the first proposals coming from the Imasaka group for using Raman scattering to induce modulation. Their simulations showed that Raman scattering in molecular hydrogen could generate pulses at the femtosecond level [17]. Around the same time, Kaplan suggested a similar approach using cascaded stimulated Raman scattering [18]. In 1998, Kawano demonstrated stimulated Raman scattering in hydrogen by pumping with a pulsed Ti:sapph laser, observing both vibrational and rotational Raman scattering, and producing spectral components that would be capable of forming ultrashort pulses [19]. In the early 2000s, the Harris group (featuring Deniz!) began experiments using two driving lasers at the pump and Stokes frequencies in hydrogen, and generated a spectrum from 195 nm to 2.94 μm [20]. In his PhD work, Deniz generated a spectrum with over 200 components from a hydrogen-deuterium mixture [21]. The Harris group went on to generate single-cycle pulses using these techniques [22]. In 2008, the Kung group produced a Raman spectrum that would

be capable of synthesizing pulses less than 0.5 fs long when corrected for dispersion [23]. Part of what made this possible is that the driving laser frequencies were chosen as exact multiples of the modulation frequency, and so the spectrum was completely harmonic. In 2012, the Imasaka group was able to generate a very broad harmonic spectrum, with components spanning from UV to NIR [24].

All of these experiments used pulsed driving lasers. So, while they are impressive achievements, they all suffer from the inherent limitations discussed in subsection 3.3.1. Work in CW molecular modulation was pioneered largely by the Carlsten and Imasaka groups. In 1998, Carlsten first demonstrated CW Raman lasing using a hydrogen gas-filled high-finesse cavity [25], which is the basic framework of many of our experiments: A single pump beam is locked to a cavity that is resonant at both the pump and Stokes wavelengths. When Raman scattering occurs, Stokes photons resonate in the cavity and can build up to significant intensities and begin lasing. The group continued to make substantial progress in the CW regime, generating both Stokes and anti-Stokes sidebands from vibrational and rotational molecular transitions [26, 27]. Carlsten has also developed formalism for the cavity Raman system that we make use of in Chapter 6 [28]. Imasaka has also demonstrated CW Raman lasing in cavities and showed that this and similarly generated spectra are coherent, and so can be useful in pulse and waveform generation [29, 30]. Deniz built on this work in 2007 with a proposal to use a system like this to modulate an independent mixing beam, not just the driving pump laser [31]. Since then, a main focus of the Yavuz Group's work on this project has been to establish high enough molecular coherence in the CW regime so that an independent beam may interact with it (i.e. a high-frequency broadband optical modulator).

4.2 Past Work in the Yavuz Group

Most of the work in the Yavuz Group prior to when I joined the lab in 2014 was performed using a similar experimental layout as shown in Figure 3.1, or slight variations thereof.

The first major result came in 2009 [32]. The main differences in the experimental design from Figure 3.1 were that a 1555 nm ECDL with a fiber amplifier was used as the pump beam, rather than one at 1064 nm. An earlier version of the optical cavity was also used, which was only 27 cm long and was not cooled by liquid nitrogen. For this experiment, the pump laser excites the $|v = 0, J = 1\rangle$ to $|v = 0, J = 3\rangle$ transition in deuterium, which has a frequency of 8.9 THz. Stokes light is produced at 1630 nm, where the cavity is also highly reflective. The pump beam is also circularly polarized to induce the preferential generation of first-order Stokes light due to angular momentum conservation rules, rather than higher-orders. A maximum of 321 mW of Stokes light was output from the cavity (the highest amount ever at the time using this sort of scheme), with conversion efficiencies from the pump to Stokes light as high as 65% at low pump powers.

For the next experiment in 2010, rather than focusing on generating one strong Stokes beam, the goal was to make a broad spectrum [15]. Here, the pump laser used is the amplified 1064 nm ECDL as described in Section 3.2, and the target transition is from the $|v = 0, J = 0\rangle$ to the $|v = 1, J = 0\rangle$ state in deuterium, at a frequency of 89.6 THz. Linearly polarized light is used, and Stokes light at 1555 nm is generated. As the cavity is highly reflective at 1555 nm as well as 1064 nm, high intensities of both pump and Stokes light build up. The molecules begin oscillating in unison as coherence is established, and the pump beam interacts with this coherence to generate anti-Stokes light at 807 nm. The cavity is not very reflective at 807 nm, so this light does not build up to high-power, and leaves the cavity after just a single pass. 10s of mW of pump and Stokes beams are emitted from the cavity, along with 10s of μ W of the anti-Stokes beam. While the power of the anti-Stokes

light was weak, it is significant that the deuterium gas pressures used (0.1 to 0.5 atm) were an order of magnitude lower than those in other CW Raman experiments. This means that there is much less dispersion in the output from the cavity, which would be important for the types of pulse generation and waveform synthesis described in Section 1.1. To show potential for these applications, the mutual phase coherence of the three spectral components was also demonstrated using an interferometer and a frequency doubling BBO crystal.

Following these results, a broadband spectrum with many components was generated in 2011 [33]. This experiment used the newer version of the optical cavity that was described in Section 3.3, and so was also cooled by liquid nitrogen. Sidebands from two transitions, $|v = 0, J = 0\rangle$ to $|v = 0, J = 2\rangle$ and $|v = 0, J = 1\rangle$ to $|v = 0, J = 3\rangle$, were observed with frequencies of 5.4 and 8.9 THz respectively. Because the cavity mirrors are still very reflective within tens of nanometers of the pump wavelength, multiple orders of these sidebands were generated with up to 12 simultaneous spectral components present.

Having established the ability to generate a broad, coherent spectrum, the next goal was to show that the system could modulate a separate laser (called a mixing beam), not just the pump beam that drives the molecular vibrations. This was achieved later that year, and the results were published in 2012 [16]. The experimental setup was again similar to the previous ones, but hydrogen gas was used, and the $|v = 0, J = 1\rangle$ to $|v = 0, J = 3\rangle$ transition at 17.6 THz was excited using circularly polarized light to suppress higher-order Stokes-modes. With the system pumped at 1064 nm, Stokes light at 1134 nm was generated. Now though, with the molecules in a coherent state, a separate laser at 785 nm was sent through the system. This light is generated from an ECDL and is amplified by a tapered amplifier to an incident power of about 100 mW. The light does not resonate with the cavity, and so travels through in a single pass. As it interacts with the hydrogen in the cavity, it is modulated by 17.6 THz, and sidebands at 750 nm and 823 nm are generated. These

sidebands are also not resonant with the cavity, and so they are transmitted through and then detected. A maximum modulation efficiency of $2.3 \cdot 10^{-6}$ is measured. While the efficiency was low, this was the first experiment in which modulation of an independent mixing beam was observed using CW light for both the mixing and pump beams.

The last major result before I began work on the experiment came in 2013 [34]. Most of the successes so far had come from using the rotational transitions in hydrogen or deuterium, but not the much higher frequency vibrational transitions. This was partially due to the ability to suppress higher-order Stokes-modes by using circularly polarized light with rotational transitions (thus increase the power in a single mode and improving the ease of detection). But after having generated a broad spectrum and showing the ability to modulate independent beams using the rotational transitions, the goal was now to use both vibrational and rotational transitions to generate as broad and as dense of a spectrum as possible. To this end, deuterium gas was used along with a linearly polarized pump beam. A "ro-vibrational" spectrum was produced containing components with frequencies of the form

$$\nu_{l,m,n} = \nu_{\text{pump}} \pm l\nu_{\text{vib}} \pm m\nu_{\text{rot1}} \pm n\nu_{\text{rot2}}. \quad (4.1)$$

In this equation l, m, n are integers, $\nu_{\text{pump}} = 282$ THz (1064 nm), ν_{vib} is the $|v = 0, J = 0\rangle$ to $|v = 1, J = 0\rangle$ transition frequency at 89.6 THz, ν_{rot1} is the $|v = 0, J = 0\rangle$ to $|v = 0, J = 2\rangle$ transition frequency at 5.4 THz, and ν_{rot2} is the $|v = 0, J = 2\rangle$ to $|v = 0, J = 2\rangle$ to $|v = 0, J = 4\rangle$ transition frequency at 12.4 THz. A trace of a typical spectrum produced in this experiment is shown below.

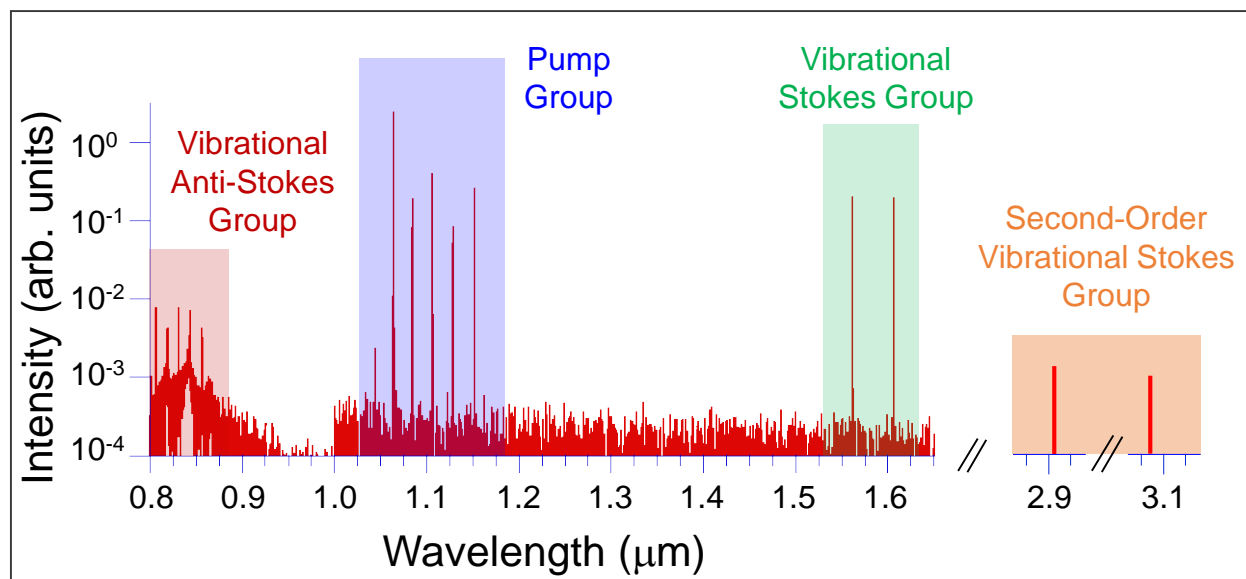


Figure 4.1: A Broadband Ro-Vibrational Spectrum

The ro-vibrational spectrum generated from a single pump beam at 1064 nm. Over 15 spectral components are present, spanning nearly two octaves of bandwidth. All components are captured on a spectrum analyzer, except for the two longest wavelengths, which were out of range and recorded separately using a photodiode. The pump group in blue contains the pump beam as well as several orders of modulated light from the two rotational transitions. Each component of this group then acts as an independent source, and is modulated by the vibrational transition, producing the vibrational Stokes and anti-Stokes groups. Note that only two frequency components are present in the vibrational Stokes group, as many of the components fall outside the range of high reflectivity of the cavity, and so are very low-power and not detected. These two high-power components are modulated by the vibrational transition again, and form the second-order vibrational Stokes group.

5 A 90 THz Modulator

To briefly summarize our previous results, we had gradually worked our way up from generating a single sideband on the pump beam, to many spectral components (but at the lower rotational frequencies of the gas), to finally a spectrum with many components at both vibrational and rotational frequencies. Meanwhile, the progress on modulating separate mixing beams had lagged behind, and had only been demonstrated using a 17.6 THz transition in hydrogen. So when I joined the lab in 2014, it was clear that the next step in the project was to modulate a mixing beam using a higher frequency transition. We then began work towards modulating a beam using the approximately 90 THz $|\nu = 0, J = 0\rangle$ to $|\nu = 1, J = 0\rangle$ transition in deuterium.

5.1 The Mixing Beam

The mixing beam was chosen as a 785 nm ECDL, like what was used in the previous mixing experiment. This diode was chosen largely for convenience—other experiments in our group had used them, and so we had some experience with working with them, as well as the necessary parts. Additionally, tapered amplifiers are available at this wavelength, and the ability to amplify the mixing beam makes the experiment easier. While it is true that the modulation efficiency is independent of the mixing beam power, because the efficiency is low, we want as much power in the mixing beam as possible to maximize the power in the sideband for detection.

The mixing beam is produced from an ECDL of the same design as that of the pump beam shown in Section 3.2. The output is sent through an isolator, and then a half-wave

plate and a polarizing beam cube, to ensure polarization purity and allow for easy power adjustments, before it is coupled to a tapered amplifier. The power is amplified to about 1 W. The output beam profile is highly elliptical, as is typical for tapered amplifiers, and so it is shaped with a cylindrical lens. The beam is then sent through a second isolator and into a single-mode fiber to give it a near perfect Gaussian profile at the output, before it is sent to the optical cavity. The beam is mode-matched to the cavity, not for resonance (as it travels through in a single pass and thus is not needed), but to ensure good spatial overlap with the pump beam so that it interacts with the region of coherence in the molecules. The beam is overlapped with the pump beam using a dichroic mirror that transmits light under 1000 nm and reflects light above 1000 nm. The mixing beam propagates in the opposite direction of the pump beam (simply due to space constraints on the optical table), and about 200 mW of the mixing beam is incident on the cavity. After traveling through the cavity and being modulated, the mixing beam and modulated light are separated from the pump beam, again using a dichroic mirror. A diagram of the major features is shown below.




Figure 5.1: *The Mixing Beam Experimental Design*

The pump beam enters from the upper-left of the diagram after having passed through the same optics shown in Figure 3.1. It then resonates in the cavity and establishes molecular coherence in the deuterium. A 785 nm beam is produced by an ECDL of the same design as the pump ECDL. The beam passes through an isolator and is amplified to 1 watt by a tapered amplifier. The beam then passes through another isolator and is coupled to a single-mode fiber to ensure good spatial-mode quality. The beam is roughly mode-matched to the cavity so that it will have good overlap with the pump beam. Dichroic mirrors, which transmit the mixing beam while reflecting the pump beam, allow the mixing beam to be combined with the pump beam and then separated after passing through the cavity, along with any sidebands put on it.

5.2 Sideband Detection

Once the experiment is set up for the pump beam to resonate in the deuterium-filled optical cavity, as it has been in many other experiments, the addition of the mixing beam is not particularly challenging. After being amplified and shaped, simply sending the mixing beam through the cavity with the pump beam locked almost certainly results in the generation of sidebands on the mixing beam. The difficulty lies in detecting these sidebands in order to demonstrate the functionality of our modulator (since the efficiency would most likely be extremely low), and in maximizing the efficiency within the limitations of the experiment. Much of the laboratory work was devoted to sideband detection and efficiency optimization. A description of the procedure we developed for this process is given below, for the sake of future students who might encounter similar challenges with this experiment.

It is worth noting that the frequency up-shifted sideband from the 785 nm mixing beam is within a couple nm (depending on the exact tuning of the ECDL) of a helium-neon (HeNe) laser wavelength of 633 nm. This is extremely convenient, as it allows us to use a HeNe to trace the wavelength-dependent path that the modulated light will take through our system. First, we adjust the mixing beam and a HeNe "alignment" beam so that they will propagate co-linearly through our system. Although we did not have the optics available to make the mixing and alignment beams propagate simultaneously, a flip mirror allows us to quickly switch between which of the two beams is directed towards the cavity. To ensure that the paths they take overlap (so that the alignment beam will accurately trace the path of the mixing beam and any sidebands), a second flip mirror is used to direct each beam in turn away from the cavity and into a single-mode optical fiber. Because the fiber is very sensitive to alignment, by separately optimizing the transmission of both the mixing and the alignment beam using independent sets of mirrors, we can be confident that they

will take the same path when either is directed through the cavity. Another set of mirrors right before the cavity will adjust the path that both beams take, and allows the beams to be overlapped with the counter-propagating pump beam that is transmitted through the cavity. This overlap can simply be done simply by eye using a fluorescent IR card, as a more careful alignment procedure is not necessarily helpful if the mixing and pump beams are not the exact same shape. A diagram of this setup is shown in Figure 5.2.

After passing through the cavity, the mixing beam and sidebands (or alignment laser if in use) are separated from the pump beam and other spectral components generated in the cavity by using a dichroic mirror as previously described. The sidebands are produced co-linearly with the mixing beam and are many orders of magnitude weaker than it, so the light is sent to a diffraction grating where we may take advantage of the large modulation frequency to get good spatial separation of the beams. The frequency up-shifted light is allowed to propagate several meters before being focused onto a photodiode, which outputs a signal to an amplifier followed by an oscilloscope. The alignment beam being at nearly the same wavelength of the modulated light is very valuable for this purpose, since it can be used to align the optics after the diffraction grating without having to make extremely accurate measurements and calculate where the detector should be placed.

A large number of optics are required before the light finally reaches the photodiode, and so losses are inevitable. Additionally, before optimizing the overlap of the mixing beam with the pump beam, the modulation efficiency may be quite low. These factors together mean that the modulated light at the photodiode may be at the picowatt level or even lower and impossible to detect, especially given the presence of scattered light from the 20 W pump beam that must be running. Isolating the beam path to the detector with PVC pipe and dark fabric helps cut down the background noise a significant amount. Eventually though, to improve the signal to noise ratio, we found we needed to send the mixing beam through a beam chopper before it reached the cavity (not shown in the diagram). The

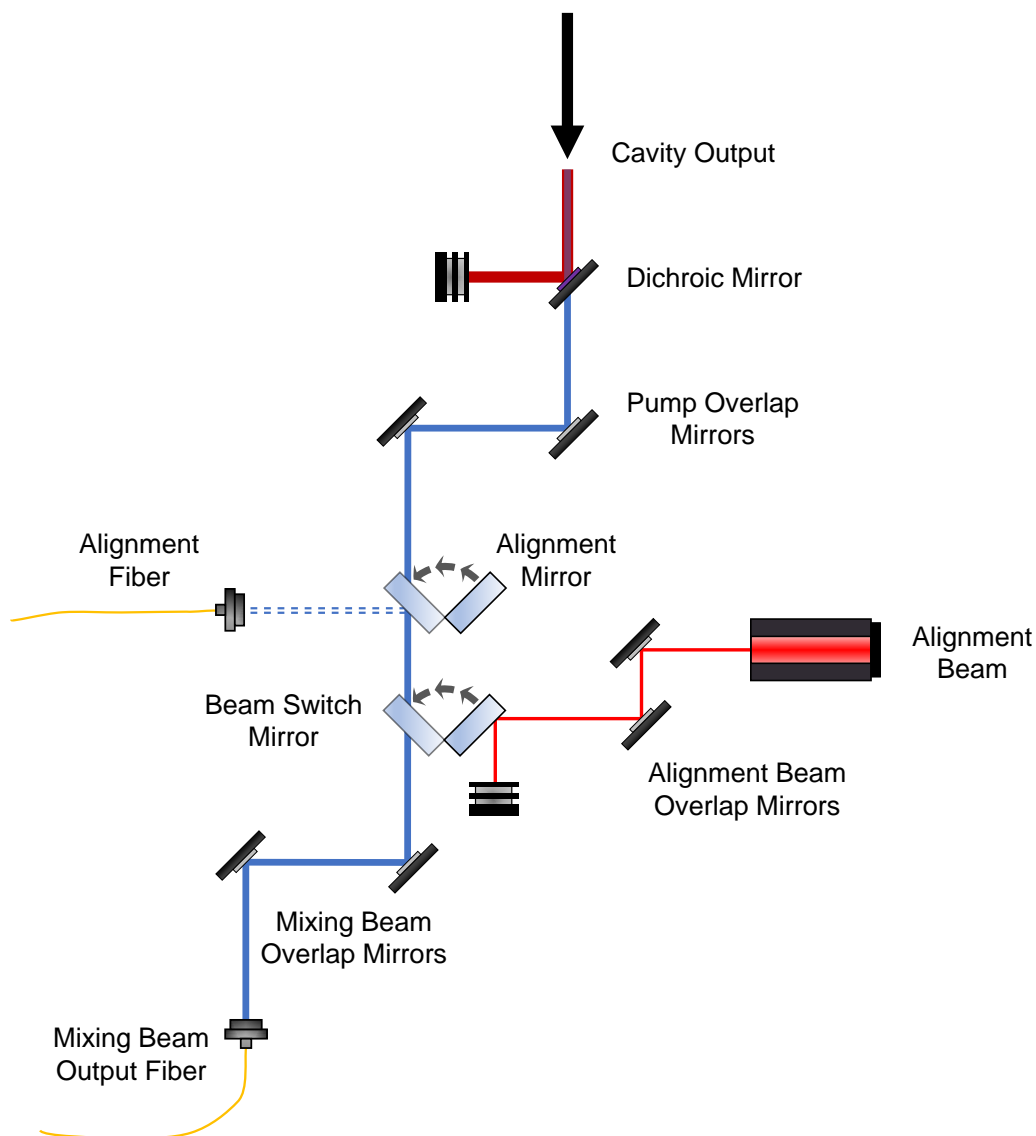


Figure 5.2: *The Mixing Beam Alignment Setup*

After coupling to the fiber as shown in Figure 5.1, the mixing beam is sent to a set of "overlap" mirrors. A HeNe "alignment" beam is sent to its own set of overlap mirrors. The "beam switch" flip mirror can be used to switch which of the two beams are directed onward to the cavity. A second "alignment" flip mirror can be used to direct each beam in turn to a single-mode fiber. By separately aligning each beam for good coupling to the fiber using the overlap mirrors, it can be ensured that the mixing and overlap beams will take the same path through the cavity. Both beams can then be roughly overlapped with the counter-propagating "pump overlap" mirrors by eye using an IR viewing card

beam chopper, essentially a small set of spinning fan blades, rapidly blocks and unblocks the beam, giving it a square-wave intensity profile in time with a 50% duty cycle and a frequency of around 1 kHz. Sending the output of the photodiode to a lock-in amplifier set to this same frequency can then greatly cut out background noise. A diagram of this detection scheme is shown in Figure 5.3.

