

**Mode-Locked Fiber Lasers: Development and  
Application**

by

**Darren D. Hudson**

B.S., Centre College, 2003

A thesis submitted to the  
Faculty of the Graduate School of the  
University of Colorado in partial fulfillment  
of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2009

This thesis entitled:  
Mode-Locked Fiber Lasers: Development and Application  
written by Darren D. Hudson  
has been approved for the Department of Physics

---

Steven T. Cundiff

---

Jun Ye

Date \_\_\_\_\_

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Hudson, Darren D. (Ph.D., Physics)

Mode-Locked Fiber Lasers: Development and Application

Thesis directed by Professor Adjoint Steven T. Cundiff

The field of mode-locked fiber lasers has grown tremendously over the last 10 years. In the last few years, in particular, this class of laser has moved from just offering a low cost, rugged and compact source of ultrashort pulses to offering state of the art ultrashort pulses. Rapid progress in fiber development has led to a variety of specialty fibers: highly nonlinear fiber for various wavelength ranges, high dopant gain fiber, double-clad high gain fiber, and dispersion compensating fiber (to name just a few). These fundamental developments resulted in higher performance fiber laser systems. For instance, the high dopant gain fiber and the nonlinear fiber resulted in the ability to make fiber frequency combs at any wavelength. The double-clad fiber has allowed researchers to push the average power of Yb fiber lasers to  $>10$  W; a level which is already above that offered from the popular Ti:sapphire system.

In this thesis, Erbium based mode-locked fiber lasers are examined from a development and application point of view. The first two chapters review some of the basic concepts that are used throughout this thesis. Chapter 3 covers a crucial advance necessary for fiber lasers to be used in precision experiments: frequency control and frequency dissemination over fiber links. The first point is accomplished with a fast intra-cavity actuator, while the second point is addressed using a stabilized fiber link. Chapter 4 then reviews two atomic physics based experiments that used stabilized fiber lasers.

The next two chapters describe and present characterization for a new method of achieving a mode-locked fiber laser based on a device known as a

waveguide array. We believe this method could yield one of the most robust and compact mode-locked fiber lasers ever created. The experiment detailed in Chapter 5 involved measuring the pulse shaping of these waveguide array devices via autocorrelation. This measurement was the first demonstration of pulse shortening in waveguide arrays. Further characterization in Chapter 6 measured the effects of multi-photon absorption on the discrete spatial soliton that is formed at high peak power in the waveguide array. This experiment showed that multi-photon absorption in the device effectively clamps the spatial soliton power distribution, with increases beyond a certain peak intensity causing virtually no change in the output distribution. The last experiment in Chapter 6 details a measurement of the full electric field shaping of the waveguide array using Frequency-Resolved Optical-Gating. Analysis of the data shows that the waveguide array has a spectral phase attraction point. Thus, any value for the input spectral phase is transformed, upon traveling through the waveguide array, into one output spectral phase. The last chapter provides a *big picture* overview of the topics covered in this thesis, and takes a look at the future directions in which this work is headed.

## Dedication

to my best friend and wife, Adrienne.

## Acknowledgements

My time here at JILA has been a great learning experience scientifically and personally. My advisors, Steven Cundiff and Jun Ye, have allowed me to learn a great deal of physics while working in two wonderful labs (and participating in two wonderful ski retreats!). From Steve I have learned to approach new concepts with a determination to understand every detail. I feel this attitude is crucial for success in any technical field. I would like to thank Steve for his willingness to not only guide me in the lab, but to send me to conferences. Meeting and talking to people in the same field has an energizing effect on research, and I am grateful for having the opportunity to experience that feeling many times over the last few years. From Jun I have learned that energetic optimism about your experiment can lead to not only great fun in the lab, but great success as well. I hope that when I leave JILA I can replicate his wonderful attitude about science with my colleagues.

I owe a great debt of gratitude to the JILA shop staff. Everyone that knows of JILA knows about the amazing support staff. They are a significant part of the progress of science in JILA. In particular, Tom Foote, David Alchenberger, Terry Brown, Paul Beckingham, Hans Greene, James Fung-A-Fat, Mike Whitmore, and Carl Sauer have all been extremely willing to educate me. For this, I am very grateful.

By working in two labs for a large part of my graduate career I was able to make a lot of friends. I would like to thank several of these friends who were

directly responsible for my early hands-on training: David Jones, Kevin Holman, Pete Roos and Seth Foreman. To them I am grateful for their willingness to discuss physics and answer my questions. I would also like to thank Seth Foreman for his great friendship in my latter years as a graduate student. Our ability to discuss physics and make each other laugh was quite fun and sorely missed when he graduated. Other members of JILA that I am extremely grateful to have had the chance to work with are: Mike Thorpe, Matthew Stowe, Jason Jones, Thomas Schibli, Eric Hudson, Brian Sawyer, Marty Boyd, Andrew Ludlow, Dylan Yost, Marcio Miranda, Jared Wahlstrand, Andy Funk, Mingming Feng, Alan Bristow, Ryan Smith, and Qing Chao. While this list is only partial, it is still quite long and does not permit me to thank each individual for their unique contribution to my training. However, I would be remiss (if not beaten up) if I didn't single out my older brother, Eric, who has proven to be one of the wisest voices throughout my life. To say I wouldn't be where I am today without him is quite an understatement.

Another group of friends I must thank are friends from my graduate year with whom I weathered many homeworks, exams, and jam sessions. This list goes something like this: David James, Matt White, Paul Parazzoli, Ted Reber, Jim Peoble, David Hume and Mike Thorpe. These people I consider great physicists, but greater friends. I knew graduate school would be a challenging experience, but I never expected it to be so fun until I met them. I believe to be successful, one has to have balance in their life and these friends provided that for me. A special thank you is in order to David James who proved to be an awesome roommate and fellow electronics tinkerer. Though our device success rate hovered around 50%, our fun rate was through the roof. I would also like to thank Mike Thorpe, who is as fun to talk to about spectroscopy as he is to talk to about golf. Although, he seems to be more of an expert in the former rather than the latter!

I would like to thank my mother and father who have always encouraged me with faith and love. If I have accomplished anything, the real credit lies with them. They are, without a doubt, the most selfless people I have ever met and for that I can't say thank you enough.

Finally, I want to thank my beautiful wife Adrienne. Her ability to make me smile and have fun no matter what mood I'm in is amazing, and her patience and willingness to put off big adventures for the sake of bringing me along is greatly appreciated.

## Contents

### Chapter

<b>1</b>	Introduction	1
1.1	Mode Locking Fundamentals . . . . .	3
1.2	A Brief History of Mode Locking . . . . .	7
1.3	Effective Saturable Absorbers: Kerr Lens Mode locking and Additive Pulse Mode Locking . . . . .	9
1.4	Erbium-doped Mode-Locked Fiber Lasers . . . . .	14
1.5	Rate Equations for Lasers Based on the $\text{Er}^{+3}$ Gain Medium . . . . .	16
1.6	Frequency Combs . . . . .	22
1.7	Conclusions . . . . .	25
<b>2</b>	Laser Stabilization	27
2.1	Feedback Control: The Basics . . . . .	28
2.2	RF locking via Phase Locked Loops . . . . .	29
2.2.1	Phase Locked Loops and Frequency Combs . . . . .	36
2.3	Optical Frequency Locks: Pound-Drever-Hall Locking . . . . .	39
2.4	Measuring It . . . . .	44
2.4.1	Phase/Frequency Noise Spectral Density . . . . .	46
2.4.2	Timing Jitter . . . . .	49
2.4.3	In-loop error vs. Out-of-loop error . . . . .	50

<b>3</b>	Tight Timing Control of a Mode Locked Erbium Fiber Laser	53
3.1	Using an EOM for fast feedback control . . . . .	56
3.2	Characterizing the performance of the EOM . . . . .	60
3.3	Synchronizing the EOM laser to a free-running mode-locked fiber laser . . . . .	64
3.4	Conclusion for EOM Locking . . . . .	66
3.5	Remote Synchronization . . . . .	66
3.6	Cross-correlator as a timing jitter detector and ultimate limits . .	74
3.7	Putting it into context . . . . .	76
<b>4</b>	Applications of mode-locked fiber lasers	77
4.1	Coupling a Frequency Comb to a High Finesse Cavity . . . . .	78
4.2	Molecular Detection Experiment . . . . .	79
4.2.1	Brief explanation of CRDS . . . . .	84
4.2.2	Broadband CRDS . . . . .	85
4.3	High Harmonic Generation: Moving the frequency comb to the UV and beyond . . . . .	89
4.3.1	The Ytterbium Fiber Laser and Amplifier . . . . .	92
4.3.2	Measurement of the intra-cavity intensity . . . . .	93
<b>5</b>	Introduction to Waveguide Arrays	96
5.1	What are Waveguide Arrays? . . . . .	97
5.2	Waveguide Arrays and Mode-Locking . . . . .	100
5.3	Testing the Saturable Absorber Properties . . . . .	101
5.3.1	Theory of coupled modes . . . . .	103
5.3.2	Experimental System and Measurements . . . . .	104
5.3.3	Dispersion Measurement of the Waveguide Array . . . . .	111
5.4	Conclusions . . . . .	111