5.2.1 Detection by Eye

Despite these efforts to improve the ease of detecting the mixing beam sideband, after any alignment changes, or several days of down-time for the experiment, the sideband was often undetectable. It would then take a great deal of time and luck to overlap the mixing beam well enough with the pump beam to generate a measurable signal. This was made especially difficult by the poor locking performance of the pump beam to the cavity, which necessitated a delicate balancing act of being careful enough to not disturb the lock, while moving quickly enough to make alignment adjustments before the lock drop on its own. Interestingly, the human eye is quite sensitive to small amounts of light and can surpass many detection schemes under the right conditions, and we were able to use this fact to develop an effective alignment procedure.

When we are unable to measure a sideband signal at the photodiode, we let the cavity output reflect off of the diffraction grating, and then direct it to the edge of the optical table using a flip mirror where we try to view it by eye (shown in Figure 5.3). While any modulated light should be well separated from other beams at this point, and beam blocks are also in place to allow only this modulated light through, additional precautions are taken, given that there is a 20 W 1064 nm beam and a 1 W 785 nm beam in use, each passing through many optics with potential for scatter. First, several dielectric filters are used to make sure the frequency up-shifted sideband is properly isolated and to block any stray




Figure 5.3: *The Mixing Beam Detection Setup*

The mixing beam and any sidebands exit from the top of the cavity, along with some transmitted pump and Stokes light and a large amount of reflected pump light. The dichroic mirror reflects most of the light over 1000 nm, which eventually reaches a beam block (not shown). The mixing beam and sidebands are transmitted through the mirror and reach a diffraction grating, which is necessary to spatially separate the much weaker sideband. A flip mirror allows the sideband to be redirected towards the edge of the table for detection by eye when needed. Otherwise, the sideband is allowed to propagate over several meters for increased spatial isolation, and is then sent through some final bandpass filters before being focused onto a photodiode. This signal is either sent to an oscilloscope for direct viewing or to a lock-in amplifier (in which case the mixing beam is sent through a beam chopper right after amplification).

scattered light. Next, the area around the viewer is always checked for scattered light using an IR card and an IR scope before viewing. Finally, the viewer wears goggles that block IR light but transmits most visible light. For additional safety tips, see the riveting lab safety video made in collaboration with former labmate Nick Brewer [35].

With the appropriate safety precautions taken, the room lights are turned off and the pump beam is locked to the cavity with the mixing beam passing through. The viewer then crouches down level with the table and looks into the flip mirror at an angle so that the reflection of the diffraction grating is visible. (Note that the cavity mirror from which the light exits will not be visible, as the system is aligned to view the first-order diffracted light.) By slowly "scanning" their head along the mirror in their field of vision, the viewer will eventually see a flash of red light. Since the light is a collimated beam, it is very directional and be quite difficult to find at first. With practice though, this becomes fairly easy. The viewer takes a mental note of the perceived brightness of the light, and makes a small adjustment to the alignment mirrors for the mixing beam. They then look again to see if the light has become brighter or dimmer and adjust the mirror again in the same or opposite direction as needed. Typically, just a couple of iterations of this procedure over the horizontal and vertical adjustments of a single mirror will be enough to increase the sideband power to where it can be once again detected by the photodiode. If possible, having a friend, or even just a begrudging colleague, make the mirror adjustments while the viewer maintains focus on the light can speed up this process. With the signal now detectable by the photodiode, it is easy to further optimize it using the alignment mirrors.

This method is actually so sensitive that an experienced viewer can even detect the 532 nm anti-Stokes sideband generated from the HeNe alignment beam, which is well below the power level that could be quantified with the detector. The HeNe beam is approximately 1 mW, and so accounting for various losses and the inherent misalignment with the pump beam, there is likely only 10s of femtowatts of 532 nm light reaching the eye. The adjustable

gain of the human eye can certainly rival that of any Thorlabs PDA10A photodiode!

5.3 Results

We achieved the best performance with a deuterium gas pressure of 0.3 atm, and reached $0.5 \cdot 10^{-6}$ conversion efficiency into the frequency up-shifted sideband. While the efficiency should be the same for the frequency down-shifted sideband, we did not directly measure this. This conversion efficiency number accounts for losses from various optics after the cavity, and so is defined as $\frac{P_S}{P_M}$ inside the optical cavity, where P_S and P_M are the optical powers of the frequency up-shifted sideband and carrier mixing beam respectively. With approximately 100 mW of mixing beam power in the cavity, we have about 50 nW of sideband generated there, with single nW level powers when traversing most of the detection setup. Although this conversion efficiency was low by most standards, this was the highest modulation frequency of an independent mixing beam in the CW regime that had ever been demonstrated, and so was an important result. We published these results in *Applied Sciences* in 2014 [36].

We also explored the effects of the pump beam power on the modulation efficiency. We found that as we increase in the input pump power, the transmitted pump power from the cavity tends to stay constant at about 1 mW, whereas the transmitted Stokes power increases linearly to a maximum of about 7 mW. This is similar to the behavior we observed in our previous experiment where the system becomes saturated with the pump beam and additional pump power is converted into Stokes light [16]. Using the known transmission of the cavity mirrors and the shape of the cavity-mode, we can calculate the intensity profiles of the circulating pump and Stokes beams. The maximum intensities are found to be 10 kW/cm^2 and 62 kW/cm^2 for the pump and Stokes beams respectively. We then calculate the predicted modulation efficiency given these intensity profiles, the Raman-gain

coefficient of deuterium, and the gas pressure in the cavity. The solid line in Figure 5.4 shows the predicted conversion efficiency using the estimated intra-cavity pump and Stokes intensities. We observe reasonable agreement between the experimental data points and the calculations. The somewhat lower measured efficiencies than the predicted ones are likely largely due to an imperfect spatial overlap between the mixing beam and pump beam in the cavity. The spatial overlap is difficult to estimate in our experiment and our calculations do not take this effect into account.

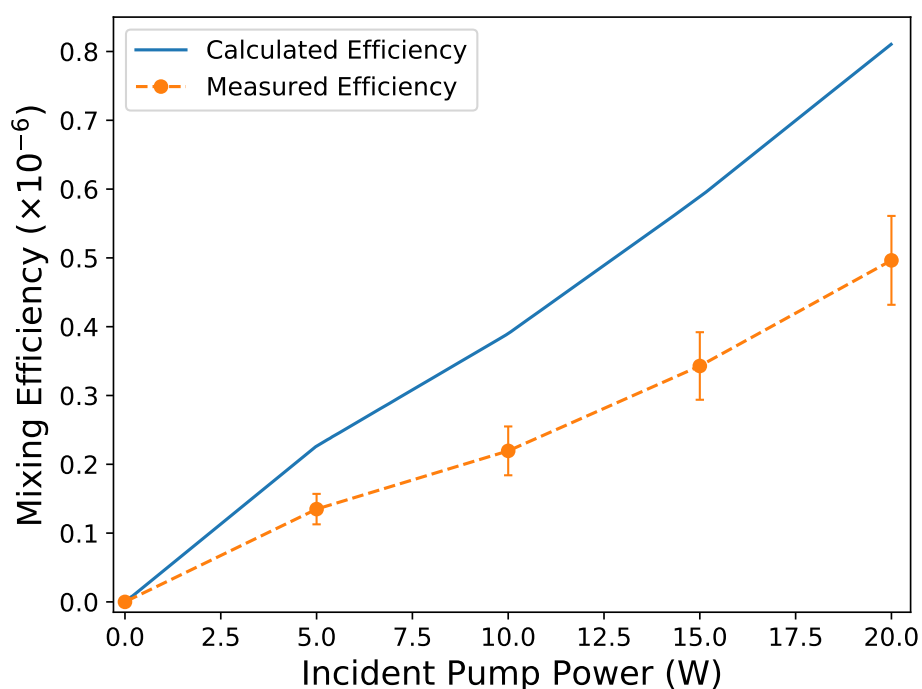


Figure 5.4: *Conversion Efficiency vs Incident Pump Power*

The conversion efficiency from the mixing beam to the approximately 636 nm anti-Stokes sideband is shown as the incident pump power is varied at a constant deuterium pressure of 0.3 atm. The solid line is a theoretical calculation, based on the measured transmitted powers of the pump and the Stokes beams through the cavity.

6 Cavity Redesign

While we were able to use our experimental setup to demonstrate modulation of an independent mixing beam at 90 THz, we felt that we were reaching its limits because detection of the mixing beam sideband had proven quite difficult. Additionally, the low modulation efficiency would likely make the more advanced experiments we wished to pursue prohibitively difficult, such as modulating an already broad laser source. To address this, we began to consider what changes we might make in our setup to increase both the efficiency and the ease of operation.

For guidance in this process, we first recall from Section 2.2 that the modulation efficiency within the cavity is approximately proportional to

$$\epsilon \propto p^2 \times I_p \times I_s \times l^2, \quad (2.36 \text{ revisited})$$

where p is the pressure of the gas, I_p and I_s are the circulating intensities of the pump and Stokes beams, and l is the length of the cavity.

The models used for most of the calculations in this chapter are based on the theoretical analysis performed by Carlsten et. al [28]. Note that our analysis contains an adjustable parameter that is chosen to scale the system's performance to be in line with what we observe experimentally. The model assumes mostly "ideal" conditions and so there are many effects that are not accounted for. We therefore do not take the results to be completely predictive of behavior we might see. Rather, we are simply interested in the trends they show, and what insights they might provide, as far as making improvements to the experiment.

The primary results we make use of from Carlsten's work are the equations for the

steady-state electric field of the pump laser inside the cavity ($E_{p_{ss}}$), the steady-state electric field of the generated Stokes beam ($E_{s_{ss}}$), and the threshold electric field of the pump beam incident on the cavity required to generate a Stokes beam ($E_{p_{threshold}}$). These are reproduced below:

$$E_{p_{ss}} = \left(\frac{L_S}{G} \right)^{1/2} \quad (6.1)$$

$$E_{s_{ss}} = \left[\frac{\omega_S(K - L_p E_{p_{ss}})}{\omega_p G E_{p_{ss}}} \right]^{1/2} \quad (6.2)$$

$$E_{p_{threshold}} = \frac{l \times L_p}{c} \left(\frac{L_S}{T_p G} \right)^{1/2}. \quad (6.3)$$

In these equations, L_p and L_S are the losses from the cavity mirrors at the pump and Stokes wavelengths; T_p is the transmission of the cavity mirrors at the pump wavelength; G is the Raman-gain coefficient, which depends upon the linear gain coefficient of the gas and the spatial overlap between the pump and Stokes cavity-modes; and K is the optical pumping constant, which depends upon the transmission of the cavity mirrors at the pump wavelength and the incident pump beam power.

6.1 Gas Pressure

Based on Equation 2.36, the easiest avenue for improvement might seem to be to simply increase the gas pressure and enjoy quadratic gains in efficiency. However, the equation is somewhat misleading in this regard, as the terms affecting the efficiency are not entirely independent of one another. A plot is shown below of the theoretical cavity output powers for the pump and Stokes beams when 1 watt of pump beam is coupled to the cavity and the deuterium gas pressure is varied. The cavity parameters are set as described in Section 3.3, with a finesse of about 23,000 and a length of 75 cm.

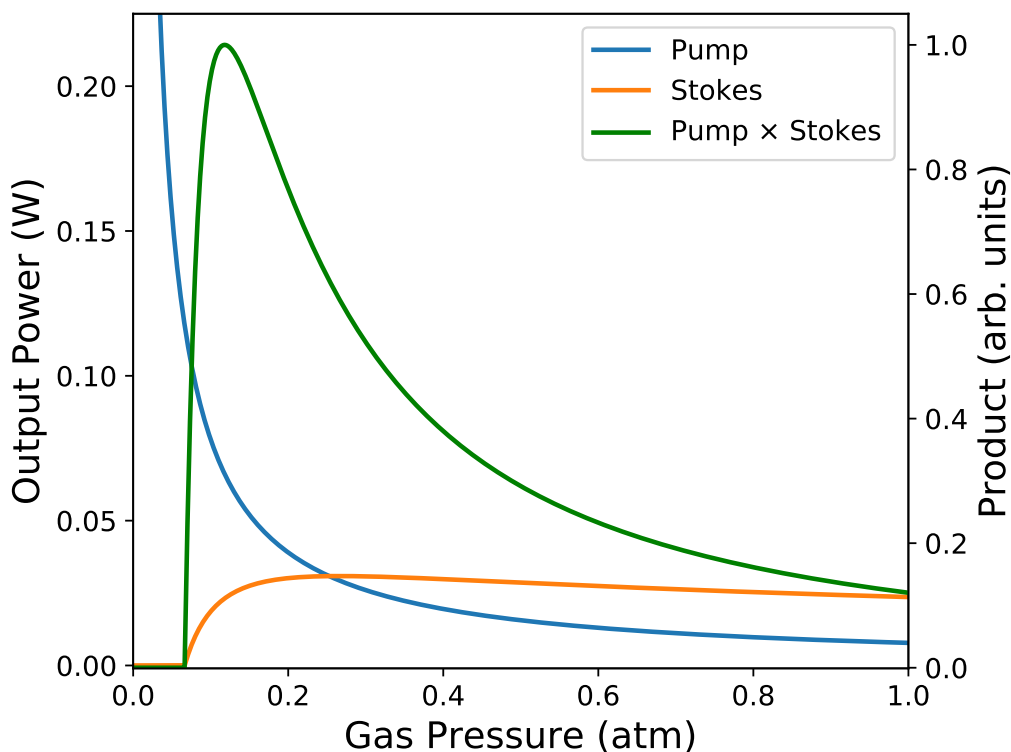


Figure 6.1: *Cavity Output Pump and Stokes Power vs Gas Pressure*

1 watt of pump power is coupled to the cavity. Keeping other performance values constant, as the gas pressure is increased the Raman lasing threshold is crossed and Stokes power builds up. At slightly higher pressure, the maximum product of the power of the two beams is reached. The pump power rapidly drops with increasing pressure as it is converted into the Stokes beam. Slightly after the maximum product is reached, additional gas pressure causes both the pump and Stokes power to drop.

Section 6.1 shows that after the pressure threshold is reached for Raman lasing with 1 W of coupled power, the product of the output power of the two beams is maximized with only a slight additional increase in pressure. Higher pressure will still increase the Stokes power somewhat after this point, but as this rapidly depletes the pump power, their overall product drops. At higher pressures still, the output of both the pump and the Stokes decreases.

We next plot the pump and Stokes output power for the same system, but now showing the estimated modulation efficiency based on Equation 2.36. The estimated modulation efficiency grows continuously with pressure after threshold is reached, since the growth from pressure outweighs the decrease in product of the beam powers. However, the efficiency only rises as approximately the square root of the gas pressure, rather than quadratically as the equation might imply.




Figure 6.2: *Conversion Efficiency vs Gas Pressure*

The same plot from Section 6.1, with the predicted modulation efficiency of a mixing beam included. After threshold is reached, the efficiency rises only as approximately the square root of the gas pressure, rather than quadratically as Equation 2.36 might imply.

Still though, this analysis would suggest that we should use as high of a gas pressure as we can, and perhaps it would be worth the effort to build the cavity to withstand extremely high gas pressures. However, not all effects related to gas pressure are accounted for in this

analysis, and given the relatively slow growth in efficiency that is predicted, this would probably not be a fruitful path. The most significant effect is likely that higher gas pressure can greatly affect laser locking performance, making it both more difficult to lock, and reducing the transmission when we do manage to so. We typically notice a big drop in coupling performance of the pump beam between calibrating the system under vacuum and then adding even just 0.1 atm of gas.

Additionally, higher gas pressure will increase the linewidth of the Raman transition due to pressure broadening. Pressure broadening will be dominated by collisional effects in this regime, and so will increase approximately linearly with gas pressure. This effectively decreases the Raman-gain of the system and so will reduce the established molecular coherence.

6.2 Cavity Geometry

We might next consider increasing the length of the cavity, which increases the interaction length of the laser with the gas. We are however constrained by the need for the beam shape to match a stable-mode of an optical cavity. This means that as we increase the length, we will generally have to increase the radius of curvature of the cavity mirrors, which will increase the overall beam size and reduce the intensities of the pump and Stokes beams. To examine the overall effect of cavity length, we can replace the constant pump and Stokes intensities in Equation 2.36 with an expression for the intensities per unit length and integrate over the length of the cavity:

$$\epsilon \propto p^2 \times l^2 \int_{-1/2}^{+1/2} \frac{I_p(z) \cdot I_s(z)}{\text{meter}} dz. \quad (6.4)$$

Note that this integration is performed on-axis with the cavity, and so due to the radial symmetry of the cavity-mode, the varying intensities in the x and y directions can simply be accounted for with a proportionality constant. $I(z)$, then, refers to the peak intensity value of a beam at position z , or, more precisely, $I(x = 0, y = 0, z)$ with respect to the optical axis of the cavity. Figure 6.3 shows the integrated product of the pump and Stokes intensities starting with 1 W of pump coupled to the cavity for various cavity lengths and mirror radius of curvatures (RoCs). We restrict the cavity geometries to a reasonably practical design range (i.e. not 10 m long cavities). For a given mirror RoC, increasing the cavity length will always increase the total integrated intensity, but shorter cavities can support smaller RoCs and smaller beam sizes. The rapid rise near the end of the curves occurs because the waist of the cavity-mode only barely increases with length at those points, and so the additional length mostly just allows for more of a beam to be enclosed. Cavities have real eigenmodes for lengths as large as $l = 2 \cdot \text{RoC}$, but tend to be less stable near this maximum value. So in practice, not all of this rapid rise is exploitable. Only values for $l < \text{RoC} - 1 \text{ cm}$ are plotted.

We next substitute Equation 6.4 into Equation 2.36 to use the integrated intensity rather than a constant value, in order to examine the effect of cavity geometry on modulation efficiency. Otherwise the same parameters as in the previous model are used, with 1 watt of pump beam coupled to the cavity. The gas pressure is chosen to be a constant 1 atm, although we note that this does not significantly affect the shape of the curves, as we are essentially examining just geometric effects here. Figure 6.4 shows the predicted modulation efficiency for the same cavity geometries that were shown in Figure 6.3, now using this integrated intensity.

The predicted mixing efficiency varies by about an order of magnitude over this range of cavity parameters. Effects such as the narrowing of cavity resonance linewidths and decreased stability as the cavity length is increased are not accounted for, as these are




Figure 6.3: *Intensity Product for Various Cavity Geometries*

Optical cavities using symmetric mirrors of various radius of curvature (RoC) are shown. Longer cavities give more total intensity, although the growth is not as rapid as one might suspect, since longer cavities result in larger mode-waists and lower peak intensities. For a given cavity length, smaller RoCs create a smaller mode-waist, and thus higher peak intensities and total intensity.

difficult to quantify. However, both of these effects would serve to reduce the modulation efficiency as length is increased, so the overall range in the modulation efficiency would likely be lower than what is shown here.

The fact that the predicted efficiency does not have a huge variance with the cavity geometry is significant. The allure of large efficiency gains with increased length would have been difficult to weigh against potential drawbacks, as we know from experience that longer cavities come with decreased locking performance and decreased laser power coupling. But the fact that the predicted dependency on cavity length is modest made

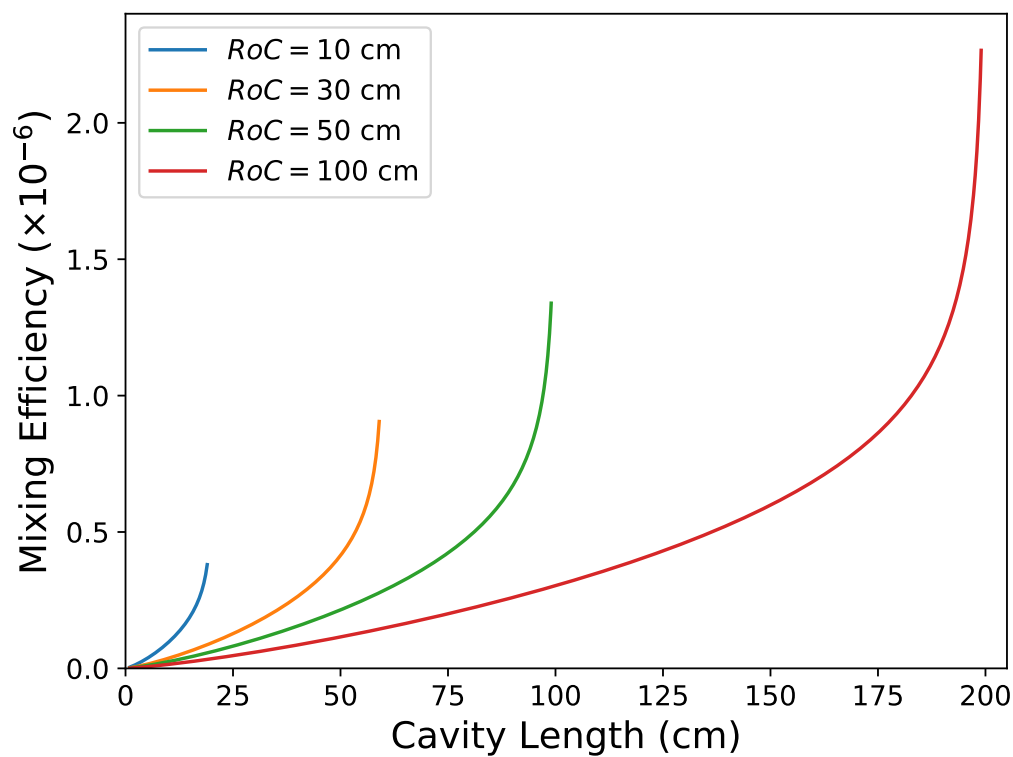


Figure 6.4: *Predicted Mixing Efficiency for Various Cavity Geometries*

The predicted modulation efficiency for cavity designs, accounting for purely geometric effects (i.e. the amount of light coupled and overall locking performance is assumed to be constant). The efficiency varies by about an order of magnitude, although unaccounted for effects would likely reduce this.

us confident that we would not lose much efficiency from reducing the length (possibly even gaining some with improved coupling), while at the same time improving the ease of operation with better locking performance. For this reason, we decided to actually switch to using a shorter cavity. Beyond this choice, we did not have much reason for a specific length, as the calculations Figure 6.4 is based on most likely do not contain the detail needed to inform this decision. So because we had a 27 cm cavity already built from previous experiments, we switched back to using it. We chose to use 30 cm radius of curvature mirrors with it, as these would provide us with a relatively high intensity product at this length, without pushing the bounds of stability on the cavity. A detail view of Figure 6.4 is

shown below, with the parameters of the original cavity and the chosen redesign marked.