<b>6</b>	<b>Multi-Photon Absorption and Spectral Phase Clamping</b>	114
6.1	Discrete Spatial Solitons . . . . .	115
6.2	Multi-photon absorption theory . . . . .	115
6.3	Experimental Setup and Results of the Three Photon Absorption Measurement . . . . .	117
6.4	Chirp effects on self-focusing . . . . .	126
6.5	Spectral Phase Clamping . . . . .	126
6.6	Conclusions . . . . .	134
<b>7</b>	<b>Summary and Outlook</b>	137
	<b>Bibliography</b>	139
	<b>Appendix</b>	
<b>A</b>	<b>Matlab codes</b>	147
A.1	A simple model of a mode-locked laser . . . . .	147
A.2	Pound Drever Hall simulation . . . . .	148
<b>B</b>	<b>Fabricating Waveguide Arrays</b>	151
B.1	Growing the wafer: Molecular Beam Epitaxy . . . . .	151
B.2	Processing the wafer . . . . .	152
B.3	Photolithography . . . . .	155
B.4	Etching . . . . .	158

## Figures

### Figure

1.1	Single-mode and multi-mode lasing . . . . .	5
1.2	Simulation of mode-locking for various numbers of modes . . . . .	6
1.3	Schematic for additive pulse mode locking . . . . .	11
1.4	Nonlinear polarization rotation . . . . .	13
1.5	Pulse shortening rates for active and passive (slow and fast) saturable absorbers . . . . .	15
1.6	Fiber loss versus optical wavelength . . . . .	17
1.7	Energy level diagram for the $\text{Er}^{+3}$ ion . . . . .	18
1.8	Optical frequency comb . . . . .	24
2.1	Phase locked loop schematic . . . . .	30
2.2	Circuit diagrams for PID control and the corresponding frequency response . . . . .	34
2.3	Optical frequency comb on a photodetector . . . . .	38
2.4	Phase Noise floors for thermal and shot noise . . . . .	40
2.5	Reflection power for PDH setup . . . . .	42
2.6	Reflection phase for PDH setup . . . . .	42
2.7	Pound-Drever-Hall error signal . . . . .	45
2.8	In loop vs. Out of loop measurement . . . . .	52

3.1	Atacama Large Millimeter Array telescope . . . . .	55
3.2	The Linac Coherent Light Source . . . . .	55
3.3	Schematic of the Electro-optic modulator . . . . .	58
3.4	EOM fiber laser layout . . . . .	59
3.5	EOM fiber laser spectrum . . . . .	60
3.6	Impulse response of the intra-cavity EOM . . . . .	62
3.7	Transfer function of the intra-cavity EOM . . . . .	63
3.8	Timing jitter spectral density (TSD) and integrated TSD for the EOM locked laser . . . . .	65
3.9	Topology of remote synchronization experiment . . . . .	69
3.10	RF Locking for Remote Synchronization . . . . .	72
3.11	Schematic of the cross correlation timing jitter measurement . . . . .	73
3.12	Optical cross correlation timing jitter results . . . . .	74
4.1	Frequency comb coupling to a passive cavity . . . . .	80
4.2	Time domain interference in an optical cavity (in phase) . . . . .	81
4.3	Time domain interference in an optical cavity (out of phase) . . . . .	82
4.4	Cavity ringdown spectroscopy setup . . . . .	83
4.5	Spectrum of the laser and mirror dispersion curves . . . . .	87
4.6	Broadband CRDS experimental setup . . . . .	88
4.7	Absorption spectrum for CO, NH <sub>3</sub> , and C <sub>2</sub> H <sub>2</sub> . . . . .	90
4.8	Similariton laser layout . . . . .	92
4.9	Coupled laser spectrum and noble gas plasma current . . . . .	95
5.1	Power versus propagation distance for two cases . . . . .	99
5.2	Waveguide model . . . . .	102
5.3	Schematic of the experimental setup for the waveguide pulse chop- ping measurement . . . . .	106

5.4	2D power distribution of the output of the waveguide array at various powers . . . . .	107
5.5	Autocorrelations of each waveguide mode . . . . .	109
5.6	Central waveguide pulse reshaping as a function of input power . . . . .	110
5.7	White light interferometer measurements . . . . .	112
6.1	Experimental setup for the three photon absorption measurement . . . . .	118
6.2	3 photon absorption measurement comparison with theory . . . . .	120
6.3	Autocorrelation and 3D power distribution for normal chirp pulse . . . . .	121
6.4	Autocorrelation and 3D power distribution for shortest input pulse . . . . .	122
6.5	Autocorrelation and 3D power distribution for an anomalous chirp input pulse . . . . .	124
6.6	Relative power in the neighboring waveguides . . . . .	125
6.7	Relative output power versus input pulse chirp . . . . .	127
6.8	Experimental setup for spectral phase measurement . . . . .	129
6.9	Electric field envelope shortening in center waveguide as a function of input peak power . . . . .	130
6.10	Electric field envelopes from the compressor . . . . .	132
6.11	Electric field envelopes from the center waveguide . . . . .	132
6.12	GDD of the input pulses and output pulses . . . . .	133
6.13	Vector plot showing the spectral phase attractor nature of the waveguide array . . . . .	135
A.1	Simulation output for low modulation frequency (PDH). . . . .	150
B.1	Lattice orientation . . . . .	153
B.2	Large view of 110 direction . . . . .	154
B.3	Schematic of the photoresist layer . . . . .	156

B.4	Schematic of the photoresist process . . . . .	157
B.5	Schematic of the etching process . . . . .	159
B.6	Etching of AlGaAs in $1\bar{1}0$ direction . . . . .	160
B.7	Etching of AlGaAs in $110$ direction . . . . .	161

## Chapter 1

### Introduction

Mode-locked lasers are an extremely useful type of laser. Simply stated, their output is a phase coherent train of very short pulses ( $< 1$  picosecond). The usefulness of such a system, while perhaps not obvious at first, is immense. Application of these lasers range from micro-machining metals [1] all the way to facilitating the most precise frequency measurements ever made [2]. Based on the many proposals for new technologies that utilize mode-locked lasers [3, 4], it is clear that these lasers will be an invaluable tool for future technologies.

This thesis deals with a particular type of mode-locked laser known as an Erbium doped, mode-locked fiber laser. These lasers have received much attention due to their low cost, low power consumption, long term robustness, and ease of long distance transmission (through single-mode fiber). The first experiments presented here cover the topic of stabilizing the repetition frequency of these lasers. Initial demonstrations of these lasers showed large amounts of high frequency noise in these systems [5]. Thus, we set out to build a laser with a fast intra-cavity actuator to cancel this frequency noise. The system that was built employed an electro-optic modulator to allow for noise cancellation up to fourier frequencies of  $> 200$  kHz. This fast actuator allowed for 2 orders of magnitude reduction in the timing fluctuations of the laser pulses [6]. We then turned our attention towards synchronizing remotely located fiber lasers using this fast actuator in

conjunction with a stabilized fiber link. Applications of this technology include coherent aperture synthesis of radio telescopes, precise frequency dissemination, and pump-probe timing in long range accelerator experiments [7]. This experiment led to extremely low timing fluctuations for a kilometer scale system [8]. These two experiments are covered in Chapter 3.

Chapter 4 of this thesis presents two atomic physics experiments that were conducted using mode-locked fiber lasers and passive optical cavities. The first experiment employed a home-built Er fiber laser for broadband molecular detection in a broadband cavity ringdown setup [9]. This laser was built with several intra-cavity actuators that allowed for long range and tight frequency control. The second project in this chapter dealt with a commercial Ytterbium doped fiber laser for use in a high-field physics experiment. In this project we coupled the high power Yb laser (13 W average power, 90 fs pulses) to a passive enhancement cavity with an intra-cavity focus to produce a peak intensity of  $3 \times 10^{14} \text{ W/cm}^2$  [10]. This experiment was the initial demonstration of a system aimed at achieving high-order harmonic generation in a noble gas, a process that could allow for shifting a frequency comb to ultra-violet wavelengths.