Figure 6.5: *Original and Redesigned Cavity Mixing Efficiencies*

A detail view of Figure 6.4, with the original and redesigned cavity geometries marked. While the plot indicates a slight reduction in efficiency, we hoped that unaccounted for effects in the efficiency and the improved ease of operation at a shorter cavity length would outweigh this.

6.3 Cavity Finesse

A final factor to consider in redesigning the experiment is the finesse of the cavity, which will affect the circulating power and intensity of the pump beam, and therefore impact both the threshold of input power where Raman lasing will occur, and how much Stokes light will be generated. We begin by using the same parameters as in the previous plots, and this time set the cavity mirror reflectivity at the pump and the Stokes wavelength equal

to one another, and examine the modulation efficiency as this reflectivity is varied. A plot of the predicted efficiency vs cavity finesse is shown below.

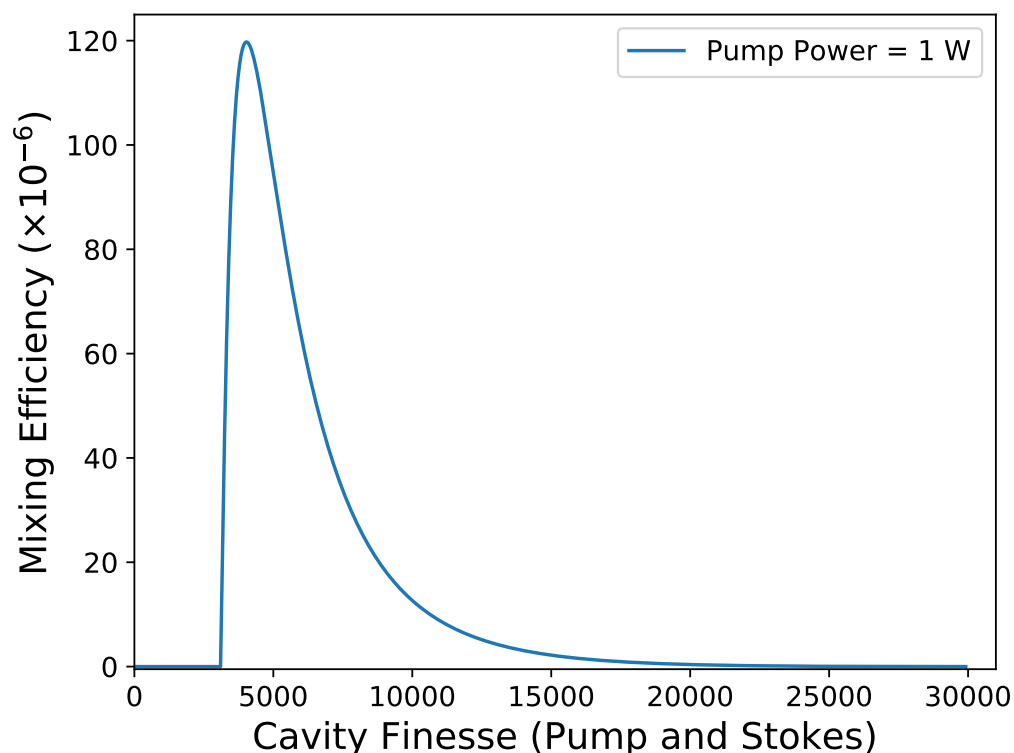


Figure 6.6: *Mixing Efficiency vs Cavity Finesse*

The mixing efficiency depends very strongly on the finesse of the cavity. In the absence of Raman lasing, a higher finesse would mean a higher circulating pump power, but because the Stokes and pump powers must be balanced for the highest efficiency, the efficiency reaches its maximum value at a finesse just slightly higher than the threshold needed for Raman lasing.

The results here were somewhat surprising, and suggest that our finesse of 23,000 might actually be too high. It seems that higher finesse is not necessarily better for high modulation efficiency, even though it would lead to higher laser powers inside the cavity. Rather, because the product of the pump and Stokes intensities must be maximized for maximum efficiency, there is an optimal finesse. The predicted efficiency spans about two orders of magnitude, approximately matching our current efficiency at our cavity finesse,

with much higher efficiencies at lower finesse. So it seems there could be great potential for improvements here!

The amount of coupled power should affect this optimal finesse and so we next examine the efficiency for several different input powers, as shown below in Figure 6.7. The curves are individually normalized so that their shapes and maxima may be more easily compared, since the un-normalized plot would span several orders of magnitude. For a given cavity finesse though, a higher coupled power always increases the predicted efficiency—increasing the pump power simply shifts the value of finesse where the maximum efficiency occurs.

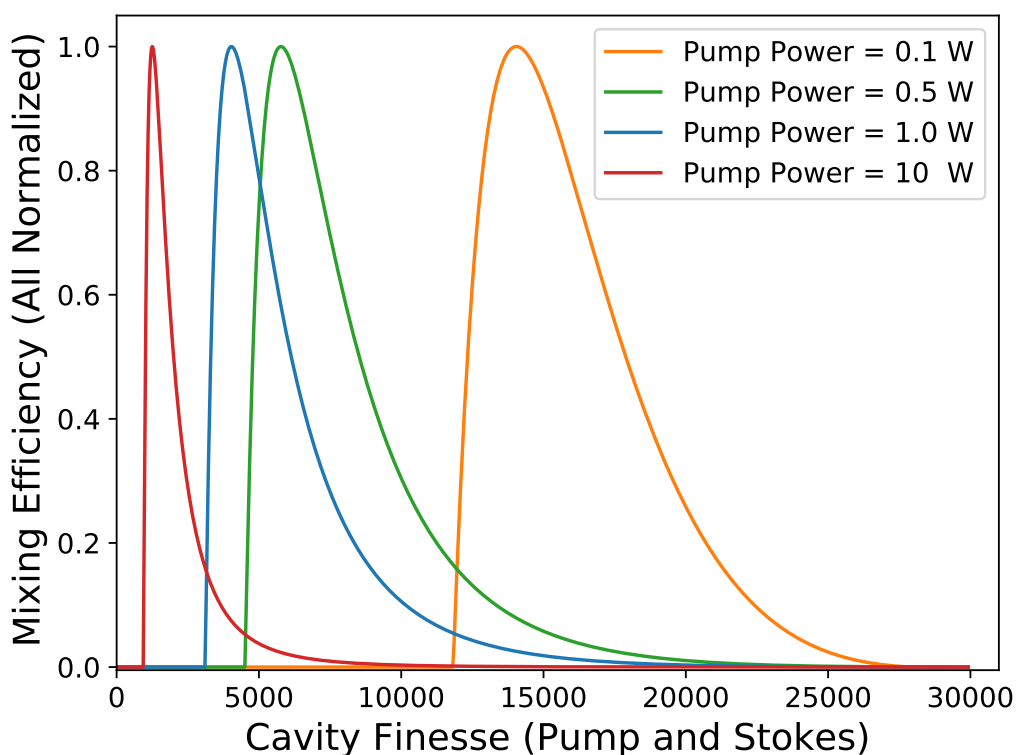


Figure 6.7: *Modulation Efficiency vs Cavity Finesse and Input Power*

Higher coupled pump power lowers the finesse where maximum modulation efficiency occurs, as well as narrows the efficiency curve. This is to be expected, as higher-power means a lower finesse is required to reach the Raman lasing threshold, and the maximum efficiency occurs shortly after threshold where the pump and the Stokes powers are properly balanced. The curves are individually normalized for ease of comparison, as the plot would otherwise span several orders of magnitude.

Note that higher coupled pump power lowers the finesse where the maximum modulation efficiency occurs, as well as narrows the efficiency curve. This is to be expected, as higher-power means a lower finesse is required to reach the Raman lasing threshold, and the maximum efficiency occurs shortly after threshold where the pump and the Stokes powers are properly balanced. To help better understand this behavior of the system, we plot the finesse at which the maximum modulation efficiency occurs against the input pump power in Figure 6.8.

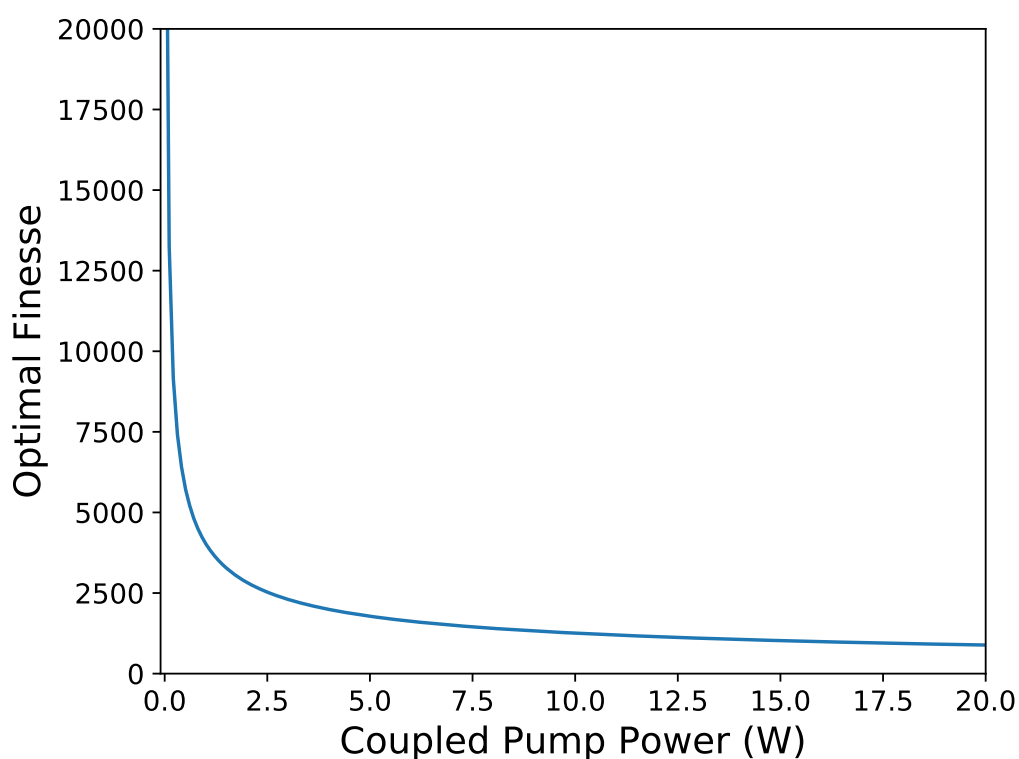


Figure 6.8: *Optimal Cavity Finesse vs Pump Power*

The plot shows the "optimal finesse"—the one that should give the highest modulation efficiency—for a given amount of coupled pump power. The optimal finesse is very sensitive to the pump power at low values, but levels off rapidly above about 3 W. We are unlikely to be able to couple this much power, and so will face this region of high sensitivity in trying to optimize the cavity configuration.

6.3.1 Choosing the Finesse

To choose the finesse for our new cavity, our best bet would seem to be to first determine how much pump power we can couple in, and then use the results of Figure 6.7. With our current setup, we usually have about 12 W of pump power incident on the cavity. When the cavity is filled with gas, and given its length and high-finesse, we typically can couple only about 1 W of pump power to it. This puts us close to the steepest part of the curve in Figure 6.8, where the optimal finesse depends very strongly on the coupled power. Unfortunately, this makes it very difficult to determine what cavity finesse might give us the best performance.

As previously mentioned, a variety of unaccounted for factors likely make the plots in this chapter informative of the trends and types of behavior to expect with our system, but not necessarily accurate in their precise values or rates of change. Additionally, the amount of power we can couple to our system is not a constant, and will likely depend strongly on the cavity finesse in ways that we cannot predict. This means that we cannot simply measure the amount of power we can couple to our current system and use this to determine the optimal finesse. Another point worth considering is that it would be risky to choose a new finesse that would give the maximum efficiency even if we could, because this point occurs only shortly after the threshold finesse. With only slight measurement error, manufacturing flaw, or drop in performance, we would be likely to end up with a system that is unable to pass threshold at all.

Finally, we note that while we have set the cavity mirror reflectivity for the pump and Stokes wavelengths, R_p and R_s , equal to each other in these calculations, this does not necessarily give the best performance. The optimal configuration lives in a rich parameter space depending on the gas pressure, coupled pump power, and R_p and R_s . Figure 6.9 below shows a plot comparing the curves for the modulation efficiency vs cavity finesse

with $R_p = R_s$ (the same curve in Figure 6.6), and one where R_p is fixed for a finesse of about 4,500 and only R_s is allowed to vary. This second configuration outperforms the first almost everywhere, with a slightly lower threshold, slightly higher peak efficiency, and most noticeably a much gentler drop in efficiency as the Stokes finesse is increased past the optimal value.




Figure 6.9: *Modulation Efficiency vs Pump and Stokes Finesse*

Both curves are generated using 1 W of coupled pump power and 1 atm of deuterium pressure. The blue curve is the same as that in Figure 6.6, and shows the mixing efficiency as both R_p and R_s are adjusted together. The orange curve has R_p fixed at 0.9993 (for a finesse of about 4500), with only R_s varying, and greatly outperforms the blue curve. (I.e., for a given pump power, the orange curve has a higher efficiency at almost all values of finesse and is less sensitive to an "error" in the finesse.) While this is promising for future directions the experiment might take, we are currently unable to implement the configuration associated with the orange curve.

While we found this to be an interesting result, and it shows that this type of system might have room for improvement in the future, we were not able to source mirrors at this time with such precise reflectivities at two different wavelengths at a reasonable cost. The best we could realistically hope for would be approximately equal pump and Stokes reflectivities with a 10% deviation or so from our specified finesse. Given the many factors that were difficult to account for, it was hard to make an entirely analytical decision for our desired reflectivity. We did suspect our finesse was higher than optimal for our amount of coupled pump beam, and we greatly valued the improved locking performance that a lower finesse might bring and believed this might improve our coupled power as well. So in the end, somewhat arbitrarily, we ordered mirrors with a reflectivity of about 0.9992 at the pump and Stokes wavelengths, for a finesse of about 4,000.

6.4 Performance

So counter-intuitive to the proper course of action that a cursory glance at Equation 2.36 might suggest, we decided to change our optical cavity for one that was about three times shorter ($l = 75$ cm to $l = 27$ cm) and with six times lower a finesse ($\mathcal{F} \approx 23,000$ to $\mathcal{F} \approx 4,000$). We ordered these new mirrors from Evaporated Coatings Inc, with the target reflectivities at the pump and Stokes wavelengths, and as transparent as possible in other regions of the spectrum to allow for single pass modulation of mixing beams. We implemented this setup using our old 27 cm cavity and these new mirrors, and were initially able to achieve modulation efficiencies of $1 \cdot 10^{-6}$: a modest increase from our previous result of $0.5 \cdot 10^{-6}$. (It should be considered though that the 27 cm cavity was not able to be cooled with liquid nitrogen, which, as discussed in subsection 3.3.4, likely reduced the performance by about a factor of six.) However, we achieved a significantly higher "ramping" efficiency of 10^{-4} . This ramping process is explained below.

6.4.1 Ramping Efficiency

When aligning the pump laser to the optical cavity, it is often helpful to send a triangle wave driving voltage to either a piezo on a cavity mirror, which will shift the cavity resonances, or to a piezo in the pump ECDL, which will shift the laser frequency. Both processes give approximately linear frequency shifts with voltage, each with a total range at the 10 GHz level. So depending on which piezo is used, the triangle wave effectively sweeps or "ramps" the cavity resonances through the laser frequency, or sweeps the laser frequency through the cavity resonances. Both methods allow the relative frequency spacing of the cavity resonances to be viewed, and can be very useful because it allows laser light to be transmitted through the cavity without it having to be locked.

When we ramp either the pump laser or the cavity piezo, there are brief moments when the laser is on resonance and high intensity builds up in the cavity. This is enough time to generate a Stokes beam, and for modulation to occur of any mixing beam that is present. During these brief periods of resonance, the pump and Stokes beams build up to much higher intensities than when we lock the pump laser, leading to much more power in the mixing beam sidebands. The timescale of the ramps are in the millisecond range, which is very much still in the CW domain compared to the other timescales in the system. Any frequency broadening due to the timescale of these millisecond bursts of transmitted light is much smaller than their linewidth when operating in a purely CW-mode.

We are not entirely sure of the reason behind this behavior, and it was not observed in with the previous 75 cm long cavity. This high efficiency is also observed briefly when the laser lock is first engaged, before dropping down to the steady-state efficiency. It is possible that our locking performance is simply not able to effectively keep the laser on resonance, and that sweeping past a resonant frequency allows more light to couple. While previously this would not be likely, as the cavity resonance linewidths were much narrower

than the laser linewidth, the new cavity linewidth is approximately 150 kHz and on the same scale of the laser linewidth. So perhaps, while the locking system cannot keep the laser frequency completely within this range, when the laser is ramped, it can temporarily couple much more light into the cavity.

Another possible cause for this effect is heating of the cavity mirrors. Due to the very high intensities inside the cavity, there is significant heating of the mirrors, and this may temporarily change their reflectivity properties and curvature and reduce light coupling. These effects would not be present when the laser is only resonant for milliseconds, as heat would not have time to build up significantly.

We look at these time dynamics in more detail in Figure 6.10. Here, a simultaneous trace of the transmitted pump, Stokes, and anti-Stokes beam signals from the redesigned cavity are shown. As the anti-Stokes beam is not resonant with the cavity, it is formed through self-modulation of the pump beam, and so functions in the same way a frequency up-shifted sideband would from an independent mixing beam. The signals are recorded while the frequency of the pump beam is scanned back and forth three times over a cavity resonance. The pump beam (blue) is resonant with the cavity over a period of about 100 μs and builds up to high-power. During each resonance a Stokes beam (orange) is generated through Raman lasing. Note the threshold behavior of the Stokes beam, which reaches peak power very rapidly after the pump passes the necessary intensity. When both the pump and Stokes beams are at high-power, the anti-Stokes beam (green) is generated, which is analogous to a sideband from a mixing beam. The anti-Stokes beam is present for a much shorter period of time, as it requires high intensity from both the pump and the Stokes beams. The peak power of each of these beams is much higher than what is achieved when the pump laser is locked, and allows us to achieve transient modulation efficiencies of 10^{-4} .




Figure 6.10: *Cavity Output Time Dynamics*

The output of the cavity is recorded as the frequency of the pump beam is scanned back and forth three times over a cavity resonance. The pump beam (blue) is resonant with the cavity over a period of about 100 μs and builds up to high power. A Stokes beam (orange) is generated through Raman lasing once the pump hits a threshold power. When both the pump and Stokes beams are at high-power, the anti-Stokes beam (green) is generated, which is analogous to a sideband from a mixing beam. The anti-Stokes beam is present for a much shorter period of time, as it requires high intensity from both the pump and the Stokes beams. The peak power of each of these beams is much higher than what is achieved when the pump laser is locked.

6.5 Looking Forward

This much higher ramping modulation efficiency was very exciting for us, as it allowed for new types of experiments that would not have previously been possible. Additionally, we found the locking performance to be greatly improved compared to the old cavity, presumably due to some combination of reduced finesse and length as we had hoped. Whereas previously the laser would often only lock for seconds at a time, we could now easily maintain a lock for minutes or even hours. (While these are rather poor durations compared to many locking systems, we remind the reader that these experiments are performed with over 10 W of power incident on large, gas-filled cavities, sometimes with a reservoir of boiling liquid nitrogen around it, and so should not be directly compared to traditional laser stabilizing cavities.) The improved locking performance made running the experiment much easier, and would allow for longer and more complex measurements, in addition to the options the higher efficiency enabled.

In preparation for future experiments, we worked to optimize the alignment and minimize the losses from the optics that the mixing beam and sidebands would pass through after leaving the cavity. Having a longer-lasting cavity lock was extremely helpful for this process. When we repeated the modulation experiment of the 785 nm mixing beam, instead of using a beam chopper and lock-in amplifier to detect the sideband, we could simply direct the now plainly visible sideband (without the need to aim it directly into the eye) onto a photodiode. While the overall efficiency within the cavity was only several times higher than in the previous experiment, much more of the modulated light was isolated and usable now. Figure 6.11 shows the 636 nm anti-Stokes sideband from the 785 nm mixing beam with 10s of nW of optical power focused onto a piece of paper. Modulating light in such a stark manner like this—by sending a mixing beam through the cavity in a single pass and having the visible beam of the sideband emerge at clearly a different

color than the input light—had been a long-time goal of mine with the experiment, and so this was a very proud moment for me, and an auspicious sign for our next round of experiments.




Figure 6.11: *A Proud Moment*

With the improved modulation efficiency and locking performance of the new cavity, we were able to isolate the 636 nm anti-Stokes sideband of the 785 nm mixing beam and focus it down so that it was easily visible by eye.

7 Broadband Modulation


With our redesigned cavity setup working well, our next goal was to modulate an already broadband source, rather than the single frequency 785 nm ECDL we had previously used. While we should be able to modulate almost any light source, we wanted to use a pulsed laser. Because pulsed lasers have an inherent spread in frequency space, this would allow us to demonstrate the broadband capability of our modulator. Additionally, using an already pulsed laser as our mixing beam (rather than an incoherent broadband source) and then modulating it, may be a good path towards reaching new limits in pulse duration. This would likely not have previously been possible for us, at least if we wanted to measure the spectrum of the modulated light and do so without substantially more power in the mixing beam—it was challenging enough to detect the previous sideband we produced that had its power concentrated over just MHz-level bandwidth, and now we would need to resolve a sideband over THz-level widths. We decided to use a titanium sapphire (Ti:sapph) laser for our mixing beam because its wavelength matches nicely with the most transparent region of our cavity, we already had a variety of optics that would be suitable for it, and, most importantly, we were able to borrow an old setup from Deniz' former colleague, Dick Prepost.