The last two chapters cover a device that could be used to provide an easy and robust method of mode-locking. Simulations of these devices, known as waveguide arrays, show that they exhibit all of the features required for mode-locking. In the first experiment we set out to test these theories by measuring the pulse shaping characteristics via autocorrelation (Ch. 5). We found that at high peak power the waveguide array effectively shortened the input pulses in the time domain [11]. This behavior is consistent with a saturable absorber, the element required for mode-locking. In Chapter 6, I present the results from a characterization of the waveguide array in terms of spatial power distribution and spectral phase. We observed several interesting phenomenon in this study includ-

ing power clamping of spatial solitons and spectral phase attractor points. This thesis concludes with a look towards the next experiment, which is to build a novel mode-locked laser based on waveguide arrays.

In the following sections of this chapter, I will lay out the fundamentals of the mode-locking process and give a brief historical account of the invention of mode-locking. Then, I will describe several *effective* saturable absorbers. In the last two sections I will introduce the Erbium fiber gain medium and examine the rate equations for such a system, and discuss a powerful tool based on mode-locked lasers known as frequency combs.

## 1.1 Mode Locking Fundamentals

The term *mode-locking* refers to the requirement of phase locking many different frequency modes of a laser cavity. This locking has the result of inducing a laser to produce a continuous train of extremely short pulses rather than a continuous wave (cw) of light. In principle, though, a continuous train of pulses can be generated from a Q-switched laser. The difference between these two scenarios lies in the optical phase of the pulses. The mode locked pulses are phase coherent with each other, while the Q switched pulses are not. This simple fact has massive ramifications in regards to the application of these two types of lasers.

To understand the mode locking process, we will begin by looking at a cw laser in the frequency domain. For a single longitudinal mode cw laser ( $\nu=c/2nL$  for a Fabry-Perot cavity), we have the situation shown in Fig. 1.1-(a). Here, only one resonant mode of the laser cavity overlaps in frequency with the gain medium. Thus, the laser emits a cw beam with a narrow range of frequencies ( $E(t) = E_1 e^{i(\omega_1 t + \phi_1)}$ ). In general, however, the gain medium could overlap with several modes. We can describe the output of such a laser in the time domain as:

$$E(t) = \sum_n^N E_n e^{i(\omega_n t + \phi_n)} \quad (1.1)$$

where the sum is over all of the lasing cavity modes,  $E_n$  is the amplitude of the  $n^{\text{th}}$  mode,  $\omega_n$  is the angular frequency of the  $n^{\text{th}}$  mode, and  $\phi_n$  is the phase of the  $n^{\text{th}}$  mode. For the single-mode laser, this sum just has one term as given above. As we will see, the phase term plays the key role in the difference between incoherent multimode lasing and mode locking.

Let's make things more interesting by increasing the gain bandwidth to overlap with more of the cavity modes (see Fig. 1.1-(b)). In this configuration, there are 3 terms in Equation 1.1. The output of such a laser depends critically on the phase relationship between the 3 modes. If each mode has a randomly varying phase with respect to the other modes, then a time domain detector on the output would show us that the laser is emitting a cw beam with a large amount of intensity noise (see Fig. 1.2-(a)), while a frequency domain detector would show us that the energy was contained in narrow spikes (with lots of intensity noise) spaced evenly by the free spectral range (FSR) of the cavity. However, if we can fix the relative phases to a set value, then the situation changes dramatically (see Fig. 1.2-(b) and (c)). With fixed phase relationships, the three modes can combine to interfere in such a way as to constructively interfere at multiples of the roundtrip time of the cavity, while they destructively interfere elsewhere. This process creates shorter pulses as the number of phase locked modes increases. A MATLAB simulation was used to create Fig. 1.2 (note: the code for this simulation can be found in Appendix A).

A natural question to ponder is exactly how do we achieve this phase locking of the longitudinal modes? To answer this question we will switch to the time domain picture of mode-locking. We know that a mode-locked laser produces

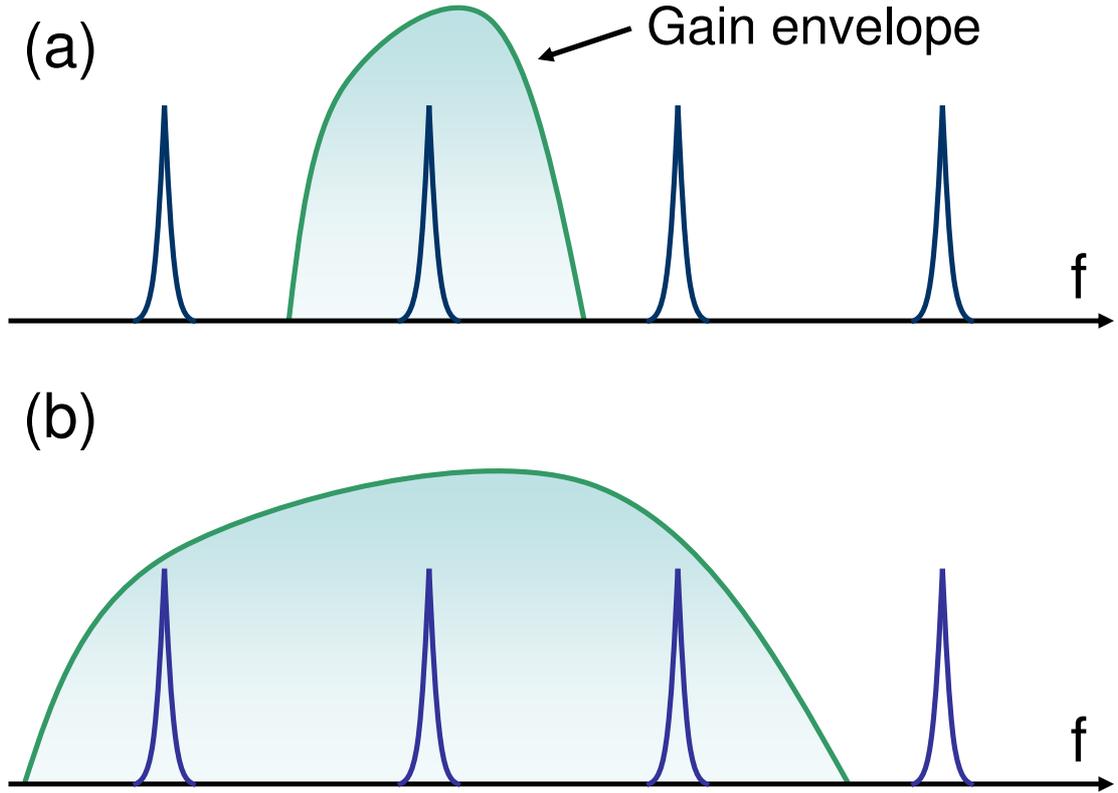


Figure 1.1: Resonant cavity modes and the gain spectrum of a laser. Part (a) shows single-mode lasing, where the gain envelope overlaps with just one cavity resonance. Part (b) shows multimode lasing, where the gain envelope overlaps with several cavity resonances.