7.1 The Ti:sapph Laser

The Ti:sapph was an old model from KM labs that housed a variety of optics in precision mounts to create a ring cavity, as well as the titanium sapphire crystal that serves as the gain medium. It is designed to operate in either a broadband pulsed-mode, with pulses at

a repetition rate of 94 MHz (corresponding to the free spectral range of the cavity), or in a CW-mode with a widely tunable narrow lasing-mode. The system needs to be externally pumped, and the manual recommended using at least 5 W of 532 nm laser light that has a very good spatial profile and high polarization purity. The Ti:sapph undoubtedly would have been a pleasure to work with, had it not been for the fact that we had no such pump laser available to us. We eventually found a very good deal on a 4 W green laser that was recently retired from performing in laser light shows. These types of lasers, while well suited for dazzling crowds with mesmerizing displays, are not built to dazzle a physicist with their spectral or spatial quality. The spatial quality was so poor, in fact, that the beam could only be focused to a diameter of about 70 microns, and less than 1% of its power could be coupled to a single-mode fiber. A diagram of the Ti:sapph setup as well as a photograph are shown in Figure 7.1 for reference.

After our best attempts at spatial-mode and polarization cleaning, the pump light enters the system. When the system is correctly aligned, the pump is first directed through a cavity mirror that is transparent at 532 nm, but reflective in the range that titanium sapphire fluoresces. The beam is focused through the crystal, and any unabsorbed light passes through another mirror and into a beam block. The crystal fluoresces and then lases in both directions through the cavity-mode path. The cavity path is designed to accommodate a pair of prisms that serve to correct for dispersion and adjust the bandwidth when the laser is operating in pulsed-mode, or for wavelength tuning when in CW-mode. The beam leaves through the partially transmissive output coupler. The crystal glows a beautiful red when in operation, as seen in Figure 7.1(b).



(a) Diagram



(b) Photo

Figure 7.1: *The Ti:sapph Laser*

The pump light at 532 nm enters the system at the top. It passes through a cavity mirror and focuses on the titanium sapphire crystal, which it primes to serve as the gain medium. The crystal lases in both directions through the cavity-mode and leaves through the output coupler. A pair of prisms serve to compensate for dispersion and adjust the bandwidth when the laser is in pulsed-mode, and is used for wavelength tuning when in CW-mode.

7.1.1 Initial Lasing

While our pump laser was not well-suited to this application, with extensive alignment adjustments we were able to achieve CW lasing at approximately 810 nm, although this was initially at very low power. We had originally hoped that we might be able to pump the Ti:sapph with some 445 nm diodes we had on hand. However, we could only get the system to lase with an output power of a few mW, since our 445 nm diode output only a

couple of watts and titanium sapphire is less absorptive at this wavelength than at 532 nm. Eventually though, we managed to reach an output of about 350 mW using the 532 nm pump beam. Some of the highlights of this progress are shown in Figure 7.2.




Figure 7.2: *Ti:sapph* CW Output Power

The main milestones in our progress with the *Ti:sapph* laser are shown. We initially tried to pump with blue diodes we had on hand, but we could only barely achieve lasing with them. Ultimately, switching to the green pump laser, reshaping the beam, thoroughly cleaning the cavity mirrors, and getting a few new optics, along with extensive alignment adjustments, allowed us get up to an output power of 350 mW.

Despite the poor quality of the pump beam, the *Ti:sapph* output has a very good spatial-mode. This should be expected, as stable lasing can only occur in one of the cavity-modes, which will in general have a clean shape. The pump beam quality simply limits the amount of light that can be pumped in the crystal in the path of the cavity-mode, and so increases the difficulty of achieving lasing. A CCD beam profiler camera was used to record the spatial-modes of the pump beam and the *Ti:sapph* output, which are shown in Figure 7.3.

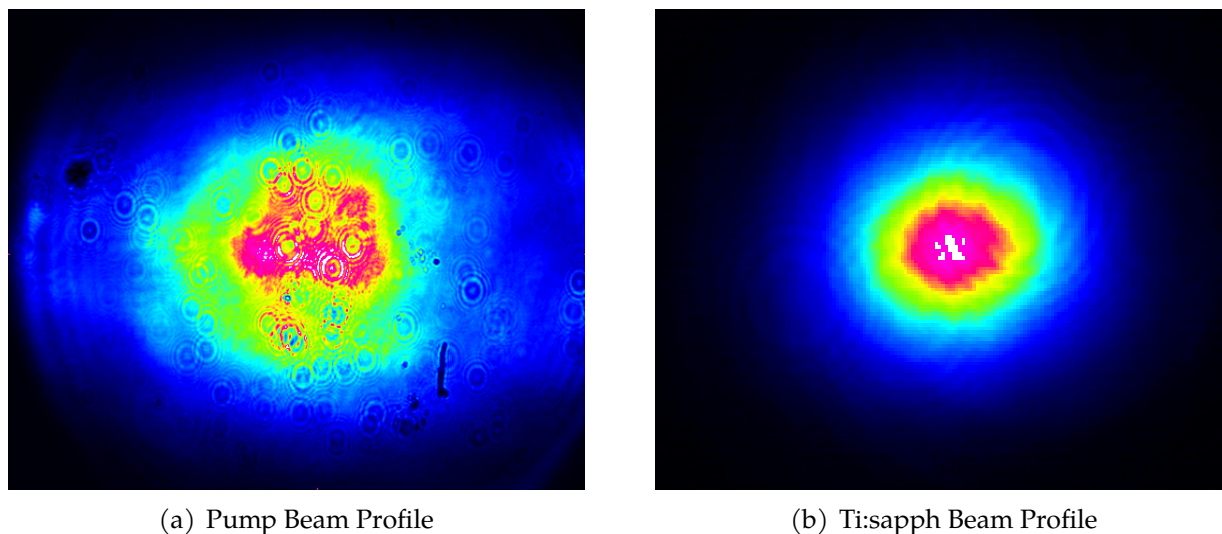


Figure 7.3: *Beam Profiles*

The beam profiles for the pump laser for the Ti:sapph (a), and the CW output beam of the Ti:sapph laser (b). The pump beam is very poor quality, with only about 1% of its power able to be coupled to a single-mode fiber. The output of the Ti:sapph is very good quality with a near perfect Gaussian profile.

7.1.2 Mode-Locking

Achieving pulsed operation with the Ti:sapph proved much more difficult than CW lasing, due to the low-power and bad spatial-mode of the pump. For pulsed operation, the Ti:sapph needs to be "mode-locked", where it simultaneously lases at many different frequencies spaced by the free spectral range of the cavity. Just like as shown in Section 1.1, the different frequencies will interfere with each other and produce short pulses of light followed by intervals of nearly 0 intensity. When mode-locked, the Ti:sapph lases in a different cavity spatial-mode than when in operating in CW-mode. Mode-locking relies on the Kerr effect in the crystal, which will cause high-intensity light to experience a slightly higher index of refraction in the crystal and to be focused more tightly within it. This calls for an adjustment in alignment of other optical elements in the cavity, and results in a change of the cavity-mode. This difference in cavity-mode shape allows the crystal to also serve as

the essential saturable absorber needed to initiate mode-locking: Low-intensity light will not be shaped by the Kerr effect to the new cavity-mode, whereas high-intensity light will be. Random fluctuations in intensity occur as different lasing frequencies come in or out of phase with one another, from which high-intensity bursts propagate more efficiently in the cavity-mode and experience higher-gain, allowing the laser to begin pulsing.

However, with the Ti:sapph lasing at less than its design power, the beam could not be sufficiently shaped by the Kerr effect to form a distinctly new cavity-mode, which lowered the effectiveness of the crystal as a saturable absorber. This made mode-locking very difficult to achieve, and addressing this is what took the majority of the time in getting the Ti:sapph operational. Yet with extensive alignment adjustments (many needing to be outside of the typical range for this setup, meaning the cavity was likely not aligned to the "design" mode), we were able to get a reasonably stable mode-lock.

In practice, to encourage these intensity fluctuations, and to nudge the laser away from operating in a CW-mode and into a stable pulsed-mode, the user bangs on one of the prism mounts with a wrench (yes, really). For an experienced user and, with a little luck, a single well placed "whack" will be enough to mode-lock the Ti:sapph. From the not so lucky days, members of the Yavuz Group will remember the repeated clang of metal on metal echoing from room 5307.

7.1.3 Beam Characterization

At best, when the Ti:sapph was mode-locked, we could get about 150 to 200 mW of output power with a spectral full-width-half-max of about 20 nm. It was easy to verify that it was mode-locked by noting the broadening on an optical spectrum analyzer, or by observing a beat frequency corresponding to the repetition rate when directing the light to a fast photodiode connected to an RF analyzer. During usage, we would commonly send the

beam through a frequency doubling BBO crystal and look for blue light (at approximately 405 nm) as a very quick check for mode-locking. Because the frequency doubling is a non-linear effect and depends on the square of the intensity of the input light, we would get virtually no light if the laser was in CW-mode, but when operating in pulsed-mode, with peak intensities millions of times higher, we would generate about 1 mW of blue light, which is easily visible.

To more fully characterize the beam and exactly measure the pulse widths, we set up an autocorrelator using the BBO crystal. A plot of a typical autocorrelation scan is shown in Figure 7.4. Note that we used an interferometric autocorrelation configuration, rather than the more common intensity autocorrelation setup. This results in visible interference fringes in the plot. Additionally, the background signal will approach 0.1 if the time-axis was scanned further, as opposed to the 0.5 background level for an intensity autocorrelation setup.

The signal recorded is (ideally) proportional to

$$I(\tau)_{\text{phd}} \propto \int_{-\infty}^{+\infty} \left| [E(t) + E(t - \tau)] \right|^2 dt, \quad (7.1)$$

where $I(\tau)_{\text{phd}}$ is the signal recorded at the photodiode as a function of the delay time τ between the two arms of the interferometer. Deconvolving this to recover $(|E(t)|^2)$ for the plot in Figure 7.4, gives a pulse width of about 50 femtoseconds.




Figure 7.4: *Ti:sapph Pulse Autocorrelation Signal*

The recorded signal for the autocorrelation of the Ti:sapph pulses. As this is an interferometric autocorrelation measurement, interference fringes are recorded and the signal has minimums approaching 0 at the center of the pulse. The tails of the pulse will approach 0.1 if the delay range is increased. Deconvolving the plot gives a pulse width of about 50 femtoseconds.

7.2 Experimental Design

With the Ti:sapph well characterized, we were ready to integrate it with the rest of our experiment setup, which is shown in Figure 7.5. The design is similar to that used for the mixing beam experiment that was shown in Figure 5.1.

The deuterium-filled cavity is prepared as before with about 0.3 atm of pressure and the pump beam at 1064 nm. The Ti:sapph beam is collimated and sent through an isolator. The isolator is especially important here, as even μWs of back-reflected light can disrupt

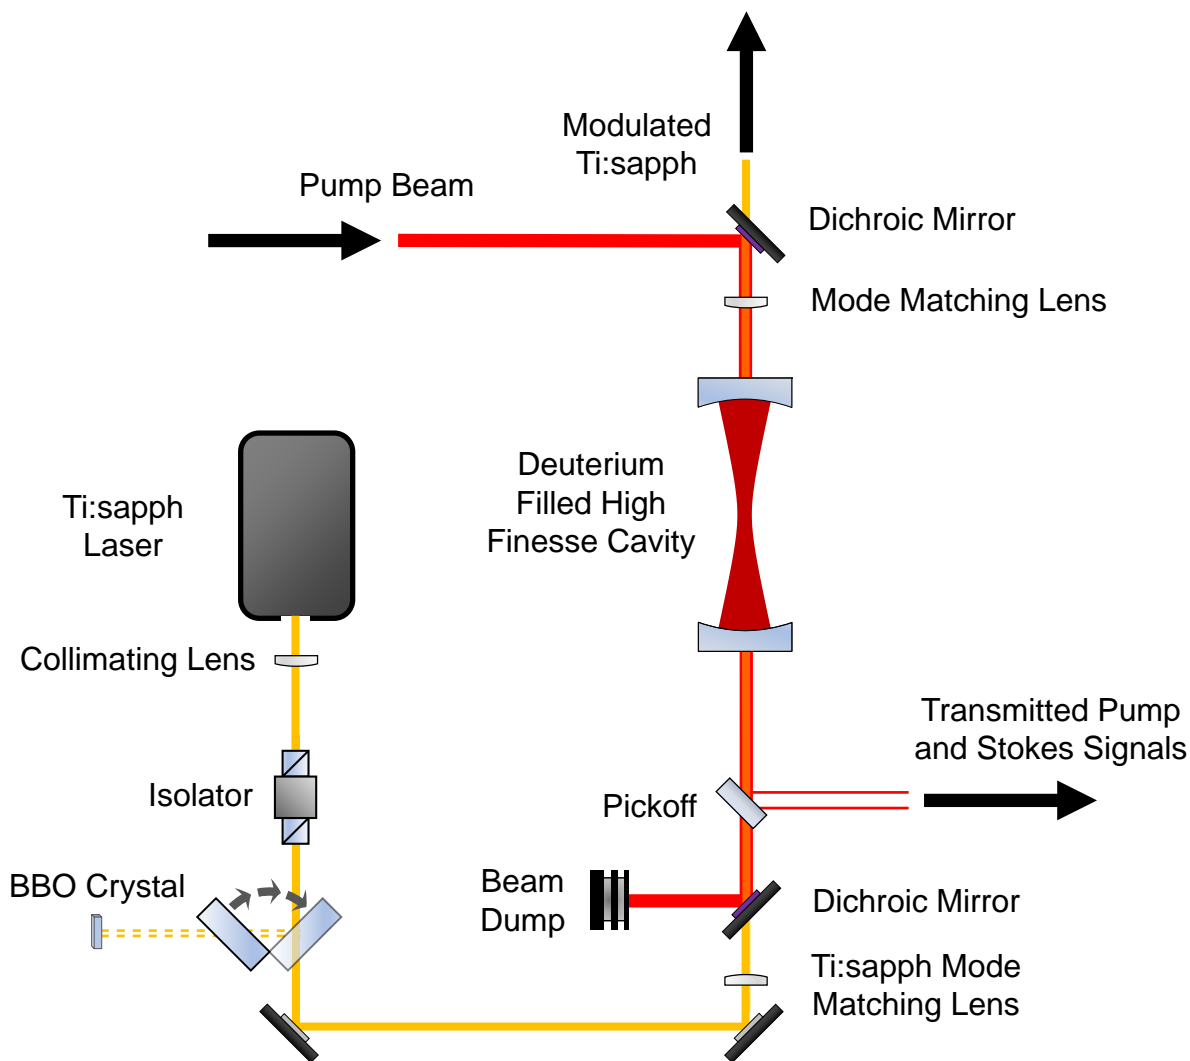


Figure 7.5: *Ti:sapph Modulation Experimental Design*

The setup is similar to that of the mixing beam experiment in Figure 5.1. The Ti:sapph beam is collimated and sent through an isolator, which is essential for preventing disruption to the mode-lock. A flip mirror allows the beam to be directed towards a BBO crystal, which lets us easily check if the laser is mode-locked. The beam is then roughly mode-matched and aligned to the cavity, and is overlapped with the pump beam using a dichroic mirror. The cavity is prepared as before using the 1064 nm pump beam. The Ti:sapph beam and any sidebands are isolated for detection using another dichroic mirror after passing through the cavity.

the mode-lock of the Ti:sapph. A flip mirror is set up right after the isolator to allow the beam to be quickly directed to a BBO crystal, which we found to be one of the easiest and

fastest ways to check that the laser is mode-locked. Past the flip mirror, the beam is roughly mode-matched to the cavity and is overlapped with the pump beam using a dichroic mirror, as in the mixing beam experiment. After passing through the cavity, the Ti:sapph and sidebands are separated from the other beams and sent on to our detection setup.

7.2.1 Detection Scheme

As with the previous mixing beam experiment, the most challenging part with this setup was detecting the modulated light. Although we had slightly more power in the Ti:sapph than in our 785 nm mixing beam, as well as higher modulation efficiency from the cavity, we could not use the same detection scheme as before since this would not allow us to verify that we had achieved broadband modulation. Using a commercial optical spectrum analyzer would be ideal, but our spectrum analyzer needed μW level input powers, which would not be nearly sensitive enough. We tried a variety of detection approaches to get the needed sensitivity in the presence of strong background signals. We had some promising results in building a very sensitive Fourier-transform spectrometer—essentially a Michelson interferometer where the length of one arm is varied and the interferogram is recorded—but ultimately had to abandon the idea due to resolution and stability issues. Stepper motors that could precisely record their position were too jittery and introduced too much noise; electromagnetic actuators were smoother, but we could not track their displacement precisely enough without expensive solutions; and piezos did not quite have the needed range of travel. We ultimately designed and built our own simple grating spectrometer, with a focus on achieving high sensitivity. A diagram of it is shown in Figure 7.6.

The Ti:sapph and any sidebands, or any other light to be analyzed, are sent to the spectrometer setup. The light hits a diffraction grating, and is spread out by wavelength. The system is aligned so that the input light will couple to a single-mode fiber as a diffrac-

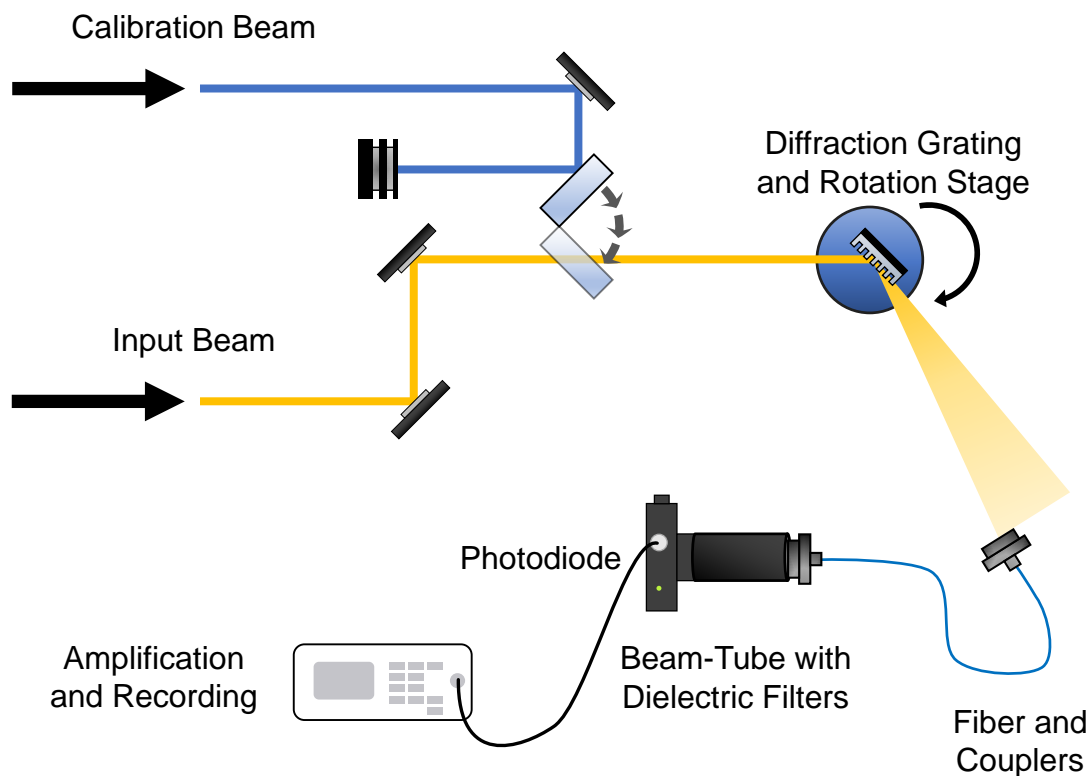


Figure 7.6: *The Spectrometer*

The input light to be analyzed is carefully aligned to the diffraction grating, so that the beam will couple into a single-mode fiber as the grating is rotated. The fiber greatly cuts down on background noise and improves spectral resolution, as only a small slice of the input beam will couple at a given grating position. At the output of the fiber, the beam may be sent through a variety of dielectric filters, if necessary, and the signal is amplified and recorded. A calibration beam of known wavelength or spectral shape may be sent to the diffraction grating as well to periodically recalibrate the device.

tion grating is rotated on a motorized stage. The single-mode fiber blocks much of the background noise, and a playful assortment of dielectric filters after the fiber can help as well. The signal from the photodiode is then amplified and continuously recorded as the diffraction grating is rotated. The rotation angle of the diffraction grating is recorded as well, which corresponds to a known wavelength of light coupling to the fiber. The system is automated and controlled via computer so that, upon initiating a scan, the rotation stage first moves to a designated starting position and then rotates at a fixed rate to a set end

position. The signals from the photodiode and rotation stage are sent to a computer and are processed to show a spectrum. A full scan usually takes a few seconds.

A variety of adjustments can be made to increase the sensitivity of the device, usually at the expense of resolution. Several irises (not shown) between the diffraction grating and fiber couple can be adjusted to change the angular spread of light that reaches the fiber, and effectively change the bin size. Additionally, a multi-mode fiber may be used in place of the single-mode one, which increases the total signal and the bin size, but does increase the background noise as well.

To calibrate the system for correct wavelength readings, a separate "calibration beam" may be sent in via a flip mirror that is aligned to propagate along the same path as the input beam. After slight alignment adjustments, it is typically sufficient to fine-tune the spectrometer by using a single calibration beam of known wavelength and applying the grating equation shown in Equation 3.1. To improve accuracy, an LED of known spectral shape may instead be used, and the resulting spectral data is fit to this known shape. For initial calibration or after major alignment changes, a white light source may also be used as the calibration beam, in conjunction with a monochromator placed between the output fiber and photodiode in Figure 7.6. While the monochromator is adjusted to allow only a narrow range of light wavelengths through, the angular position of the grating for maximum transmitted intensity is recorded for various wavelengths.