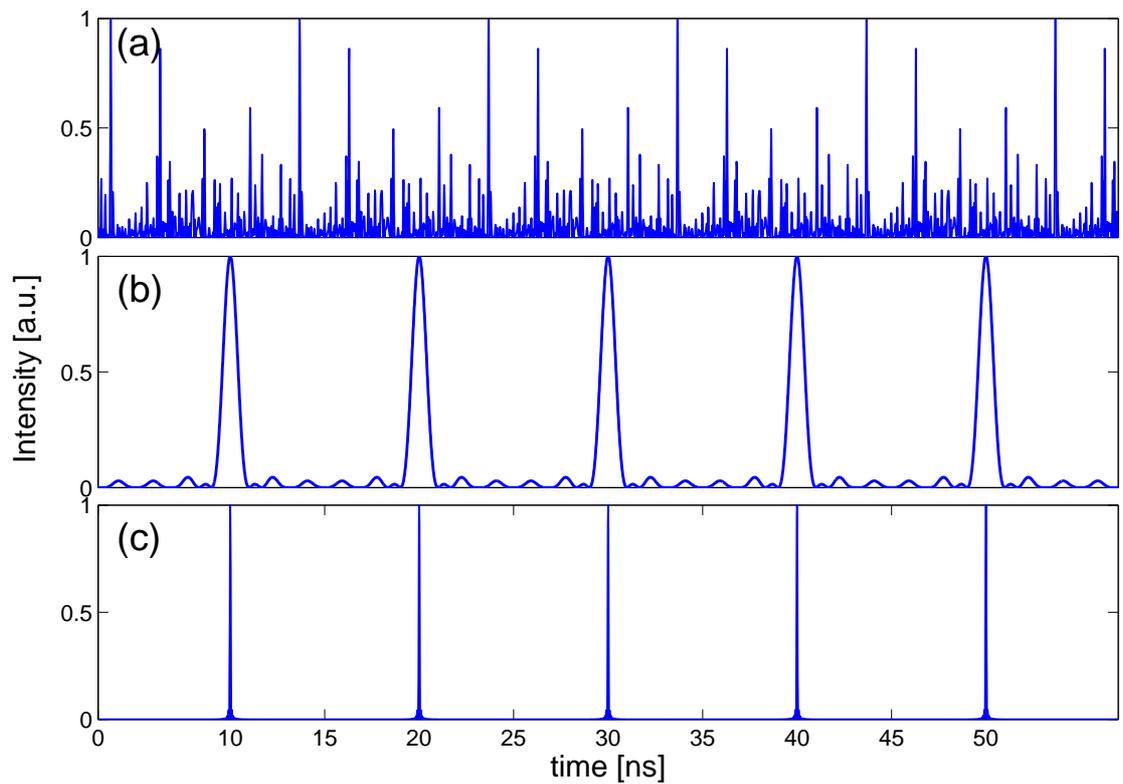


Figure 1.2: Simulation of modelocking. In this simulation, the cavity has  $\text{FSR} = 100$  MHz. (a) No phase coherence between the multiple modes (80 modes simulated), (b) 5 phase coherent modes, and (c) 80 phase coherent modes. The pulse train emerges as ultrashort intensity spikes (fs scale) spaced by longer ns intervals (roundtrip time of the cavity).

ultrashort pulses at a rate equal to the round trip time of the optical cavity. This means there has to be some part of the laser that prefers to make the laser produce pulses over cw radiation. This statement equates to saying that we need some element that provides high loss at low intensity (cw radiation) and lower loss at high intensity (pulsed operation). Such a device is known as a saturable absorber. As will be discussed in the next section, real atoms were the first saturable absorbers used. The operational principles of atoms or molecules as saturable absorbers are straightforward: low intensity light is absorbed by the atoms and re-emitted into  $4\pi$  steradians (i.e. out of the laser cavity), while high intensity light fully excites the atoms and passes most of its photons through the medium. Clearly, the main feature of the saturable absorber is its decreasing loss with increasing intensity. As will be seen in later sections, this behavior can be mimicked with optical processes that have nothing to do with actual atomic or molecular resonance absorption.

## 1.2 A Brief History of Mode Locking

The history of mode-locked lasers began not long after the first demonstration of a continuous wave lasing in 1960. While Maiman's [12] ruby laser was created at Hughes Research Laboratory in California, the creation of the first mode-locked laser would occur on the opposite coast of the United States at Bell Laboratories in New Jersey. In 1964, Hargrove *et al.* [13] used an extremely clever acousto-optic technique to provide a loss modulation in a Helium-Neon laser cavity, which led to the laser being actively mode locked. While the pulses were still relatively long by today's standards (several nanoseconds), this demonstration opened the door for many more researchers to push the boundaries of ultrashort optical pulses. Only one year later, in 1965, Mocker and Collins showed that they could achieve transient locking of the modes of a multimode Q-switched laser [14].

Since only a few modes were involved in this process, the pulse widths were on the order of 10s of ns. Their technique, however, required no active modulator, and thus was the first demonstration of passive mode locking. The component that locked the modes in their laser was a saturable Q-switching dye (cryptocyanine in methanol). The downside to this dye was that it required the laser to be Q-switched in order to saturate and thus the laser emitted mode locked pulses only at the Q-switched intervals. The transient nature of the mode locked pulses proved to be problematic in practical applications (ultrafast spectroscopy, nonlinear optics, etc.). This problem was solved in 1972 when Ippen *et al.* introduced a laser based on the saturable dye (Rhodamine 6G) that could mode lock continuously [15]. The pulses from this laser were found to have pulse widths of only 1.5 picoseconds. After this demonstration, researchers pushed the gain bandwidth further with other types of saturable dyes, and developed external cavity pulse compression techniques (based on adding new spectral content through nonlinearity, then recompressing through chromatic dispersion). This led to the 1987 demonstration by Fork *et al.* of a 6 fs pulse [16]. That's 6 orders of magnitude in 22 years! Few fields in physics can claim that sort of progress in such a short time.

This discussion of the development of mode locked lasers would not be complete without a look at *effective* (fast) saturable absorber systems. Although, atomic/molecular transitions in dyes were the first method discovered to passively mode-lock a laser, this method is by no means the only way. A landmark paper in 1984 by Mollenauer *et al.* demonstrated this fact [17]. This experiment showed that the saturable absorber effect can be simulated by optical phenomenon. This approach has several advantages including the fact that the recovery time of an optically based saturable absorber can be extremely fast ( $\approx$  a few optical cycles) since it does not depend on an atomic/molecular resonance. These types of ab-

sorbers, which have lead to the shortest pulses ever created, are discussed in the next section.

### 1.3 Effective Saturable Absorbers: Kerr Lens Mode locking and Additive Pulse Mode Locking

This section deals with a type of saturable absorber known as an *effective* saturable absorber. For this special class of saturable absorbers, some process other than atomic/molecular absorption is used to achieve a saturable absorption effect. The fact that these absorbers do not have to rely on actual atomic transitions means that the recovery time for the saturable absorber can be much faster than atoms. Slow saturable absorbers produce  $< 1$  ps pulses by shortening the leading edge of the pulse via saturable absorption and the trailing edge via gain saturation. However, if the saturable absorber recovers fast enough, it can shorten both sides of the pulse using the saturable absorber effect. The most common method of achieving such an absorber in the lab is to exploit the intensity dependent index of refraction:  $n(I) = n_0 + n_2I$ , where  $n_0$  is the index of refraction,  $n_2$  is the nonlinear index coefficient, and  $I$  is the optical intensity. Since non-resonant optical processes are extremely fast, the nonlinear index responds on the order a few optical cycles. Thus, the recovery time for a saturable absorber based on this effect is essentially instantaneous. In the following paragraphs I will briefly review two types of mode locking based on effective saturable absorbers: Kerr Lens Mode Locking (KLM) and Additive Pulse Mode Locking (APM).

In most realizations of Kerr Lens Mode locking (KLM) [18] the nonlinear index of refraction creates a lens into the gain medium, which then causes self-focusing of the beam. Combined with an intra-cavity aperture, this effect creates a situation where the cavity prefers pulsed operation: if the laser is in cw operation, there is a high loss due to the aperture, while in pulsed operation the beam focuses

through the aperture with minimal loss. The most popular type of mode-locked laser based on this is the Titanium-doped Sapphire laser. In these lasers, mode-locking has been achieved with *hard* apertures (i.e. irises put into the beam path) and with *soft* apertures (i.e. aperturing in the gain medium). Due to the large gain bandwidth of the Ti:Sapphire medium (650-1100 nm) and the ultrafast recovery time of the saturable absorber mechanism, optical pulses less than 2 optical cycles ( $\approx 5$  fs for Ti:Sapphire wavelengths) have been observed [19].

The operation of APM is slightly less intuitive since it relies on interference of circulating pulses. In the first realization of APM [17], this interference was between pulses in two coupled cavities (see Fig. 1.3). The main cavity has the gain medium and an output coupler, while the secondary cavity has a nonlinear section (i.e. an optical fiber). Pulses that are coupled to the nonlinear cavity experience an intensity dependent phase shift (this is how the system differentiates the peak from the wings). When these pulses are coupled back to the main cavity they can be made to overlap with the normal pulses in such a way as to constructively interfere at their peaks, while destructively interfering at their wings. Thus, the addition of multiple pulses (with interferometric control) results in pulse shortening on every round trip, just like a real saturable absorber.

One special type of APM based on nonlinear polarization rotation (P-APM) [20] is particularly useful in a fiber laser cavity. In this scheme, the pulses that are added are not from separate cavities, but are co-propagating pulses with different polarization. To achieve this, elliptically polarized pulses propagate in a Kerr medium to produce nonlinear polarization rotation. Experimentally this situation can be produced by inserting a quarter-wave plate into the fiber cavity (in a free space section), so that linear polarization can be turned into elliptical (see Fig. 1.4-(a)). The highest intensity part of the pulse (i.e. the peak) undergoes a nonlinear phase shift and thus rotates its polarization some amount. The wings of

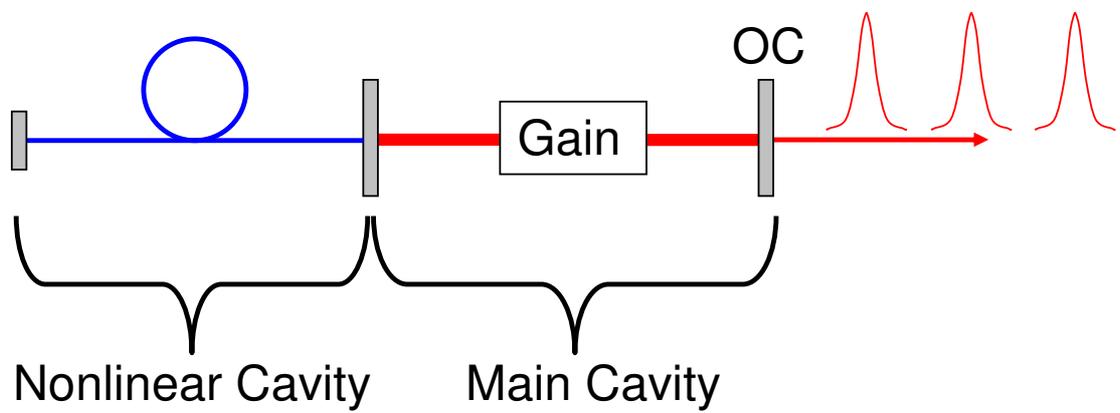


Figure 1.3: Additive pulse mode locking. The nonlinear cavity has a nonlinear section that provides the *effective* saturable absorber. When the cavity lengths are adjusted correctly, the pulses from nonlinear and main cavity interfere constructively at their peaks and destructively elsewhere. This results in a low loss for high intensity, and a high loss for low intensity.

this pulse, which have low intensity, do not undergo this phase shift and thus experience no rotation (see Fig. 1.4-(b)). A quarter-wave plate and linear polarizer at the output of the Kerr medium (fiber) turn the intensity dependent polarization into an intensity dependent transmission (i.e. saturable absorber). This type of mode locking can produce pulse widths that are close to the gain bandwidth limit of Erbium ( $\approx 100$  fs).

I will conclude this discussion on saturable absorbers with a look at the pulse shortening rates of each of these types of mode-locking. Active mode-locking introduces a large loss modulation in the cavity. This strong interaction between the modulator and the cavity means that the pulse shortening rate ( $\delta\tau/\tau$ ) is very high in the initial stages of pulse formation, and is given by  $\delta\tau/\tau = m\omega_m^2\tau^2/4$  (where  $m$  = modulation depth and  $\omega_m$  = modulation frequency) [21, 22]. Thus, as the pulse gets shorter, the pulse shortening rate decreases as the square of the pulse width (see Fig. 1.5). This is the reason that active mode-locking produces relatively long pulses (ps regime). A slow, passive saturable absorber has a pulse shortening rate given by  $\delta\tau/\tau = m_s/2$ , where  $m_s$  is the modulation depth of the absorber. Since the modulation depth only depends on the pulse energy, the pulse shortening rate is independent of pulse width. Finally, the pulse shortening rate for a fast saturable absorber is given by  $\delta\tau/\tau = \gamma W/2\tau$ , where  $\gamma$  is the self-amplitude modulation coefficient and  $W$  is the photon flux density. Thus, the pulse shortening rate is proportional to  $1/\tau$ . This means shorter pulses experience larger pulse shortening rates. This runaway process is eventually limited by the gain bandwidth filter and dispersion in the laser cavity.

We can think of these pulse shortening rates in terms of self-starting as well. The large pulse shortening rate of active mode-locking means that even an initially long pulse will initiate the mode-locking process since there is large pulse shaping per pass. However, since the pulse shortening decreases with pulse width, the

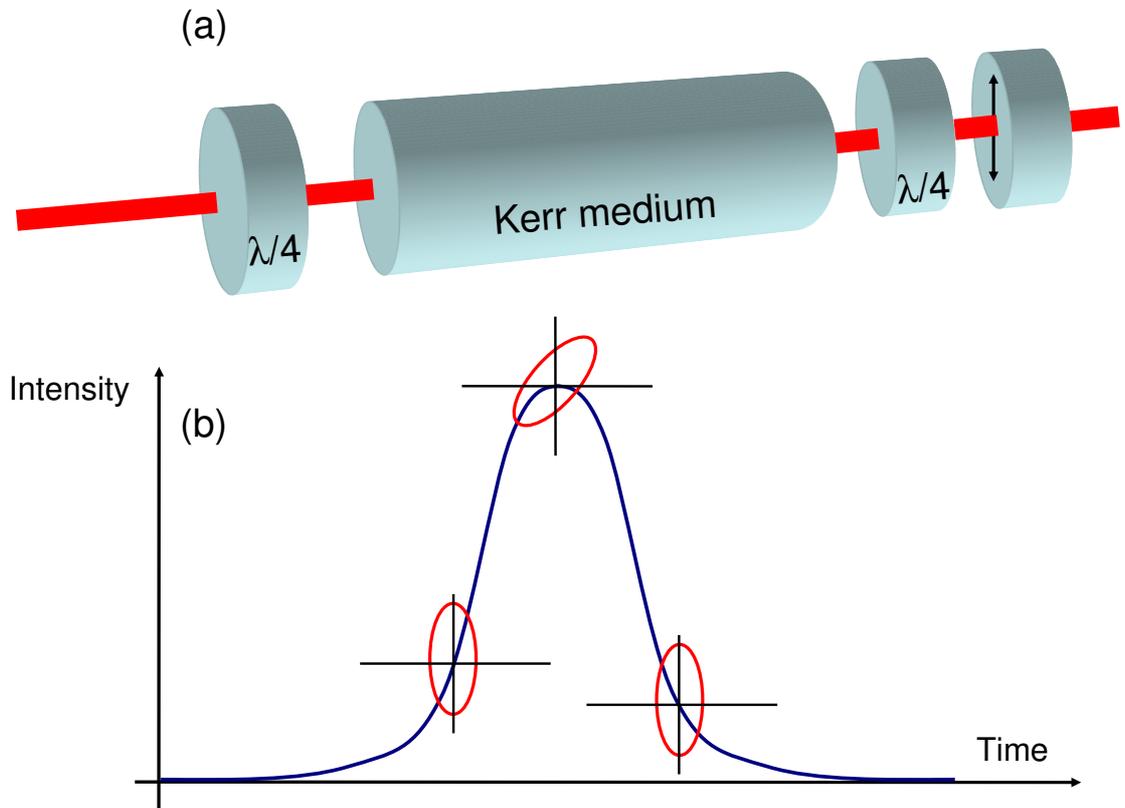


Figure 1.4: Nonlinear Polarization Rotation. (a) Experimental components needed for NPR. Typically, the Kerr medium is a single-mode optical fiber. A quarter-wave plate transforms the elliptical polarization after the fiber back into a linear polarization. The polarizer after this wave plate selects out the peak of the pulse. (b) Schematic of the nonlinear polarization rotation process. The wings of the pulse undergo little or no rotation relative to the peak of the pulse. This polarization difference is turned into a loss when the pulse is analyzed at the polarizer.

ultimate pulse width is limited. The slow saturable absorber has a lower initial pulse shortening per pass, but the performance of the absorber does not fall off with shorter pulse widths. Thus, slow saturable absorber mode-locking has been demonstrated to self-start and produce shorter pulses than active mode-locking. The fast saturable absorber is clearly the hardest to self-start due to its initial pulse shortening rate. A long initial pulse undergoes very little shaping due to the weak effect of the fast saturable absorber. However, once a fluctuation is produced that is short enough to initiate the mode-locking process, the ultimate pulse width is limited only by the broadening effects mentioned above.

Now that we have developed a basic understanding of mode-locked lasers, I would like to examine a particular type of mode-locked laser that has a gain medium based on Erbium doped fiber. Erbium based mode-locked fiber lasers are particularly interesting due to the emission wavelength of Erbium and its ties to the telecommunications industry. In the next section I will discuss the benefits to researchers due to this relationship.