Our spectrometer usually has an absolute wavelength accuracy to within a nanometer and a resolution of about half a nanometer, which is adequate for our purposes of detecting modulated light. More importantly though, it is sensitive at the picowatt level of input light. Figure 7.7 shows the recorded spectrum from a HeNe laser that is attenuated to approximately 5 picowatts of power at the input of our spectrometer, in the presence of our pump beam lasing at its full power of 20 W and directed towards the cavity as usual. The peak power occurs at 632.9 nm, which agrees well with the typical value of 632.8 nm

for a HeNe laser. The small secondary peak is likely due to a distorted spatial profile in the beam, which is translated to another local maximum in transmitted power by the fiber. Further closing the irises in the system eliminate this (but causes the signal-to-noise ratio to become too high when at such a lower power). This distortion is possibly introduced by the attenuating optics, as no such extra peaks are present when the laser is not as substantially attenuated.




Figure 7.7: *Spectrum From a 5 Picowatt HeNe Beam*

The spectrum measured using our spectrometer setup from a HeNe laser attenuated to approximately 5 picowatts of power. While background noise is prevalent at this power level, we correctly measure the wavelength of the beam to within 0.1 nm, and its narrow spectral spread is readily apparent. The spectrum is captured with our 20 W pump beam running to match the normal conditions when we operate the experiment. The small secondary peak is a measurement artifact and likely due to a slight spatial distortion of the beam, and was not present during higher-power scans.

We were now confident in our spectrum analyzer's ability to accurately characterize light, and we believed we would have more than enough sensitivity to resolve the characteristics of the Ti:sapph and sidebands. One potential issue though is the dynamic range of the spectrum analyzer, which is limited by our amplifier to a factor of about 100 in power between signals without saturating and while maintaining reasonable resolution. We could record the spectra in two separate scans over a period of minutes, but we wished to record them in a single scan to more definitively show that spectral characteristics of the carrier beam are imprinted on the sideband. Since we expect the sidebands to be between 10^4 and 10^6 times weaker than the carrier beam, we choose to attenuate the carrier beam to approximately the same power level as the sideband. A diagram of the spectrum of Ti:sapph and sidebands as it propagates through the system is shown in Figure 7.8. The unmodulated Ti:sapph first enters the cavity and acts as the carrier beam. Inside the cavity, both frequency up-shifted and down-shifted sidebands are added to it. Only the carrier and the frequency up-shifted sideband exit the cavity however, as the cavity is not very transmissive at the down-shifted sideband wavelength. The beam then passes through a dielectric short-pass filter, which transmits the remaining sideband but significantly attenuates the carrier beam. With the powers between the carrier and the sideband more balanced, the light is sent into our spectrum analyzer.




Figure 7.8: *Ti:sapph Power Balancing*

The power of the Ti:sapph and sideband must be somewhat balanced for simultaneous detection in our spectrometer. The initial broadband Ti:sapph beam enters the cavity, where frequency up-shifted and down-shifted sidebands are imprinted. The cavity mirrors are not very transmissive at the down-shifted sideband frequency, and so only the original carrier beam and up-shifted sideband exit the cavity. The two beams then pass through a filter, which transmits the sideband but significantly attenuates the carrier beam to approximately the same power as the sideband.

7.3 Results

With the spectrum analyzer working well, and the details of the data recording process worked out, we lock the pump laser to the cavity and overlap the Ti:sapph beam with the pump. We were able to get a reasonable amount of modulated power for detection without having to "ramp" the pump beam as described in subsection 6.4.1, so the spectra were captured with the pump in a true CW-mode. To optimize the alignment of the Ti:sapph and the modulation efficiency, we first ran the Ti:sapph in CW-mode to concentrate the sideband power into a much narrower wavelength range for easier detection. After optimizing the setup, we recorded the Ti:sapph and frequency up-shifted sideband in a single scan, shown as the blue curve in Figure 7.9 (a), with the Ti:sapph centered at around 800 nm and the

sideband near 645 nm. The sideband actually has a higher-power than the carrier beam here due to the limited control we have over attenuation, so a separate scan of the unattenuated carrier is superimposed in orange and normalized to the same power as the sideband, to more clearly show the input spectral features. Both beams show up as narrow peaks, within the resolution of our spectrum analyzer, and spaced by 90 THz just as would be expected.

We were now ready for the main experiment of modulating a broadband source. We mode-lock the Ti:sapph beam, and also adjust the Ti:sapph prisms to give a small spectral peak in addition to the main broad peak. This is to provide an additional reference feature so that it is clear that the modulator fully copies the spectral properties of the carrier beam. We repeat the same procedure as before, and the result is shown in Figure 7.9 (b). Note that the broad spectrum of the carrier beam and the sideband, as well as the additional peak, and that the relative intensities between the primary and secondary peaks are preserved. The up-shifted peak appears closer to the primary peak in the sideband, but this is only an artifact of plotting versus wavelength. In a frequency view, all the original spacing are preserved.

To verify that the sideband in the mode-locked case is not just broad but has maintained its discrete spectral components, we focus the sideband onto a fast photodiode and examine the output on an RF spectrum analyzer. The result is presented in Figure 7.10, which shows a strong and narrow beat at about 94.12 MHz. As expected, this corresponds to the free spectral range of the Ti:sapph cavity, which sets the repetition rate of the Ti:sapph laser.

Finally, we also investigate the modulation efficiency of the system as a function of deuterium gas pressure in the cavity. The Ti:sapph is set in CW-mode and the maximum measured modulation efficiency is recorded at a given gas pressure. Adjusting the gas pressure will affect the locking performance of the system, and requires a readjustment of a variety of parameters. Additionally, a change in gas pressure will slightly misalign the

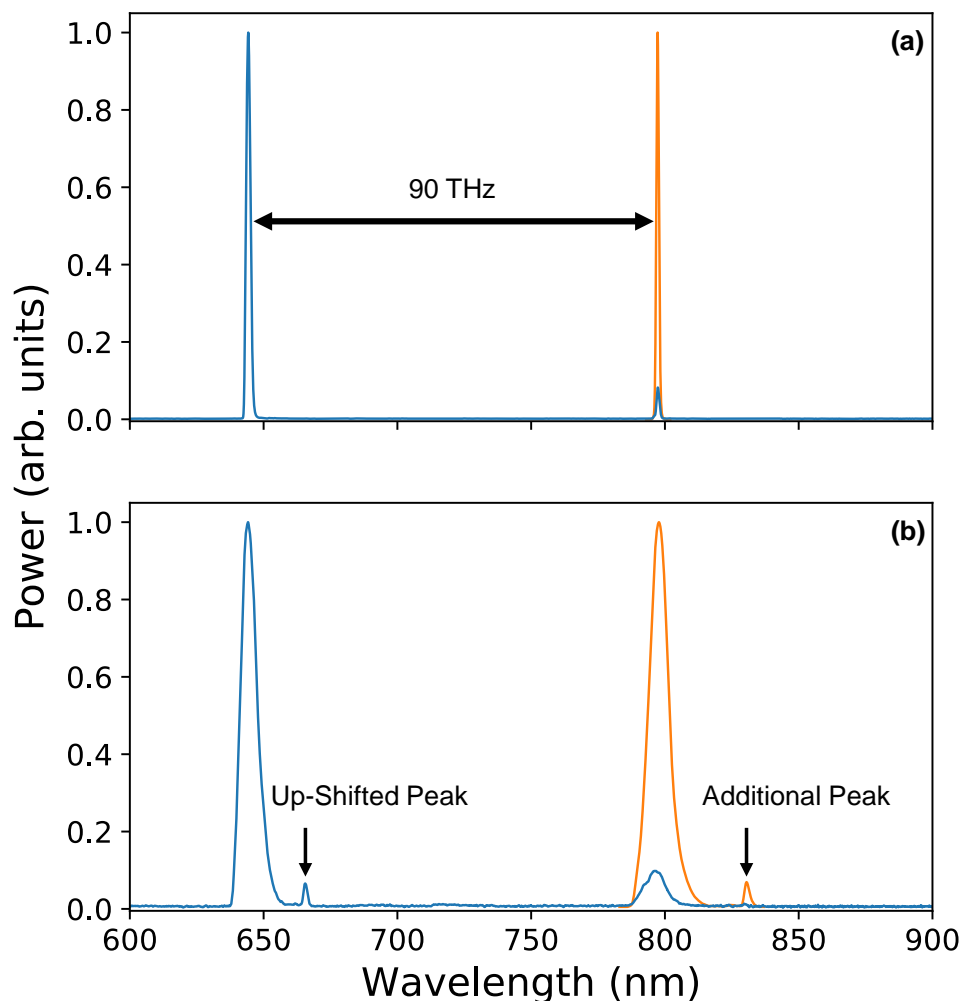


Figure 7.9: *Modulated Ti:sapph Spectrum*

The Ti:sapph and sideband spectra when in CW-mode (a), and when mode-locked (b). In each plot, a separate scan of the unattenuated Ti:sapph beam is superimposed and normalized to the sideband power to allow for more easy comparison of their spectral features. The spectral width of the carrier Ti:sapph beam are preserved in the sideband, which is frequency up-shifted by 90 THz.

cavity from beam, which upon correcting results in a slight misalignment of the Ti:sapph beam from the pump. Correcting these effects can be very time consuming, and it can be difficult to verify that the system has been fully optimized. As a result, the efficiency values recorded are considered a lower bound for a given pressure. We have found that extensive




Figure 7.10: *Ti:sapph Sideband Repetition Rate*

The signal on an RF spectrum analyzer from the frequency up-shifted portion of the spectrum shown in Figure 7.9 (b), when detected on a fast photodiode. Due to the discrete-mode nature of the spectrum, there is a strong beat at 94.12 MHz. Within our measurement accuracy, this value coincides with the measured repetition rate of the Ti:sapph laser.

tuning can typically give efficiency improvements of about 30%, and so the error bars are one-sided to reflect this.

The Raman lasing threshold occurs at a pressure of about 0.9 atm. The mixing efficiency peaks at slightly higher pressure and stays relatively constant until a gas pressure of about 2 atm, at which point the efficiency slowly declines with increasing pressure. This is in disagreement with the predictions made in Figure 6.2, which suggest a continuous increase in efficiency. However, as noted in Chapter 6, those kinds of efficiency gains are unlikely to be realized due to decreased locking performance at higher pressure, which is likely the primary cause of the declines in efficiency shown here.

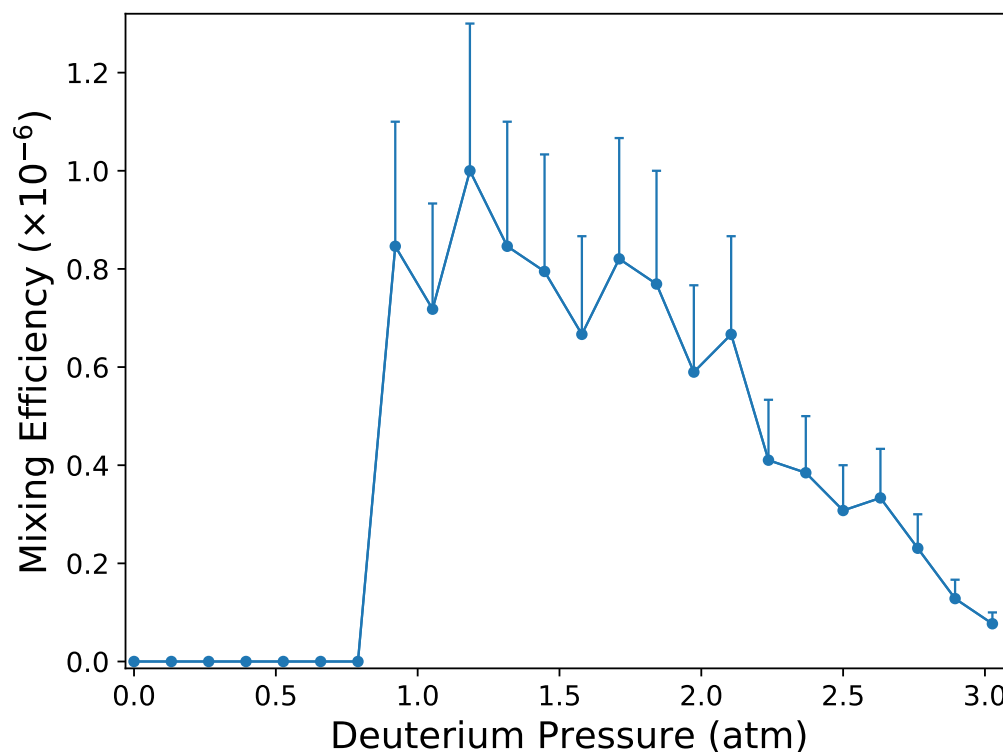


Figure 7.11: *Ti:sapph* Mixing Efficiency vs Deuterium Pressure

The mixing efficiency for the *Ti:sapph* beam when the pump beam is locked to the cavity.

After crossing the threshold pressure needed to induce Raman lasing, the efficiency rapidly rises before somewhat plateauing and decreasing. As extensive adjustments are needed to maximize efficiency at a given pressure, the maximum efficiency observed is recorded, and the error bars are one-sided to reflect a 30% increase in efficiency that is typical upon extensive tuning.

We were very excited about these results, as they showed that our modulation techniques could be successfully applied to broadband pulsed sources. Additionally, while we were somewhat limited by our *Ti:sapph* laser, we note that the spectra of Figure 7.9 have immediate important implications for ultrafast waveform synthesis. Below in Figure 7.12, we plot the calculated intensity of a pulse for the modulated *Ti:sapph* laser along with that of an unmodulated pulse. For these calculations, we take the carrier beam to be centered at 810 nm with spectral components spaced by 94 MHz and with powers given by

a Gaussian profile with a full-width-half-max of 20 nm. Components corresponding to the anti-Stokes sideband, frequency up-shifted by 90 THz and centered at 652 nm, are added in, and we also include the downshifted Stokes components (centered at 1070 nm) that would be present if the cavity was transmissive at those wavelengths. Finally, we balance the powers of the three separate beams and set the phases to be equal, corresponding to amplitude-control and dispersion-compensation.




Figure 7.12: *Calculated Ti:sapph Pulse Shapes*

The calculated pulse shape for the unmodulated Ti:sapph beam (a) and the modulated beam (b). The modulated beam contains components corresponding to both the Stokes and anti-Stokes sidebands with amplitude control and dispersion compensation. The modulated beam shows narrow features about 2.5 fs wide underneath the same envelope of the unmodulated case.

Figure 7.12 (a) shows the calculated pulse shape with no modulation, which has a width of about 50 fs. Plot (b) shows the modulated pulse shape. While the envelope remains

the same, there are now many clusters of narrow features underneath it with widths of about 2.5 fs and spaced by about 5.5 fs. In both cases, the pulses are spaced by about 10.6 ns, coinciding with the repetition rate of the cavity.

It is also worth noting that the unmodulated pulse length of 50 fs is relatively modest for a Ti:sapph laser, and so the sideband components were far from the carrier in frequency space. However, many commercially available Ti:sapph lasers can have pulse widths of less than 10 fs. Using one of these as the input beam would result in the sidebands and carrier beam overlapping in their spectral widths and completely covering the 550 to 1250 nm spectral range. With appropriate amplitude control and dispersion compensation, such a coherent spectrum would have the capability to produce sub-cycle, 2 fs long pulses, which would approach the shortest pulses produced in the optical region of the spectrum. We published these results and a description of the experiments presented in this chapter in *Optics Letters* in 2018 [37].

8 Microresonators: A New Direction

Although we were excited by the results we had produced in modulating the broadband Ti:sapph laser, we had begun to feel that any utility this might have would be fundamentally limited by the design of our modulator. The optical cavity that was necessary to achieve the high intensities required for modulation was large and needed many supporting optics, covering most of an optical table. It was not clear how this could one day be compacted into a reasonably sized device that could serve as the basis for an optical waveform generator. Additionally, the modulation efficiency was low, and, even after a redesign of the cavity, detecting the sidebands that the system produced remained challenging. To address these issues, in 2017 our group started a collaboration with Professor Dan van der Weide in the Electrical Engineering Department at UW-Madison. Dan and one of his students, Atul Bhadkamkar, had been working with glass microresonators—small spheres that support whispering gallery modes via total internal reflection—which we believed we could extend our modulation techniques to.

These microresonators had a few key properties that could help address many of the challenges and limitations of our optical cavity approach to modulation. First, they have very small mode-volumes that are on the order of only one wavelength wide and deep as they travel around the circumference of the resonator. This is much smaller than can be achieved in optical cavities, and so can lead to much higher intensities and potentially higher modulation efficiencies. And second, their resonances are based on the geometric effect of total internal reflection, rather than reflections off of mirrors via dielectric coatings. This leads to high-finesse across a wide range of wavelengths, and so a single resonator could work for virtually any pump or mixing laser, rather than requiring many different resonators

with separately engineered dielectric coatings. This inherent broadband nature, along with their small size, gives them good potential to be developed into a useful laboratory tool in the future.

At the same time, our group member Josh Karpel began work on simulating the interactions of electric fields with these types of resonators. He found that with reasonable experimental parameters, we might achieve modulation efficiencies of more than 1%. A plot adapted from his thesis is presented in Figure 8.1, which shows a predicted modulation efficiency heatmap as a function of available pump power and separation of the resonator from an optical fiber (which influences the coupling of the pump to the resonator) [38].

With the allure of all these features and high efficiency in mind, we began the process of adapting our experiment to work with microresonators and developing the necessary skills and techniques to make use of them. Atul's help during our work with microresonators was invaluable, as he fabricated resonators to meet our design criteria, as well as supplied us with very useful knowledge and equipment. We originally began working with glass microspheres that he supplied for us, which require a high-power CO₂ laser to produce. The laser is focused to the tip of a bare fiber optic cable, which heats up significantly and begins to melt. As the glass melts, surface tension naturally pulls it into an approximately spherical shape, with diameters in the range of 100 microns, and can be very smooth and have a high finesse. These resonators primarily support whispering gallery modes, which are discussed below.




Figure 8.1: *Predicted Modulation Efficiency vs Pump Power and Fiber Separation*

Adapted from Josh Karpel's thesis for a spherical microresonator of 50 micron radius and quality factor of 10^8 . The pump is set at 1064 nm and the mixing beam is 800 nm, similar to previous experiments we have performed. The mixing beam is modulated to a "target" mode about 12 THz away. The colored dashed lines indicate the fiber separation where critical coupling occurs for the corresponding mode, but due to the interaction of many different modes with different couplings, there is no special behavior at these points. A wide range of parameters allow for modulation efficiencies of greater than 1%.

8.1 Whispering Gallery Modes

Whispering gallery modes can describe the resonances of a variety of types of waves as they travel around a curved surface. Originally named for sound waves as they traverse the dome of the St. Paul's Cathedral in London, they have recently generated great interest for their application to light in microresonators. Microresonators can come in a variety of shapes that all support different forms of whispering gallery modes. The basic physics of them is the same though, and so we restrict this discussion to the spherical case. Unlike linear

optical cavities that have a fixed path length and whose mode-shape can be described by two mode-numbers (generally l and m for Hermite-Gaussian modes), whispering gallery modes have a variable path length depending on the radius at which they propagate. As a result, they are defined by three mode-numbers— q , l , and m , as well as TE or TM polarization. Modes are typically denoted as $TE(M)_{q,l,m}$. q corresponds to the number of intensity maxima in the radial direction in the equatorial plane, l is the number of wavelengths as light travels around the mode-circumference at the propagation radius, and m is the number of maxima in the azimuthal direction and so must be between $-l$ and l (a negative m indicates clockwise propagation). $l - |m| + 1$ gives the number of maxima in the polar direction. So modes where $l = |m|$ correspond to a mode with a single maxima along the equatorial plane. (To see this, consider that for $l \neq m$ there must be at least 2 values of the polar angle θ where an integer number of wavelengths fit in the meridian around the sphere at θ .)

In practice, whispering gallery modes in optical resonators are generally thought of as $q = 1$ and $l, m \gg 1$, with l and m close or equal in value. $q = 1$ is typically what is easiest to couple light into, since it has the maximum amount of the mode-intensity near the surface of the sphere where it is accessible by most coupling schemes. $m = l \gg 1$ are modes that are well confined to near the equator, which is desirable for many applications including ours, since it minimizes the mode-volume and maximizes the intensity of coupled light. Outside the surface of the sphere, mode-amplitudes experience rapid exponential decay. Inside the sphere, the mode-shape is described by

$$E_{\zeta}(r, \theta, \phi) = E_0 P_l^m \cos(\theta) e^{im\phi} j_l(k_{1,\zeta} r), \quad (8.1)$$

where P_l^m are Legendre polynomials, j_l are spherical Bessel functions, $k_{1,\zeta} = \frac{n\omega_{\zeta}}{c}$, and ζ is the index $P_{q,l,m}$ with $P = TE(M)$ as the polarization. Applying appropriate boundary

conditions to the electric field and expanding the Bessel function to second-order lets us find the mode-frequencies as

$$\omega_\zeta \approx \frac{c}{an} \left[\left(l + \frac{1}{2} \right) + a_q \left(\frac{l + 1/2}{2} \right) - \Delta^P + \frac{3a_q^2}{2^{2/3} \cdot 10(l + \frac{1}{2})^{1/3}} \right]. \quad (8.2)$$

Here, a is the radius of the microsphere, a_q is the q th zero of the Airy function, and $\Delta^P = \frac{1}{\mu_r} \frac{n}{\sqrt{n^2-1}}$ for $P = \text{TE}$ and $\Delta^P = \frac{1}{\epsilon_r} \frac{n}{\sqrt{n^2-1}}$ for $P = \text{TM}$. Note that m does not affect the frequency and so there is $2l + 1$ degeneracy. In the lab, though, a microsphere will always have some degree of eccentricity, which lifts this degeneracy. If we additionally consider the case where $q = 1$ and $l \gg 1$, we can drop the second-order terms for a_q . We can then write the frequencies as

$$\omega_\zeta \approx \frac{c}{an} \left[\left(l + \frac{1}{2} \right) + a_1 \left(\frac{l + 1/2}{2} \right) - \Delta^P \pm \epsilon^2 \left(\frac{l - |m|}{2} \right) \right], \quad (8.3)$$

where $a_1 \approx 2.338$ and ϵ is the eccentricity of the sphere, with the $+$ sign used for an oblate sphere and the $-$ sign used for a prolate sphere.