#### **1.4 Erbium-doped Mode-Locked Fiber Lasers**

Mode locked fiber lasers have been the subject of intense research and application over the last 20 years [17, 23, 24, 25, 26, 27, 28, 6]. Erbium based fiber laser systems, in particular, have many advantages over solid-state or gas cavity lasers. An all fiber cavity needs no realignment and the telecommunications components that comprise the active devices (i.e. diode pumps) are specified to last 20 years of continuous use. Also, the components needed to build a fiber laser are relatively cheap due to their mass production in the telecommunications industry. As of 2009, a nonlinear polarization mode-locked fiber laser could be built for an expense of less than 4,000 USD. For comparison, in the same time period a typical solid-state Titanium doped Sapphire mode locked laser could be

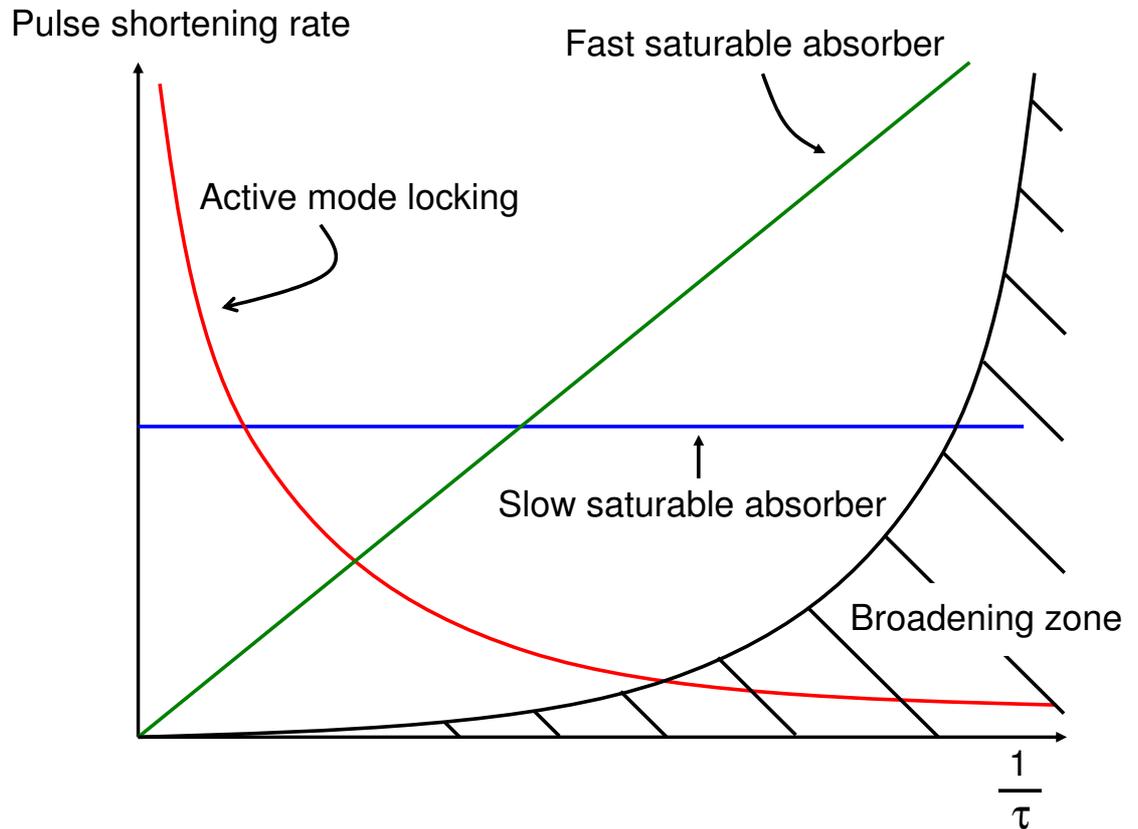


Figure 1.5: Pulse shortening rates for active and passive (slow and fast) saturable absorbers. The poor self-starting ability of most fast saturable absorber lasers is due to their ineffectiveness at long pulse widths. Thus, in many systems a slow saturable absorber starts the mode locking process or some sort of large perturbation is introduced into the system to jump start the mode locking with a shorter pulse than the unperturbed cavity could yield.

purchased from a vendor for around 100,000 USD. While the typical output power of a Ti:Sapphire system is roughly an order of magnitude larger than that of a mode-locked Er fiber laser, it is straightforward and inexpensive to build an Er amplifier that allows the Er-based system to reach average power levels close to those of the Ti:Sapphire oscillator. Using a frequency doubling crystal, one can even transform the 1550 nm centered Er laser to Ti:Sapphire wavelengths around 750 nm. Finally, the relatively small gain bandwidth of the Er gain medium can easily be converted into an octave of spectrum using highly nonlinear fiber. All of these factors have played a part in the rapid emergence of fiber lasers in the world of ultrafast physics in the last 10 years.

Erbium doped fiber is particularly useful over other rare-earth doped fibers (i.e. Ytterbium, Neodymium, Thulium, etc.) due to silica glass's low loss window in the telecommunications C band (Conventional band: 1530–1565 nm). As can be seen in Fig. 1.6, the lowest loss (due to fundamental process of Rayleigh scattering) in the glass fiber is around 0.2 dB/km. This low loss level allows for light from an Er laser to be transmitted over several hundred kilometers, a fact that will be exploited in an experiment in Ch. 3.

### 1.5 Rate Equations for Lasers Based on the $\text{Er}^{+3}$ Gain Medium

To understand the operation of an Er doped fiber laser, we begin with the energy level diagram (see Fig. 1.7). The  $\text{Er}^{+3}$  ion is a quasi-3 level system, meaning that although the lowest state in the lasing scheme is not the true ground state it is still low energy enough that it has some population due to thermal excitation. The two most popular methods of pumping this ion are to use either 980 nm light or 1450 nm light generated from a semiconductor laser (e.g. InGaAs for 980nm). The 980 nm method, which will be discussed in the following paragraphs, and the 1450 nm scheme result in similar efficiencies.

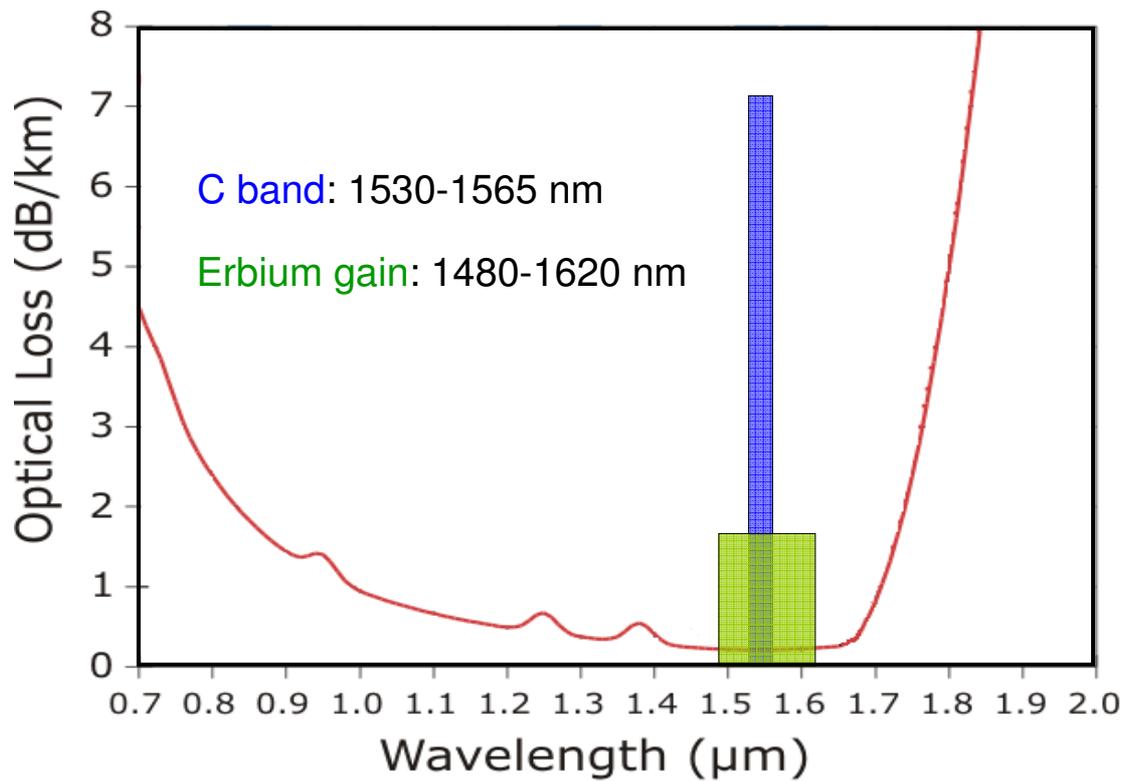


Figure 1.6: Fiber loss versus optical wavelength [29]. The minimum loss is around 1.55  $\mu\text{m}$ . The peaks near .95  $\mu\text{m}$ , 1.25  $\mu\text{m}$ , and 1.39  $\mu\text{m}$  are due to OH absorption. High quality fiber production techniques can minimize the amount of OH that is trapped in the fiber during the drawing process, thus pushing the peaks down to the floor set by the fundamental Rayleigh scattering loss.