A analogous free spectral range to that of an optical cavity would be $\text{FSR} = \frac{c}{2\pi a}$ and would correspond to a change in l by 1. This is typically tens to hundreds of GHz in our experiments, and is larger than we can smoothly ramp the laser over. However, unlike a cavity that can be aligned in such a way that only the fundamental modes are excited and are spaced by 1 FSR, there is not an easy way to selectively excite certain modes to this precision in the microresonators. Generally then, we would view only a small portion of a free spectral range as we scan our laser frequency, but there will still be many resonances present due to m splitting. This splitting is in fact very advantageous, since it could be quite difficult to tune a laser enough to find one resonance in a range of several hundred GHz. Additionally, unlike cavities where higher-order modes significantly distort the

mode-shape, there is not nearly as much of a change in shape or of a drop in peak intensity as m changes at high l .

8.2 Tapered Fibers

Efficiently coupling light into microresonators can be challenging, since the modes have such small volumes and have such different spatial profiles than the typical Gaussian laser beam. As a result, we found that we could couple less than 1% of the light from a free space laser beam directly into a resonator, and so it is not a practical coupling method. Instead, prisms or tapered optical fibers are most commonly used to produce an evanescent field, which allows light to tunnel into the resonator. While we briefly experimented with prism coupling early on, we quickly realized that tapered fibers would be better suited for our setup, given the small size of our resonators. Tapered fibers, which we affectionately refer to as "tapers", are bare single-mode fibers that are stretched so that they transition from a typical core and cladding diameter of approximately 100 microns to just several microns. For a laser propagating through such a fiber, this allows much more of the evanescent wave that normally would rapidly decay in the cladding to reach an appropriately placed nearby resonator.

In addition to allowing significant amplitudes of evanescent waves to leave the fiber, the tapering process also greatly increases the numerical aperture in the tapered region of the fiber, which can be problematic. Numerical aperture (NA) is given by

$$\text{NA} = \sqrt{n_{\text{core}}^2 - n_{\text{cladding}}^2}, \quad (8.4)$$

with typical values of about 0.1 to 0.2 for single-mode fibers. For tapered fibers though, the core and the cladding fuse, and the surrounding medium (typically air) acts as the

cladding. For glass having an index of refraction of 1.45 and air of 1.00, we find the NA of the tapered fiber to be about 1. This leads to many more propagation-modes, M , in the fiber being supported. For a typical step-index fiber that would approximately model the fiber-air interface for a tapered fiber, M is given by

$$M = 2 \left(\frac{\pi}{\lambda} r \text{NA} \right)^2, \quad (8.5)$$

where r is the radius of the fiber core and λ is the wavelength of propagating light. It is generally desirable to avoid the propagation of higher-order modes, as they can result in worse coupling to the resonator (typically, each mode cannot simultaneously have its coupling optimized) and cause interference effects when light couples back into the fiber. To avoid this, the tapered fibers are ideally produced "adiabatically" (i.e. with a small angle from the edge of the fiber where the tapering region begins, to the edge of the fiber where the diameter is at a minimum). This ensures that light in the original spatial-mode of the fiber will remain relatively unperturbed as it enters the tapered region, with minimal excitation of higher-order modes. We found as well that having strongly adiabatically tapered fibers typically led to higher transmission. We would often get as high as 90% transmission through our best fibers (meaning, approximately 10% scattering losses primarily due to the tapered region of the fiber), whereas lower quality ones might only get 60%.

8.2.1 Bi-Tapered Fibers

Tapered fibers may be mono-tapered, where they end in a small point at the minimum diameter, or bi-tapered, where they then transition back to their full un-tapered size. In the case of mono-tapering, two fibers are needed, one which couples light into the resonator, and another which allows light to couple back out. The main advantage of this is that

nearly all the light that is transmitted through the second fiber will have interacted with the microresonator (i.e. there is no transmission when the laser is off-resonance and maximum transmission when on-resonance), which can simplify optimization procedures and reduce background noise. There is no practical way to otherwise isolate this "light of interest".

On the other hand, only a single bi-tapered fiber is needed for coupling, but maximum transmission occurs when the laser is off-resonance (light simply travels through the fiber unimpeded and is detected), and minimum transmission occurs on-resonance (where large amounts of light will couple into the resonator and slowly radiate out; some of the light returns to the fiber where it will destructively interfere with the other propagating light and cause transmission dips). This can complicate initial coupling detection and subsequent optimization, and will have much higher background noise even when working well. However, the coupling is much more stable since there will not be two loose fiber ends which can easily move, and only opto-mechanics to move one single fiber are required, rather than two independent sets. Additionally, as we hoped to eventually couple large amounts of power into the resonators, we were not thrilled with the idea of several watts of laser light shooting into the air if the laser moves off-resonance. With a bi-tapered fiber, off-resonant light safely continues down the fiber where it may easily be disposed of (please consult your local waste authority for information on how to properly dispose of laser light). For these reasons, we chose to work with bi-tapered fibers in our experiment, and so "tapers" and "tapered fibers" refer to a bi-tapered fiber from here on, when not otherwise noted.

Finally, although we did not initially consider this, we suspect that we are able to achieve a higher fraction of coupling using bi-tapered fibers than we would with mono-tapered fibers. While it is well known that the theoretical maximum coupling of a fiber to a microresonator will occur at the critical coupling point, where this is a small air-gap between them (typically hundreds of nanometers to a few microns), we found that we had

the best coupling in the over-coupled regime, where the fiber and resonator were in contact. Additionally, we found that moving the two even further past this point, so that the fiber bends and wraps around the resonator somewhat, resulted in even better coupling.

We believe that the wrapping of the fiber increases the length over which it is in contact with the resonator (or perhaps the length where this an appropriately sized air-gap), and so increases the probability of a given photon coupling and the overall coupling efficiency. Similar effects have been reported by other groups and studied to an extent, although usually in the context of transmission losses due to fiber bending [39]. While we observe this in our own setup as well, generally the increase in coupling far outweighs the drop in absolute transmission. A plot of our observations for one of our resonators is shown in Figure 8.2. Much like in the pressure vs mixing efficiency plot in Figure 7.11, we know that the coupling efficiency cannot be lower than the maximum observed value, and so we use one-sided error bars. As before, extensive adjustments can usually increase this efficiency further. We roughly estimate the uncertainty as 30% and display it only on the relative-depth curve for clarity. Several hundred microns of additional translation are needed after the resonator and fiber first make contact before significant coupling occurs. This suggests that the fiber wrapping around the resonator is very important for coupling in our system. This method of coupling would likely not have been possible for mono-tapered fibers, or at least not to the same degree. So recording these results, long after our initial decision to use bi-tapered fibers, helped to further strengthen our view that bi-tapering was the right choice for our setup.

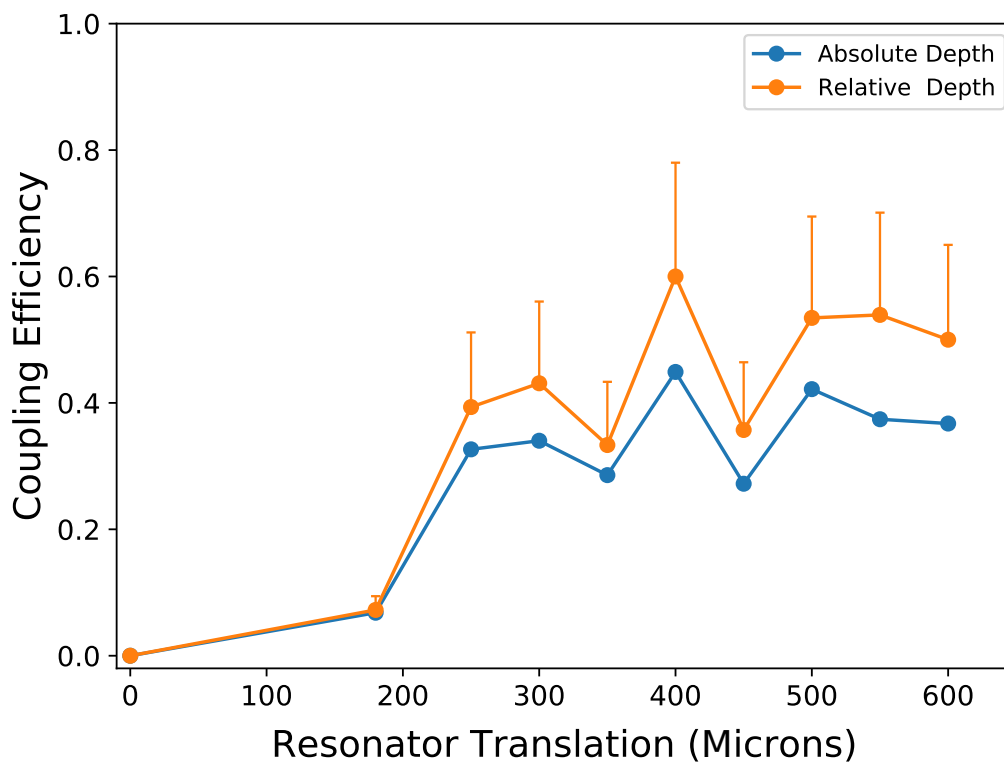


Figure 8.2: *Coupling Efficiency vs Resonator Translation*

The coupling efficiency for a microresonator as it is translated into the fiber, causing the fiber to slightly bend and wrap around. 0 microns refers to the position where the resonator first makes contact with the fiber. The absolute resonance depth is the ratio of coupled power to transmitted power before the resonator was brought near the fiber. The relative resonance depth accounts for the slight loss in transmission as the fiber is bent.

8.2.2 Fabrication

Tapering fibers is ideally done with specialized equipment such as fiber fusion splicers, which can precisely apply heat and tension to a fiber to soften and stretch it to the appropriate diameter. We did not originally have access to this equipment, so in an effort to keep our progress moving forward, we experimented with techniques to manufacture tapered fibers ourselves with simpler equipment. We found that standard butane torches can easily produce a hot enough flame to soften a fiber. Applying the correct amount of tension,

however, is more difficult. Less than 1 N of force is needed to start a fiber stretching when it is sufficiently heated, and the necessary force drops once stretching begins and the fiber diameter starts to decrease. Once the fiber diameter has reduced to only a few microns, too much tension can easily break it. Stopping the tapering process is also a delicate affair, as removing tension too rapidly often results in a slight kink forming in the fiber, which then typically causes the fiber to snap when heat is removed and it becomes more brittle. To avoid this, once the appropriate amount of tapering has occurred, tension must quickly but smoothly be adjusted to keep the fiber taut, but stop any further stretching. The tapered region may then be secured to a mount to protect it and keep it taut.

Some techniques with which we had relative success involved anchoring the fiber to a post on an optical table and applying tension to the other end via a pulley and a small (approximately 50 g) mass off the edge of the table. The operator then applies a very small butane flame to the center of the fiber, smoothly sweeping it over several cm. The fiber heats up in a few seconds and begins to stretch. It is important to not let any of the stretching region cool, as reheating and stretching again typically causes the fiber to break; Thus, the entire process needs to be done in one sequence. When the operator judges that the fiber has been appropriately tapered, they remove the heat and secure the edges of the tapered region to a mount using epoxy, which greatly reduces the risk of breakage during transport and use. A similar technique involves using a hand-operated translation stage to apply continuous tension to the fiber. In both cases, sending a visible laser through one end of the fiber can be very helpful for judging the length and diameter of the tapered region of the fiber. When the fiber diameter approaches the 10 micron and lower range, significant amounts of light can be seen leaving the fiber, as it begins to evanescently leak to the surface and scatters off of impurities (this is somewhat visible in Figure 8.5). The finished taper can be measured under a microscope to verify the diameter.

Surprisingly, we had the most success in producing these fibers when working entirely

by hand. With this method, one end of the fiber is anchored as before, and the operator applies the flame with one hand while applying tension with the other. While certainly less precise than the other techniques, this method does allow an experienced operator to rapidly adjust the amount of tension applied in response to the rate at which the fiber is stretching. However, this is a very time consuming skill to develop, very stressful to perform, and arguably has very little long-term utility. Any future students of this project would be advised to avoid resorting to any of these methods if possible. The following theses were found to have helpful advice on fiber manufacturing and coupling techniques as well [40, 41, 42].

Fortunately, after several months of using these techniques, we learned that Professor Randall Goldsmith of the Chemistry Department had a fiber splicer, and was willing to let us use it regularly. Atul became a skilled operator of the fiber splicer, and has supplied our group with tapered fibers for over a year, for which we are extremely grateful. Without access to a fiber splicer in the future, a fruitful technique to pursue might be to set up a motorized translation stage or winch to apply tension to the fiber. While ideally a variety of sensors would allow for the rapid adjustment of tension as in a fiber splicer, simply setting the motor to run for a set amount of time so that the operator only needs to learn to apply the correct amount of heat during that interval, would likely be able to produce reasonably consistent results after some practice.

Different resonators are very sensitive to the diameter of tapered fiber that is used to couple light to it, and so a variety of factors affected the ideal width we wanted over time. We generally did not have precise control over the diameter of fibers produced, but we would measure the diameter of all our fibers to narrow down the range of what to aim for. A photo of a typical tapered fiber as viewed under a microscope is shown in Figure 8.3. The reticle scale is approximately 24 microns per major tick mark, so the fiber is about 2 microns in diameter.




Figure 8.3: *A Bi-Tapered Fiber*

A typical tapered fiber as viewed under a microscope. The narrowest section of the fiber is in the frame, which was usually about 0.5 cm long. The reticle scale is approximately 24 microns per major tick mark at this magnification, so about 300 microns of the fiber is visible and the fiber is about 2 microns in diameter. Out of frame, the fiber gradually widens at both ends to its untapered diameter of about 120 microns.

8.3 Microspheres

After securing a steady supply of tapered fibers to work with, and developing our methods to characterize them, we focused our efforts on coupling to microresonators. We began using microspheres because they were easier to manufacture, although we suspected they would not be suited to our purposes in the end—being made of silica, they have a relatively low Raman-gain coefficient compared to other materials—but would be fine for developing our coupling techniques and measurement methods.

Since the tapered fibers are very fragile, they need to be securely held in place to avoid damage. Additionally, for efficient coupling, a resonator must be able to be freely moved with respect to the taper. To this end, Matt Dwyer of Professor van der Weide's group expertly designed and machined several reusable "fiber forks" that the tapers may be adhered to, as well as a mount for them on the optical table. A render of the final design is shown in Figure 8.4. Tapers are adhered to the fingers of the fiber fork, so that the entire tapered region is in the inner region of the fingers and is relatively taut. This helps to keep

the tapered region stable with respect to the resonator, and provides significant protection, since the epoxy prevents most tension on the ends fiber from being transmitted to the delicate taper. When a taper has reached the end of its life-cycle, the epoxy can be melted away, the fiber removed, and the fork is reused. A commercial 3-axis translation stage with a resolution of about 1 micron in each direction is mounted at the other end of the holder and supports the resonator stage. A resonator is affixed to the stage and is translated with respect to the taper, which is fixed in place. A microscope view of a microsphere resonator that is coupled to a tapered fiber using this setup is shown in Figure 8.5.




Figure 8.4: *The Tapered Fiber and Resonator Holder*

Tapers are adhered to the fingers of the fiber fork so that the entire tapered section is in the inner region of the fingers. A commercial 3-axis translation stage with a resolution of about 1 micron in each direction is mounted at the other end of the holder and supports the resonator stage. A resonator is affixed to the stage and can be translated with respect to the taper, which is fixed in place.




Figure 8.5: *A Silica Microsphere*

A silica microsphere and tapered fiber as viewed under a microscope. At this magnification, the major tick marks on the reticle scale are each 100 microns, so the microsphere has a diameter of about 200 microns and the fiber is less than 10 microns in diameter. A red laser is transmitted through the fiber, and some scatter can be seen as light leaks to the surface of the fiber and also from the surface of the microsphere as it resonates inside.

8.3.1 Experimental Design

Most measurements and experiments involving the microresonators were made using slight variations of the setup shown below in Figure 8.6.




Figure 8.6: *Microresonator Experimental Design*

The pump laser is generated from a 1064 nm diode, which passes through an isolator, a fiber amplifier, and a second isolator. A "monitoring" signal is then picked-off for comparing with the transmitted light signal. The laser is coupled to a bare fiber, which leads to the fiber fork and resonator holder. The output of the fiber is collimated, and a small portion is picked-off to generate the transmitted signal. The remaining light is blocked or sent on for additional analysis or for use in other experiments.

As in our previous experiments, the pump beam is generated from a 1064 nm diode, which is then sent through an isolator and fiber amplifier. Unlike with the optical cavities however, the microresonators could be pumped with almost any wavelength of light, as they will have numerous resonances across nearly the whole optical spectrum. We chose to pump with 1064 nm, simply because we had an amplifier at that wavelength and were experienced with the diodes. The beam is sent through a half-wave plate for polarization control, and is aligned to a fiber coupler. A small fraction of the input beam is picked-off right before the fiber coupler to generate the "monitoring signal" for comparison with the transmitted signal. The beam is coupled into the bare fiber, which leads to the fiber fork and resonator holder shown in Figure 8.4. The output of the fiber is collimated, and a

portion is picked-off to generate the transmitted signal. The remaining light is blocked or sent on for additional analysis or for use in other experiments.

To couple light into a microresonator, the pump laser piezo is ramped so that the laser frequency is scanned. The resonator is then brought close to the fiber using the 3-axis translation stage, and slowly lowered down. Touching the taper with the resonator with too much force risks snapping the taper, so it is important that the resonator is lowered very slowly and gently, and stopped as soon as contact is made. Typically, carefully watching the transmitted signal while lowering the resonator is enough precaution, because once they make contact the transmitted signal drops by at least 5% due to scattering losses. Another good method is to use an IR viewer and look at the resonator while lowering it. As soon as contact is made, the crystal will visibly glow with scattered light. At this point, although some light couples out of the fiber, very little light will typically resonate in the crystal—most is simply scattering on the surface of the resonator, and so only introduces a DC drop in transmitted signal with no frequency response. A variety of parameters must be iteratively adjusted to get good coupling and detect resonances on the transmitted signal—most significantly, the polarization of light and the x , y , and z position of the resonator with respect to the fiber (the rotational orientation is coarsely adjusted by hand, but does not typically need additional adjustment). This is one disadvantage of the bi-tapered fibers we use, as there will be a large DC background that can make initially detecting and optimizing resonances very difficult. Using two mono-tapered fibers can allow the transmitted signal to be almost entirely resonant light, but this introduces a variety of other complications. With practice, the difficulties of the DC background can be mitigated, and as previously discussed, we think the benefits of bi-tapering far outweigh this and other disadvantages.

Unlike the optical cavities we used in prior experiments, wherein we could be very confident that resonances were spaced by the free spectral range of the cavity, the microresonators are not so predictable and are very sensitive to perturbations. We must therefore

rely on a known frequency response of the laser to determine the response of the micro-resonators. To do this, we send the laser into our previous optical cavity, which will have resonance spacings of 555 MHz. We can then map the voltage on the piezo in the ECDL to a frequency shift of the beam. We periodically repeat this process to recalibrate the mapping, as the response of the laser may drift over time. We can typically get about 10 to 15 GHz of tuning with our laser, although there are sometimes mode-hops in this range, so we generally scan over a smaller region where the laser is well-behaved. A typical frequency response curve of our laser is shown in Figure 8.7.




Figure 8.7: *Frequency Shift of 1064 nm ECDL vs Applied Voltage to Piezo*

A typical frequency tuning curve for our laser. The response is generally very linear with a slope of about 100 MHz/v . A few mode-hops of unknown magnitude tend to occur in this range, so a somewhat smaller voltage scanning range is used in experiments. Any retuning of the laser requires remeasuring this curve, although the results are generally very similar.

With the frequency response of the laser characterized, and once good coupling with the microresonator is achieved, we ramp the laser and record the transmitted signal. Figure 8.8 shows a typical resonance scan of one of the microspheres. Two main dips are visible, but the microsphere has similar resonances every few GHz throughout the entire optical spectrum. The "depth" of the resonance, which we define as the ratio of the difference between maximum and minimum transmitted power values of the resonance to the off-resonance transmitted value, is strongly affected by the distance of the resonator and fiber, as well as the diameter of the fiber and polarization of light. The right most resonance in the plot shows a depth of about 0.5. This is a fairly average value for us, with careful alignment and high quality tapers letting us reach up to 0.90. The resonances have a width of about 300 MHz, which corresponds to a quality factor of about 10^6 at 1064 nm, where the quality factor is given by

$$Q = \frac{\omega_0}{\Delta\omega}, \quad (8.6)$$

with ω_0 as the laser frequency and $\Delta\omega$ as the resonance width. Because our microspheres are all produced under similar conditions, they all tend to have a Q value of about 10^6 , although some have been about 5 times higher. Over time though, they all tend to degrade, and a Q of 10^6 seems to be the long-term stable value. A literature search suggests this is primarily due to water vapor absorption [43, 44, 45], and so the degradation stops when the resonator water content has reached equilibrium with the ambient air. Storing the resonators under vacuum or periodically baking them at low temperatures can stave off and or somewhat restore degraded resonators, but the highest Q value resonators typically degrade very quickly and so must be measured and used right away. In the future, it would be interesting to test whether there is less degradation or a significantly higher Q value with a pump laser closer to 1555 nm, where there is a minimum in absorption by water.

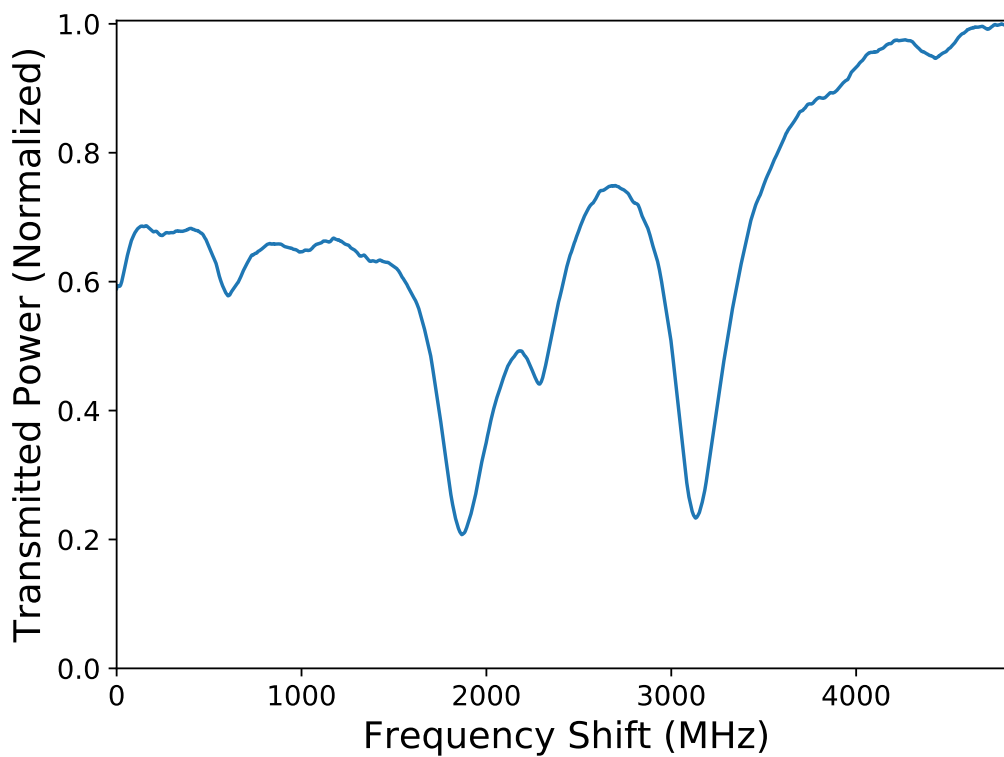


Figure 8.8: *Microsphere Resonances*

The transmitted signal from the tapered fiber as the laser frequency is scanned, which shows the resonances in the microsphere. The frequency response of the laser as the piezo voltage is varied is previously recorded. The resonances have a width of 300 MHz which corresponds to a quality factor of about 10^6 .

Figure 8.9 shows another frequency scan of a microresonator, normalized to a subsequent frequency scan where the microresonator is removed from the system (shown in orange). When in contact, the peak transmitted power drops by about 50%, which is largely due to losses that occur as some of the light that leaves the taper scatters along the surface of the microresonator, rather than coupling into it. Resonators with smoother surfaces that we produced later could show a much smaller power drop, often of about 5%. A weak frequency response is visible in the orange scan, which is partially due to the power output of the ECDL fluctuating as the piezo is scanned and the optical feedback changes, and

partially due to the response of the tapered fiber (as the tapering process tends to introduce strain into the fiber and give a slight frequency response and strong polarization response).

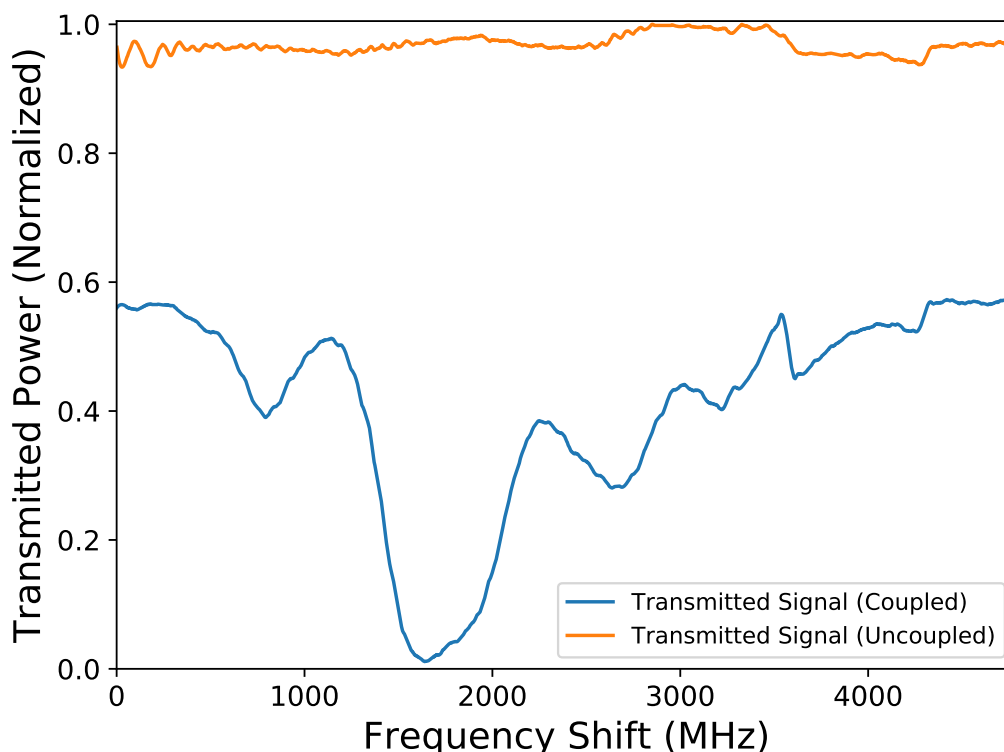


Figure 8.9: *Microsphere Resonances and Fiber Transmitted Signal*

The transmitted signal from the tapered fiber as the laser frequency is scanned when the microsphere is in contact with the fiber (blue) and when removed (orange). There is about a 50% drop in power when the microresonator is in contact, which is primarily due to scattering losses of light that leaves the fiber but does not couple into the resonator. There is a weak frequency response in the orange curve as well, which is due to the laser's response to frequency tuning and the frequency response of the tapered fiber that is caused by strain introduced during the tapering process.

8.3.1.1 High-Power Coupling

These initial experiments and tests were conducted with fairly low pump powers of a few tens of mW. Eventually though, we would need to couple in as much power as we could, potentially many watts, in order to achieve high conversion efficiency. As we

had learned from our cavity experiments, high-power coupling can cause a number of complicating effects, and so we wanted to begin testing this early on. Our main concerns were that at higher-power the transmission through the fiber might drop (due to more light leaking out in the tapered region), and that the coupling efficiency to the resonator would drop. There are also issues of the fibers breaking at high-power, which we discuss in Section 9.2. For these tests, we infer the circulating power in the microsphere by recording the transmitted power through the fiber without the microsphere present, and then again with the microsphere coupled and the laser resonant. We additionally record the laser power at the input to the fiber to determine the overall coupling and transmission efficiency of the fiber without the microsphere coupled. The results are shown in Figure 8.10. There is a drop in transmission and resonance depth at about 50 mW of circulating power, although this quickly levels off and there are no significant further drops. Up to 1 W of circulating power is reached, with approximately 20% transmission through the fiber and 50% resonance depths (so 10 W is incident upon the fiber, 2 W can be transmitted through, and on-resonance 1 W of that light couples and circulates in the microsphere with 1 W continuing on unperturbed).

For a set circulating power, transmission through the fiber is easy to optimize and has negligible uncertainty. For a given alignment of the microsphere and fiber, the coupling efficiency measurements are also very accurate. The microsphere, however, is extremely sensitive to its position relative to the tapered fiber, and different fibers and microsphere combinations can yield very different performance too. Slight changes in alignment can significantly affect the resonance depth, and different resonances may exhibit different depths for a given coupling configuration. For this reason, we perform a reasonable search of the coupling parameter space, and record the highest observed depths that we feel we would be able to frequently reproduce. We roughly estimate the uncertainty as 30% of the difference between the observed resonance depth value and the maximum value of 1.0.

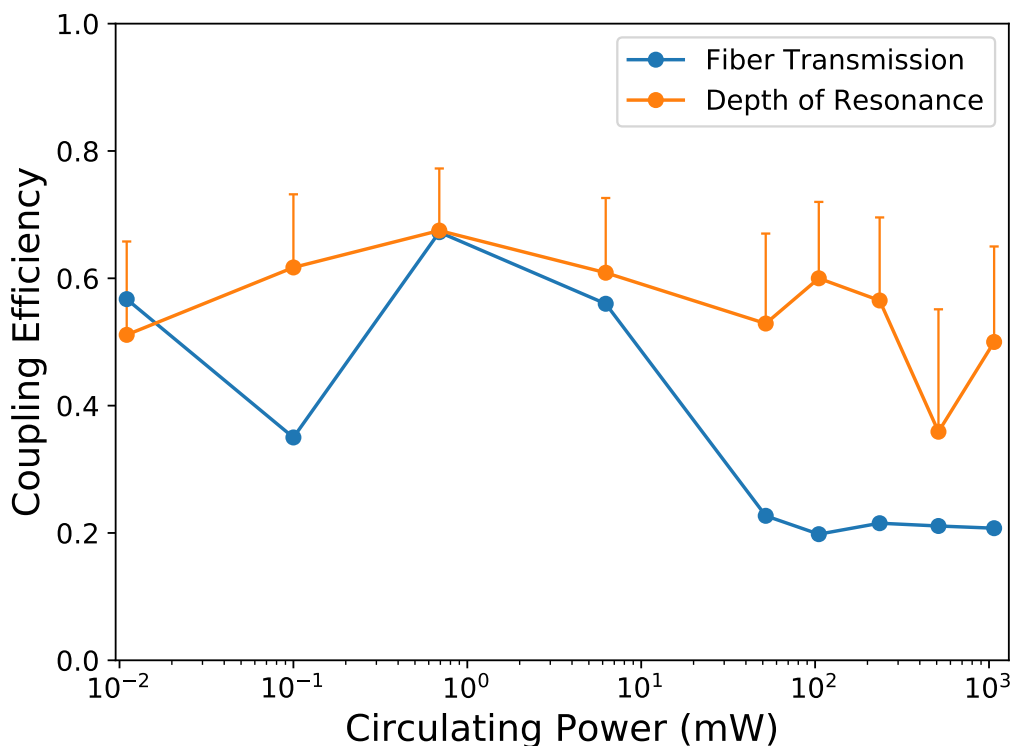


Figure 8.10: *Microsphere Coupling Efficiency*

The fraction of light coupled to and transmitted through a tapered fiber (blue) and the maximum depth of resonance observed with the microsphere coupled (orange) vs the inferred circulating power. There is somewhat of a drop in fiber transmission and resonance depth at about 50 mW of circulating power, but no clear continued drops at higher-power.

These were promising initial results for us, since they suggested that we might not immediately run into issues at high laser power as we had feared. In fact, at 1 W of circulating power, we calculate the radiation intensity inside the microsphere to be 0.1 GW/cm^2 , which is, to the best of our knowledge, the highest ever CW intensities produced inside a glass microresonator. This was about as much as we could do with the microspheres though, as we did not feel silica has a high enough Raman-gain coefficient to make it an efficient modulator. We next switched our focus to more highly Raman-active materials, as discussed in the next chapter.

9 Microdisks

We had begun initial planning for switching to Raman-active crystals while we were still in the early stages of working with microspheres. We identified calcium-fluoride (CaF_2) and magnesium-fluoride (MgF_2) as good potential materials, due to their high Raman-gain, THz level vibrational transitions, ease of availability, and the previous success of other groups in achieving high quality factors with them [46, 47]. The techniques used to create the silica microspheres will not work with these materials however, as they are not available in thin fibers that can then be melted down. Instead, we start with commercially available small cylindrical disks with diameters of 5 to 10 mm and widths of approximately 1 mm. These disks must then be highly polished along the circumference, which was performed by Atul using a micro-lathe, along with some custom mounts designed and machined by Matt.

Besides these differences in manufacturing, most of the experience and knowledge we had gained through working with optical cavities and then silica microspheres could be applied to Raman-active crystals. The basic physics of modulation was very similar to our previous cavity experiments—pump light travels through a material with a vibrational transition, and Stokes photons are generated down-shifted by the transition frequency. These are also resonant in the system and build up to form a Stokes beam, establishing a time-varying index of refraction in the material which separate beams may be sent through and modulated. But the key difference here was that, with a much smaller mode-volume and a higher density material, we hoped higher efficiencies would be possible. The basics of coupling was similar to the microspheres. Disks also support whispering gallery modes, which travel around their circumference, and coupling can also be achieved using tapered

fibers. Because these disks were much larger than the spheres though and generally with higher quality factors, the resonances were much denser in frequency space.

We went through several iterations of refinement in the manufacturing process of these disks before finding a method that worked for us. A photo of one of the early prototypes is shown below, which is an MgF_2 disk bonded to a metal rod. The disk has been polished through many steps of diamond grit paste on a micro-lathe.




Figure 9.1: *An Early Prototype of a Microdisk*

One of the first disks we produced and worked with. We typically bond the disks to metal rods for ease of transport and alignment to a fiber. We were not able to observe any resonances with this disk, likely due to both its quality factor and mode-shape.

9.1 Detecting Resonances

We attempted to couple light into the microdisks using a very similar setup as that shown in subsection 8.3.1. We were not able to observe resonances with this disk or ones of similar designs, however, only a DC drop in transmission when the fiber is brought in contact with the disk, ranging from about 5 to 50%. Although we were not sure of the reasons at the time, we now believe this was because the quality factor might have only been on the order of 10^6 . As we will see, the resonances of these disks tend to be very tightly spaced, often by only 10s of MHz. However, at this quality factor, the linewidths of these resonances would be on the order of 100s of MHz, and so they would exhibit significant amounts of blending and not be resolvable. A similar effect might be caused by the thickness of the disk. While we initially thought that light from the fiber would couple to low-order modes of the disk with narrow mode-shapes, we now believe that once resonating in the disk, light rapidly couples to higher-order modes and fills the entire width of the cylinder. This results in the accessible modes of the disk once again blurring together in frequency space and being unresolvable. This is also undesirable because the intensity inside the disk will be much lower than if the light stayed well confined.

Our later microdisks were polished with finer grades of grit in hopes of achieving a higher quality factor. Additionally, coarser grits are used to grind away excess material to create a wedge shape at the edge of the disk. The sharp angle of the wedge helps to confine light to a smaller region, increasing intensity and reducing coupling to higher-order modes. Photos of one of these wedge shaped disks are shown below.




Figure 9.2: *A Wedge-Cut Microdisk*

Cutting and polishing the circumference of the disks into a wedge shape greatly confines coupled light to a smaller volume and fewer modes. The disk is about 1 mm thick, and coupling is possible along a region of about 100 microns.

These changes allowed us to observe resonances in these disks for the first time. These frequency scans are recorded in the same way as described in subsection 8.3.1. One of these scans is shown in Figure 9.3. Note that the horizontal scale is similar to that of the microsphere scan showed in Figure 8.8, and yet many tens of resonances are visible: The resonances of the disk are much denser and narrower than those of the spheres.




Figure 9.3: *A Microdisk Resonance Scan*

The transmitted signal from the tapered fiber as the laser frequency is scanned, which shows the resonances of the microdisk. Compared to the microspheres, these resonances are much narrower and more densely spaced.

A zoomed view of a different frequency scan is shown in Figure 9.4. The deepest resonance here shows about 70% depth and a width of about 2 MHz, corresponding to a quality factor of about 10^8 .




Figure 9.4: *A Zoomed View of Microdisk Resonances Scan*

Looking at the resonances of the microdisk more closely shows maximum depths of about 70% and widths of about 2 MHz, corresponding to a quality factor of about 10^8 .

9.2 Searching for Raman Lasing

Our initial calculations suggested that even with substantial coupling from our pump laser, we might only be near the threshold for Raman lasing. This was not necessarily a major problem, since, as explored in detail in Josh Karpel's thesis, peak modulation efficiency should occur at a pump power only slightly higher than that needed for threshold—one feature in which the microresonators do differ substantially from the optical cavities, where we would expect a continuous rise in efficiency with increasing pump power [38]. To look for Raman lasing, we take the transmitted light from the system (the same setup as shown in subsection 8.3.1) and send it to our homemade, high sensitivity spectrometer that was

shown in Figure 7.6.

Unfortunately, as we increased the amount of light we coupled to the microdisk in hopes of reaching threshold, we encountered consistent problems with the fibers breaking. Our initial thought was that dust in the air might heat up from the evanescent field from the fiber and may cause the waist of the fiber to melt and break. Some sources suggest this is likely as well [41, 48]. However, setting up a vacuum chamber in which to run the experiment had no noticeable effect on fiber durability or their ability to withstand high-power, so we do not believe that dust is the issue in our case. We also tried sending the pump beam through an AOM so that we could use 1 μ s long pulses with a low duty-cycle to allow time for the fiber to cool. We found though that the fibers still tended to break when the peak laser power was at about the same level as in the CW case. It may be that because the fibers are so thin they heat up too much even during this short time. We did not wish to go to shorter pulse lengths though, as this would potentially interfere with the timescale of the Raman processes we were interested in, and also would go against everything we stood for. We are a CW group through and through.

As not many research groups send such high power through tapered fibers, this is a somewhat unexplored area. Our current theory is that the issue is with the fibers themselves, and that even with a clean fiber surface and surrounding environment, impurities in the fiber will scatter light. At high laser powers, this scatter can generate enough heat to soften the waist and break it. Fiber breakage was a major hurdle in our work with the microdisks: Producing the tapered fibers was a time and labor intensive process, but we needed to keep testing them at high laser power to look for Raman lasing, or at least to try to understand why we were not observing it. There was a fair amount of variance in the manufacturing process, however, and some fibers occasionally seemed more robust and could withstand higher amounts of power.

Unfortunately, it seemed that the microdisks were susceptible to environmental degra-

dation like the microspheres. By the time we had a few fibers that we thought could handle the high laser power we needed, the microdisks often no longer had the high quality factor we needed to produce high intensities of light. Figure 9.5 shows a frequency scan of a microdisk where the quality factor has decreased from 10^8 to about 10^7 over about two months. A few weeks later, the quality factor had further decreased to about 6.5×10^6 , where it seemed to stabilize, as shown in Figure 9.6. We believe the reason for this is due to water vapor absorption, as can be the case in silica microspheres. Repolishing the crystals at the last few grit levels can often somewhat improve the quality factor, as shown in Figure 9.7. This suggests that water vapor does not penetrate the crystal very deeply, or at least that it would take a long time to do so. However, repolishing the crystals does not generally completely restore them. We experimented with baking the crystals under vacuum, generally for about 24 hours at 100°C , but this did not make definitive improvements. Note that these frequency scans are taken using different fibers and with different coupling alignments between the fiber and disk than the previous ones, and so the depth and density of resonances are not necessarily indicative of the quality of the crystal. The width of the resonances though tends to be fairly consistent as the coupling is changed and indicates a drop in quality factor.




Figure 9.5: *A Degraded Microdisk*

After about two months in the lab, resonances have substantially broadened compared to their initial width. The quality factor has decreased to about 10^7 , presumably due to water vapor absorption.

Gradually, we managed to hone in on the requirements for robust fibers. Measuring the diameters under a microscope and recording the performance of the fiber let us determine that the sweet spot was between 2 and 2.5 microns. While we did not have precise control over the diameter of the fibers, measuring each fiber let Atul somewhat adjust his techniques as we went, allowing us to frequently produce fibers in this diameter range. Finally, we managed to assemble the elusive combination of a high quality factor microdisk, a fiber of the right diameter, and a reasonably well behaved pump laser. With several watts of input pump power, we observed light at our spectrometer at approximately 1100 nm, corresponding to a 10 THz down-shift from our pump beam, as predicted for the Stokes




Figure 9.6: *A Substantially Degraded Microdisk*

The resonances have further broadened, giving a quality factor of about 5×10^6 . The degradation seems to slow down at about this level.

beam. We do not currently have a locking system set up for the microspheres; However, as the narrowest linewidth resonances we have observed are only 2 MHz and the pump laser is about 0.5 MHz, this is not particularly necessary. The laser can be kept reasonably well on-resonance just through small adjustments to the laser piezo. To better observe the dynamics of the system though, and since different resonances can have very different performance, we continue to ramp the pump laser when looking for Raman lasing. Figure 9.8 shows the transmitted pump beam and transmitted (generated) Stokes beam simultaneously recorded. The pump beam is recorded on the "transmitted signal" photodiode shown in Figure 8.6. The rest of the beam is sent onward to our spectrum analyzer, where the Stokes beam is separated using the diffraction grating before being focused on a photodiode.




Figure 9.7: *The Repolished Crystal*

Repolishing the crystal at the last few finest grit levels often somewhat restores the quality factor. This one improved from approximately 5×10^6 to about 5×10^7 .

To view more detailed time-dynamics, we also run scans with a reduced amplitude of the ramp to the laser (a lower amplitude voltage is applied to the piezo, resulting in a smaller frequency shift in one cycle). We view the signals in the time-domain, as shown in Figure 9.9.

Within our time-resolution, the peak power of the Stokes beam occurs within 0 to 4 microseconds after the lowest dip in the pump beam, which is in agreement with the predicted time of a few hundred nanoseconds in simulations performed by Josh Karpel. The frequency shift of the laser during this scan is $15 \text{ kHz}/\mu\text{s}$, and so is negligible compared to the laser and resonance widths on these timescales.

At the high coupled-powers needed to induce Raman lasing in the microdisks, we




Figure 9.8: *Raman Lasing from the Microdisk*

The pump and Stokes signals are simultaneously recorded as the pump laser frequency is scanned. The Stokes light rapidly builds up when the pump laser is on-resonance (at a minimum in transmission), before quickly decaying as the pump laser moves off-resonance. The Stokes beam is separately normalized to a lower value than the pump beam so that they may be more easily viewed together.

noticed that the disks started to glow a striking blue-green, and would sometimes suddenly change shades, while other times they would smoothly shift along the color spectrum. This was unexpected and somewhat alarming the first time it occurred, but it soon became a fan favorite non-linear optical effect among the group, with labmates lining up out the door to view it and to have their photo taken with it. The glowing crystal is shown in Figure 9.10, which unfortunately does not entirely do it justice.