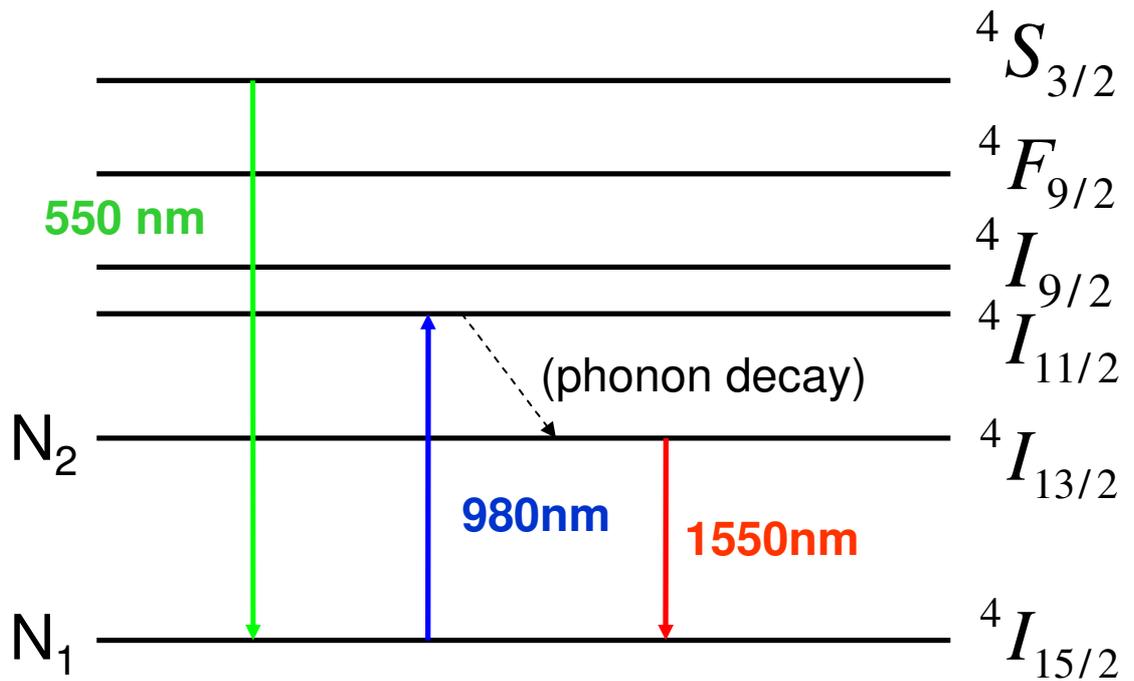


Figure 1.7: Energy level diagram for  $\text{Er}^{+3}$  ion. The  ${}^4I_{13/2}$  state is the metastable state (lifetime  $\approx 10$  ms). The transition from the  ${}^4S_{3/2}$  state (populated by two-photon absorption from the pump) is responsible for the familiar green glow of Er fiber lasers [30].

The trivalent erbium ion, when pumped with 980nm light, is excited to the  ${}^4I_{11/2}$  state, which decays to  ${}^4I_{13/2}$  (see Fig. 1.7). The decay between  ${}^4I_{11/2}$  and  ${}^4I_{13/2}$  is non-radiative (multiple phonon decay) and occurs within a few  $\mu\text{s}$ , while the metastable state ( ${}^4I_{13/2}$ ) has a lifetime of  $\approx 10$  ms. Since the  ${}^4I_{11/2}$  state has such a short lifetime, we can make the approximation that this highest excited state has zero steady-state population (i.e. no population accumulates). This approximation reduces the number of participating energy levels to two. We can now write down the relevant rate equations that describe the number of erbium ions in the upper ( $N_2$ ) and lower ( $N_1$ ) energy levels:

$$\frac{dN_1}{dt} = A_{21}N_2 + (N_2\sigma_e^s - N_1\sigma_e^s)\frac{I_s}{h\nu_s} + (N_2\sigma_e^p - N_1\sigma_a^p)\frac{I_p}{h\nu_p} \quad (1.2)$$

$$\frac{dN_2}{dt} = -A_{21}N_2 + (N_1\sigma_a^s - N_2\sigma_e^s)\frac{I_s}{h\nu_s} + (N_1\sigma_a^p - N_2\sigma_e^p)\frac{I_p}{h\nu_p} \quad (1.3)$$

where  $A_{21}$  is the Einstein A coefficient (inverse lifetime) for spontaneous emission,  $\sigma_{e(a)}^{s(p)}$  is the cross section for stimulated emission (absorption) at the signal (pump) wavelength,  $I_{s(p)}$  is the signal (pump) intensity, and  $h\nu_{s(p)}$  is the energy of each individual signal (pump) photon. By dividing the beam intensity by the photon energy of that beam, we get the total number of photons passing through a given area (i.e. photon flux).

To achieve lasing, we must have population inversion such that  $N_2 > N_1$ . The threshold for this condition occurs when the ion density in  $N_2$  just equals  $N_1$ . By setting the eqns. (1.2) and (1.3) equal and solving for the pump intensity we find the threshold intensity for population inversion:

$$I_{pth} = \frac{h\nu_p}{\tau(\sigma_a^p - \sigma_e^p)} \quad (1.4)$$

For a pump wavelength of 980 nm, this intensity is roughly 6 kW/cm<sup>2</sup>. Since

the mode field area of a single-mode Er fiber is around  $20 \mu m^2$ , the pump power needed to achieve inversion is on the order of a few milliwatts. This calculation is for a lossless cavity, however. Due to losses in fiber splices, the output coupler, and losses in the coupling of the pump diode to the fiber, the actual pump power required for lasing threshold is on the order of 10s of milliwatts (typical 980 nm pump diodes reach average powers beyond 600 mW). It is also instructive to look at the evolution of the signal beam as it propagates through the gain (amplifying) section of the laser cavity. A simple differential equation governs the signal in the presence of a 2-level gain medium:

$$\frac{dI_s(z)}{dz} = (N_2\sigma_e^s - N_1\sigma_a^s)I_s(z) \Rightarrow I_s(z) = I_0e^{gl} \quad (1.5)$$

where  $I_0$  is the intensity entering the gain section,  $g$  is the gain (given by  $N_2\sigma_e^s - N_1\sigma_a^s$ ), and  $l$  is the total length of the gain section. For our analysis, we will consider the absorption of the signal beam to be zero, thus  $g = N_2\sigma_e^s$ . The gain is then dependent only on the density of excited atoms  $N_2$  and the emission cross section of the excited Er atoms at the signal wavelength ( $\sigma_e^s$ ). The emission cross section is a constant, thus to determine the gain we only have to find  $N_2$ . Using eqn. 1.3, we have:

$$\frac{dN_2}{dt} = -A_{21}N_2 + (-N_2\sigma_e^s)\frac{I_s}{h\nu_s} + (N_1\sigma_a^p)\frac{I_p}{h\nu_p} \quad (1.6)$$

In the small signal limit, the pump intensity is much larger than the signal intensity ( $I_p \gg I_s$ ). Using this approximation along with the fact that we are analyzing a steady-state scenario ( $d/dt \rightarrow 0$ ) we can ignore the  $I_s$  term and set the left hand side of eqn. 1.6 equal to zero. Solving for  $N_2$  yields:

$$N_2(I_s \ll I_p) = \tau N_1\sigma_a^p \frac{I_p}{h\nu_p} = \tau R \quad (1.7)$$

where  $R$  is the rate at which ground state atoms are excited to the metastable state. This equation shows that the density of excited atoms in the small-signal limit is simply given by the lifetime of the excited state ( $\tau$ ) multiplied by excitation rate  $R$ . Using the fact that  $g = N_2\sigma_e^s$ , the small signal gain is  $g_0 = \tau N_2\sigma_e^s R$ . As the signal beam is increased to higher intensity, however, we must take into account the term in eqn. 1.6 that involves  $I_s$ . Solving for  $N_2$  yields:

$$N_2 = \frac{N_2(I_s \ll I_p)}{1 + I_s/I_{sat}} \quad (1.8)$$

And the large signal gain is thus:

$$g = \frac{g_0}{1 + I_s/I_{sat}} \quad (1.9)$$

where  $I_{sat} = 1/\sigma_e^s\tau$  is the saturation intensity. And finally, the differential change in signal intensity per length of gain in the strong pump regime is:

$$\frac{dI_s}{dz} = \frac{I_s g_0}{1 + I_s/I_{sat}} \quad (1.10)$$

The picture of the signal evolution is now complete. At low signal levels, there is an exponential increase in the number of signal photons in the gain medium. However, as the signal level is increased further the gain begins to saturate and asymptotically approaches a value of  $\approx I_{sat}g_0 = R$ . Thus, at high signal levels, the signal intensity increases linearly with the pump intensity.