A full wavelength scan of the transmitted light using a commercial spectrum analyzer, shown in Figure 9.11, helped to somewhat explain this mysterious glowing. The spectrum




Figure 9.9: *Raman Lasing from the Microdisk (Narrow Scan)*

Scanning over a narrow frequency range lets us clearly see that the generated Stokes light occurs very close to the pump beam in time and matches our predictions.

analyzer is very dated, and seems to be somewhat miscalibrated in its power measurements, but gives reasonably accurate wavelength readings. The pump beam is visible near 1060 nm. There are a variety of spectral features in this region, but they are not visible at the available resolution over this wavelength range. Some fluorescence is visible around 1200, 1400, and 1800 nm. The peaks from around 650 to 800 nm might be due to 2 photon processes of the 1800 and 1400 nm fluorescence. These types of processes, which are typically observed, can occur here since almost all light produced in the microdisk will resonate in it. The mysterious blue-green light could be seen to change color over a few seconds, often smoothly transitioning between blue and green, but sometimes quickly jumping. The peak on the spectrum analyzer would shift between about 480 and 550 nm.




Figure 9.10: *A Glowing Microdisk*

At high coupled-powers, the microdisk would begin to glow, often shifting between blue and green and flickering as the laser moved on and off-resonance.

This seems likely to be a 2-photon process of combinations of pump photons and Stokes photons. There is also significant self-phase modulation on the pump (not visible in this full scan), which could contribute photons to fill what spectral gaps might otherwise be left in the blue-green region from this process. Finally, there is a large peak in the ultraviolet region. There is likely a calibration issue with the spectrum analyzer, as we were not able to isolate nearly as much ultraviolet light as this plot would suggest, but frequency tripling of the pump beam matches well with the observed spectrum to explain this peak.




Figure 9.11: *Broad Spectral Scan of the Microdisk Transmitted Light*

The pump beam can be seen near 1060 nm, but the Raman lasing cannot be seen at this resolution. Other notable features include fluorescence at approximately 1200, 1400, and 1800 nm, as well as a variety of 2-photon processes between 480 and 550 nm, and between 650 and 800 nm. Frequency tripling of the pump beam likely explains the peak in the ultraviolet region.

Performing a higher resolution scan over a smaller range let us view the Raman lasing in more detail, which is shown below. The pump beam is seen near 1060 nm. There is also significant broadening from self-phase modulation. While this typically occurs with pulsed lasers, the relatively high CW power of the beam and the extremely small volumes of the tapered fiber and microdisk allow the required intensities to be reached. A Stokes beam is generated from Raman scattering off the first vibrational transition in the CaF_2 disk, at 10 THz as expected. This Stokes beam is at sufficient intensity that it interacts with the same transition again to produce a second-order Stokes beam another 10 THz away. Due to the broad linewidth of the transition (approximately 1 THz) and the dense mode-spacing of the microdisk, these beams all resonate in the crystal, allowing them to reach significant powers.

We achieved a maximum of 300 mW of Stokes light and 10 mW of second-order Stokes light, with 1.5 W of pump light coupled to the tapered fiber and 1 W of light coupled and circulating in the microdisk (meaning 67% resonance depths). This corresponds to 30% conversion efficiency of the pump into Stokes by our measure of efficiency as the output Stokes power divided by the circulating pump power. To the best of our knowledge, 300 mW of Stokes light is about two orders of magnitude higher than had ever been previously observed in these types of systems.

Unfortunately, shortly after achieving these high Stokes powers, our tapered fiber broke and in the few weeks it took to produce more, it seemed that the microdisk had somewhat degraded. We were not able to reproduce the same level of conversion efficiencies, instead reaching a maximum of about 3%. Still, we were able to explore the effect of circulating pump power on conversion efficiency. This plot is shown in Figure 9.13. Threshold is reached at about 300 mW of circulating power and mostly rises after that. We do not take the drop in efficiency at around 800 mW of power to be particularly significant. At these high circulating powers, the microdisk can rapidly heat up and expand when the laser is

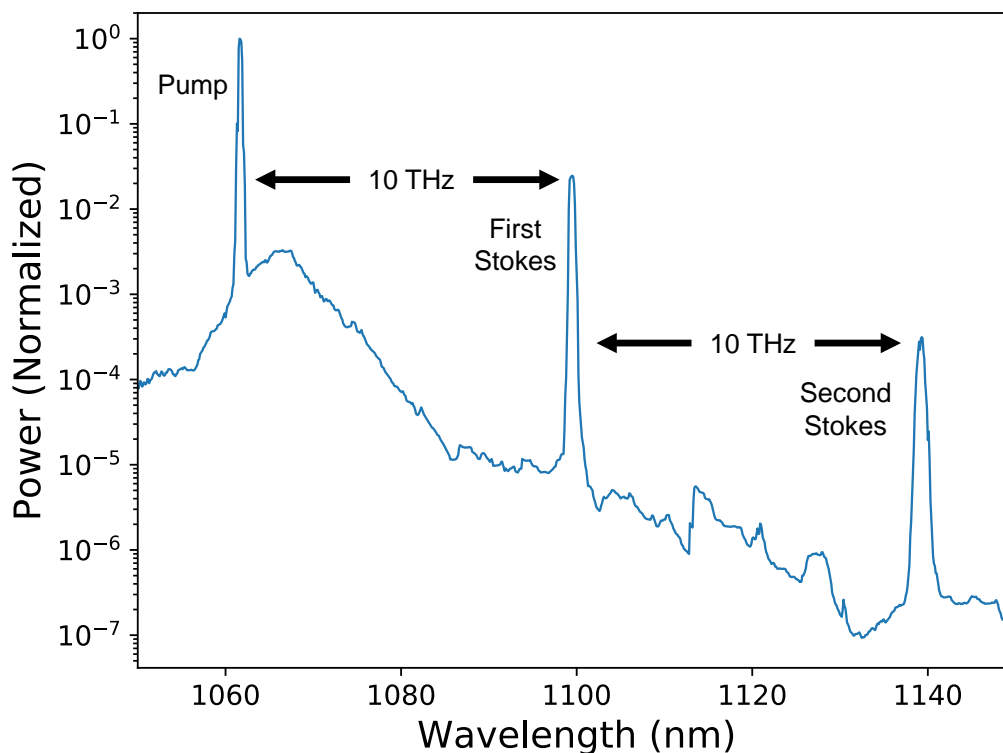


Figure 9.12: *High-Resolution Spectral Scan of the Microdisk Transmitted Light*

A high-resolution spectral scan of the transmitted light from the microdisk near the pump beam wavelength. The pump beam can be seen near 1060 nm, with broadening from self-phase modulation evident as well. A Stokes beam is generated 10 THz away, which in turn generates a second-order Stokes beam another 10 THz away.

on-resonance, causing the resonances to shift. This can make measuring the conversion efficiency challenging, and likely reduces the efficiency somewhat. Adding a heat sink and thermoelectric cooler to the microdisk in the future would likely reduce these issues. For the same reasons as discussed in regard to previous figures, the error bars are one-sided and estimated at 30%.




Figure 9.13: *Pump to Stokes Conversion Efficiency*

The conversion efficiency of the circulating pump power to transmitted Stokes power as the circulating pump power is increased. Threshold is first crossed at about 300 mW of pump power and the efficiency mostly increases from there.

9.3 Future Directions

Our collaboration with van der Weide group has become quite proficient in producing high quality microresonators and tapered fibers, as well as in the techniques to couple light to them and induce and observe Raman lasing. We should continue to refine and optimize these methods. The immediate next steps for the experiment are clear: After having observed modulation on the pump beam, we wish to now demonstrate modulation on an independent mixing beam, as in our previous cavity experiments. Doing so will require two separate tapered fibers to be coupled to a microdisk, as the optimal coupling

parameters are very wavelength dependent. We have already made progress on developing a setup to enable this.

The basic principles for this mixing beam experiment are similar to the implementation in the cavity—just as before, the mixing beam will be able to interact with the coherence produced by the strong pump beam. However, while in the cavity setup the mixing beam would not be resonant and would be modulated in a single-pass, there is no direct "single-pass" analogue here. We believe that the mixing beam will need to couple into the microdisk for efficient modulation. However, due to the geometric nature of the microdisk resonances, there will be resonances available every few tens of MHz in the optical region of the spectrum, and so any mixing beam will be near a resonance. Alternatively, tuning of the microdisk resonances via temperature control would also allow for applications where slight frequency tuning of the mixing beam is not desirable. Depending on the performance of the modulator, it might then be useful to improve the mechanical stability of the setup and work to make the system function closer to a "turn-key operation" device.

Finally, for future students who work on this project, here is a short list of ideas that may be useful one day for improvements or changes to make.

- As manufacturing techniques for dielectric mirrors continue to advance, it may become possible to source mirrors that are highly reflective at the pump, Stokes, and anti-Stokes wavelengths in deuterium, with high transparency throughout the rest of the optical spectrum. In this case, revisiting a cavity setup might be beneficial, as the coherence and modulation efficiency could be substantially increased.
- Our modulators work by producing a time-varying index of refraction in a small localized region of space, which creates frequency sidebands on any light passing through this region. It may also be possible to create sidebands by instead using a spatially-varying index of refraction that is constant in time. For light to experience

THz level modulation, this would likely require a material with very consistent structure at the 10s of microns level over a length of at least a few centimeters. This sounds difficult, but does not seem impossible. A single-mode fiber and lithography techniques would likely be useful.

- Performance might be improved in the microdisk setup by tuning the resonances of the disk rather than tuning the laser. Temperature control could provide slow, broad tuning and would improve stability as well. For more rapid scanning to view resonance shapes, a piezo-controlled compression mount could be built to squeeze the microdisk at a rate of a few tens to hundreds of Hz. By Equation 8.2, it can be seen that this would allow for significant tuning through changing the radius of the disk and its eccentricity. Other groups have successfully done this with silica microspheres and achieved tuning of over 100 GHz [49]. Several GHz of tuning would be quite adequate for our application.
- If we could make a very round microdisk so that the splitting due to eccentricity is less than the linewidth of a resonance, we would have something analogous to a confocal optical cavity. While resonances would now be far apart in frequency (spaced by about the free spectral range), they would likely be very deep and alignment insensitive, which could greatly improve stability. For our 7 mm diameter disks with $Q \approx 10^8$, we would need an eccentricity of $\lesssim 0.015$, meaning that the major and minor axes differ by about 1 part in 1000 or a few hundred nanometers.
- A simple locking system between the pump laser and microdisk would be useful. Unless we reach significantly higher quality factors, slow feedback via the laser piezo would be sufficient. However, feedback via a compression tuner on the microdisk might work better. At high laser powers, having a TEC on the microdisk for temperature stabilization would be very helpful. A passive heat-sink might be sufficient

though, and an easy starting point would be to replace the steel rod bonded to the microdisk with an aluminum or copper one. It might also be possible to have the entire microdisk and tapered fiber submerged in a liquid for cooling.

- Putting the microdisk setup under vacuum to reduce dust did not help with tapered fibers breaking at high laser power. However, it may be the case that dust is still the issue, and that the fibers are simply already too dirty by the time we get them. I have had some preliminary success with very gently cleaning the tapered region of the fibers with isopropanol, and this improved transmission through the fiber and, I believe, improved high-power durability. This should be explored more. An ultrasonic cleaner might work well.
- Depending on the exact mechanism for the degradation of the microdisks' quality factor over time, it might help to switch to using our 1555 nm laser as the pump. If water vapor absorption is a contributing factor, we might see a higher Q at this wavelength since it is absorbed much less by water than 1064 nm is.
- One or two small USB microscopes would be easy to install to view the microdisk and tapered fiber interface while making alignment adjustments. This will be especially useful as we move toward wedge-cut microdisks that come to narrower points. Switching to a micrometer that has a few microns of piezo travel for fine control of the microdisk would also be very helpful.

The molecular modulation project has come a long way since Deniz' original proposal, spanning three generations of graduate students over more than a decade of effort from the group. It has been a privilege to contribute to this project and to be involved in some important milestones—the first modulation of a mixing beam at 90 THz, reaching new modulation efficiencies two orders of magnitude higher than our previous records,

modulating a broadband source, and, most recently, observing Raman lasing with our new microresonator setup. In what little time I have left in the lab, I hope to help teach and train our group's next generation of students, so that they might one day modulate a mixing beam with our microresonator. Even if I won't be around to see it myself, when it happens, somehow, I think I'll know. At the very least, I'd like to instill in them an admiration and respect for high-frequency optical modulation. I hope my thesis has been sufficient testament to the same sentiment that I hold.

Bibliography

- [1] S. Rausch, T. Binhammer, and L. Gloster. Shortest CEP stabilised Ti:sapphire pulse. Technical report, Laser Quantum, 2013.
- [2] T. Gaumnitz, A. Jain, Y. Pertot, M. Huppert, I. Jordan, F. Ardana-Lamas, and H. J. Wörner. Streaking of 43-attosecond soft-x-ray pulses generated by a passively CEP-stable mid-infrared driver. *Optics Express*, 25:27506, 10 2017.
- [3] B. W. J. McNeil and N. R. Thompson. X-ray free-electron lasers. *Nature Photonics*, 4(12):814–821, 2010.
- [4] G. Krauss, S. Lohss, T. Hanke, A. Sell, S. Eggert, R. Huber, and A. Leitenstorfer. Synthesis of a single cycle of light with compact erbium-doped fibre technology. *Nature Photonics*, 4:33–36, 2010.
- [5] M. Chaciński, U. Westergren, B. Stoltz, L. Thylén, R. Schatz, and S. Hammerfeldt. Monolithically integrated 100 GHz DFB-TWEAM. *Journal of Lightwave Technology*, 27:3410–3415, 2009.
- [6] A. Brimont, D. J. Thomson, P. Sanchis, J. Herrera, F. Y. Gardes, J. M. Fedeli, G. T. Reed, and J. Martí. High speed silicon electro-optical modulators enhanced via slow light propagation. *Optics Express*, 19:20876–20885, 2011.
- [7] S. E. Harris and A. V. Sokolov. Broadband spectral generation with refractive index control. *Physical Review A*, 55:R4019–R4022, 1997.

- [8] J. J. Weber. *Continuous-Wave Light Modulation Using Stimulated Raman Scattering*. PhD thesis, University of Wisconsin-Madison, 2014.
- [9] J. T. Green. *Optical Cavity-Based Continuous-Wave Raman Generation*. PhD thesis, University of Wisconsin-Madison, 2010.
- [10] A. Franzen. ComponentLibrary (a vector graphics library for illustrations of optics experiments). Online, 2009. <http://www.gwoptics.org/ComponentLibrary/>.
- [11] D. A. Russell and W. B. Roh. High-resolution CARS measurement of Raman linewidths of deuterium. *Journal of Molecular Spectroscopy*, 124:240–242, 1987.
- [12] H. G. M. Edwards, D. A. Long, and G. Sherwood. Handbook of Raman spectroscopy. *Journal of Raman Spectroscopy*, 22:607, 2001.
- [13] R. W. P. Drever, J. L. Hall, et al. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B*, 31:97–105, 1983.
- [14] R. W. Fox, C. W. Oates, and L. W. Hollberg. Stabilizing diode lasers to high-finesse cavities. In *Experimental Methods in Physics Sciences*, volume 40, chapter 1, pages 1–46. Elsevier Science, 2003.
- [15] J. T. Green, J. J. Weber, and D. D. Yavuz. Continuous-wave light modulation at molecular frequencies. *Physical Review A*, 82:011805–011809, 2010.
- [16] J. J. Weber, J. T. Green, and D. D. Yavuz. 17 THz continuous-wave optical modulator. *Physical Review A*, 85, 2012.
- [17] S. Yoshikawa and T. Imasaka. A new approach for the generation of ultrashort optical pulses. *Optics Communications*, 96:94–98, 1993.

- [18] A. E. Kaplan. Subfemtosecond pulses in mode-locked 2π solitons of the cascade stimulated Raman scattering. *Physical Review Letters*, 73:1243–1246, 1994.
- [19] H. Kawano, Y. Hirakawa, and T. Imasaka. Generation of high-order rotational lines in hydrogen by four-wave Raman mixing in the femtosecond regime. *IEEE Journal of Quantum Electronics*, 34(2):260–268, 1998.
- [20] A. V. Sokolov, D. R. Walker, D. D. Yavuz, G. Y. Yin, and S. E. Harris. Raman generation by phased and antiphased molecular states. *Physical Review Letters*, 85:562–565, 2000.
- [21] D. D. Yavuz. *A Raman approach for generating ultrashort pulses*. PhD thesis, Stanford University, 2003.
- [22] M. Y. Shverdin, D. R. Walker, D. D. Yavuz, G. Y. Yin, and S. E. Harris. Generation of a single-cycle optical pulse. *Physical Review Letters*, 94:033904, 2005.
- [23] Z. M. Hsieh, C. J. Lai, H. S. Chan, S. Y. Wu, C. K. Lee, W. J. Chen, C. L. Pan, F. G. Yee, and A. H. Kung. Controlling the carrier-envelope phase of Raman-generated periodic waveforms. *Physical Review Letters*, 102:213902, 2009.
- [24] O. Shitamichi and T. Imasaka. High-order Raman sidebands generated from the near-infrared to ultraviolet region by four-wave Raman mixing of hydrogen using an ultrashort two-color pump beam. *Optics Express*, 20:27959–27965, 2012.
- [25] J. K. Brasseur, K. S. Repasky, and J. L. Carlsten. Continuous-wave Raman laser in H_2 . *Optics Letters*, 23:367–369, 1998.
- [26] J. K. Brasseur, P. A. Roos, L. S. Meng, and J. L. Carlsten. Co-linear anti-Stokes generation from a CW Raman laser. In H. Gibbs T. Gallagher, H. Van Driel and D. Wineland, editors, *OSA Technical Digest. Quantum Electronics and Laser Science Conference*, Optical Society of America, 2000.

- [27] L. S. Meng, P. A. Roos, and J. L. Carlsten. Continuous-wave rotational Raman laser in H_2 . *Optics Letters*, 27:1226–1228, 2002.
- [28] J. K. Brasseur, P. A. Roos, K. S. Repasky, and J. L. Carlsten. Characterization of a continuous-wave Raman laser in H_2 . *Journal of the Optical Society of America B*, 16:1305–1312, 1999.
- [29] K. Shinzen, Y. Hirakawa, and T. Imasaka. Generation of highly repetitive optical pulses based on intracavity four-wave Raman mixing. *Physical Review Letters*, 87:223901, 2001.
- [30] K. Ihara, C. E. Zaitsev, S. Ichi, S. Kamitomo, K. Shinzen, Y. Hirakawa, and T. Imasaka. Molecular-optic modulator. *Applied Physics Letters*, 88, 2006.
- [31] D. D. Yavuz. High-frequency modulation of continuous-wave laser beams by maximally coherent molecules. *Physical Review A*, 76:011805–011809, 2007.
- [32] J. T. Green, D. E. Sikes, and D. D. Yavuz. Continuous-wave high-power rotational Raman generation in molecular deuterium. *Optics Letters*, 34:2563–2565, 2009.
- [33] J. T. Green, J. J. Weber, and D. D. Yavuz. Continuous-wave, multiple-order rotational Raman generation in molecular deuterium. *Optics Letters*, 36:897–899, 2011.
- [34] J. J. Weber and D. D. Yavuz. Broadband spectrum generation using continuous-wave Raman scattering. *Optics Letters*, 38:2449–2451, 2013.
- [35] D. C. Gold and N. R. Brewer. Laboratory safety guide. <https://youtu.be/SBLDrGqGs5w>, 2016.
- [36] D. C. Gold, J. J. Weber, and D. D. Yavuz. Continuous-wave molecular modulation using a high-finesse cavity. *Applied Sciences*, 4:498–514, 2014.

- [37] D. C. Gold, J. T. Karpel, D. D. Yavuz, and E. A. Mueller. Continuous-wave modulation of a femtosecond oscillator using coherent molecules. *Optics Letters*, 43:1003–1006, 2018.
- [38] J. T. Karpel. *Numerical Simulation of Intense Ultrafast Quantum Phenomena*. PhD thesis, University of Wisconsin-Madison, 2019.
- [39] D. H. Lee, K. H. Kwon, J. W. Song, and J. Park. Variable optical fiber attenuator using bending-sensitive fiber. *Journal of the Optical Society of Korea*, 8:83–89, 2004.
- [40] J. P. Rezac. *Properties and Applications of Whispering-Gallery Mode Resonances in Fused Silica Microspheres*. PhD thesis, North Dakota State University, 1995.
- [41] G. A. Cohoon. *Fabrication, Characterization, and Applications of Microresonators and Resonant Structures*. PhD thesis, The University of Arizona, 2016.
- [42] Y. Lang. *Fabrication and Characterization of Microlasers by the Sol-Gel Method*. PhD thesis, California Institute of Technology, 2005.
- [43] M. L. Gorodetsky, A. A. Savchenkov, and V. S. Ilchenko. Ultimate Q of optical microsphere resonators. *Optics Letters*, 21:453–455, 1996.
- [44] D. W. Vernooy, V. S. Ilchenko, H. Mabuchi, E. W. Streed, and H. J. Kimble. High-Q measurements of fused-silica microspheres in the near infrared. *Optics Letters*, 23:247–249, 1998.
- [45] E. Hosseinian, P.-O. Theillet, and O. N. Pierron. Temperature and humidity effects on the quality factor of a silicon lateral rotary micro-resonator in atmospheric air. *Sensors and Actuators A: Physical*, 189:380–389, 2013.
- [46] I. Grudinin, A. Matsko, A. Savchenkov, D. Strekalov, V. Ilchenko, and L. Maleki. Ultra high Q crystalline microcavities. *Optics Communications*, 265:33–38, 09 2006.

- [47] F. Sedlmeir, R. Zeltner, G. Leuchs, and H. G. L. Schwefel. High-Q MgF₂ whispering gallery mode resonators for refractometric sensing in aqueous environment. *Optics Express*, 22:30934–30942, 2014.
- [48] J. Ward and O. Benson. WGM microresonators: sensing, lasing and fundamental optics with microspheres. *Laser Photonics Reviews*, 5:553–570, 2011.
- [49] V.S. Ilchenko, P.S. Volikov, V.L. Velichansky, F. Treussart, V. Lefevre-Seguin, J.-M. Raimond, and S. Haroche. Strain-tunable high-Q optical microsphere resonator. *Optics Communications*, 3:86–90, 1998.