The fundamental characteristics of lasing, small-signal gain, and gain saturation have now been covered. This should provide the reader with a basic working knowledge of  $\text{Er}^{+3}$  fiber lasers. This type of laser is used extensively throughout this thesis. In the last section of this chapter, I will briefly discuss a very important technology based on mode-locked lasers: the frequency comb.

## 1.6 Frequency Combs

With the advent of the frequency comb [31] in the late 1990s, mode-locked lasers began to receive much attention from the frequency metrology community. In principle, the frequency comb could have been realized much earlier. The long delay between the first mode-locked laser and the first frequency comb can probably be attributed to the vast gap between the two fields that were necessary to create the frequency comb. On one side, the field of precision measurements was focused on creating actuated lasers that would have the narrowest possible spectral linewidth. On the other side, the field of ultrafast spectroscopy was mainly interested in creating extremely short time domain bursts of electric field, which necessarily require that the pulses have a large spectral bandwidth. These two goals, which seem to be in direct opposition of each other, can be achieved simultaneously with a frequency comb.

In the following paragraphs, I will give a description of the frequency comb from a graphical viewpoint [32]. This description will not touch on many advanced frequency comb topics that could be covered such as fixed points [33], nonlinear effects, or octave spanning oscillator techniques [34]. As was mentioned, the frequency comb is based on mode locked lasers. In fact, many researchers use the two words interchangeably. This is not quite right, however, since technically a frequency comb really refers to a mode locked laser that has been carrier-envelope phase stabilized. To understand the meaning of this statement, let's look at a time domain picture of the pulses emitted from a mode locked laser (see Fig. 1.8). The time domain output of the laser can be viewed as the multiplication of the fast electric field oscillations and an envelope function (the ultimate limit on the envelope width would be an envelope that encompasses only 1 cycle of the electric field). It can be shown that the envelope travels at a speed known as the *group*

velocity ( $v_g = \frac{c}{n - \lambda \frac{dn}{d\lambda}}$ ), while the fast oscillations travel at the *phase* velocity ( $v_p = c/n$ ). These two velocities are, in general, not equal and thus lead to a walk-off or slippage between the two entities, known as carrier-envelope offset phase. Using the shift theorem of Fourier transforms [35], we see that the Fourier transform turns this time domain phase slip into a frequency offset,  $f_o$ . Thus, the optical frequencies of the comb can be written in terms of two radio frequencies as

$$\nu_n = n f_{rep} + f_o \quad (1.11)$$

where  $\nu_n$  is the optical frequency of the  $n^{th}$  comb mode and  $f_{rep}$  is the repetition frequency of the laser. Clearly, a random variation of the offset frequency would smear out the comb in frequency space and make it useless for any sort of precision measurement. An analogy to this sort of measurement would be like trying to measure the length of something with a ruler that is always moving back and forth slightly. Thus, it is clear that to do any sort of precision measurement with a mode-locked laser, we need to stabilize this offset frequency (and thereby produce a frequency comb).

The first technique that achieved the ability to measure (and thus stabilize)  $f_o$  relied on the so-called f-2f interferometer (quite a technical achievement by itself!). This technique is based on a simple manipulation of equation 1.11. In this scheme, light from a Ti:sapphire laser was sent through a highly nonlinear fiber with low net dispersion to broaden the bandwidth of the pulses to an octave [31]. The octave spanning pulses were then coupled into an interferometer where in one arm the light was passed through a second harmonic crystal and underwent sum-frequency-generation (SFG). The two beams were then recombined on a beamsplitter, sent through an optical filter, and detected onto a photode-

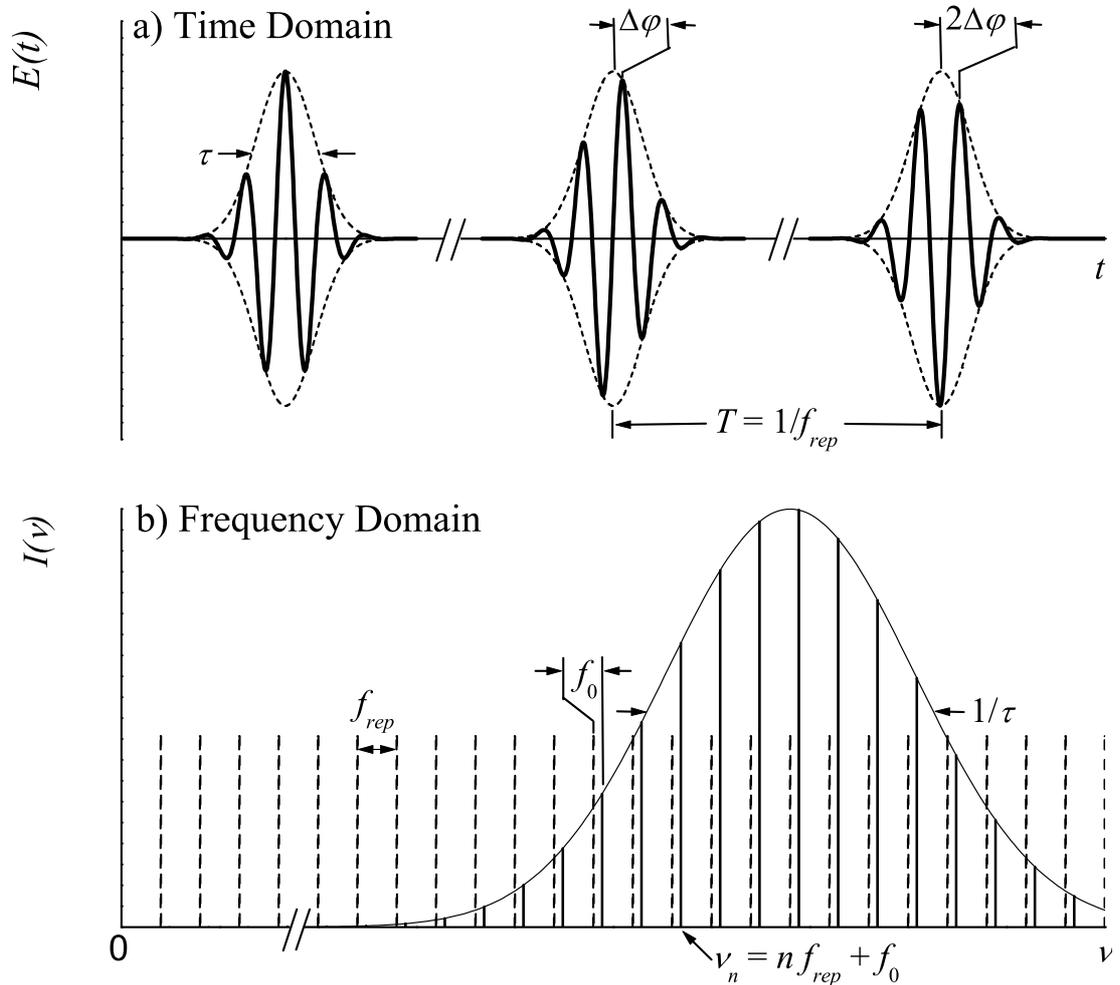


Figure 1.8: Optical frequency comb (courtesy of S. T. Cundiff). (a) Time domain picture of the electric field pulse-to-pulse evolution showing the carrier-envelope phase slip. (b) Frequency domain picture of the frequency comb. The carrier envelope phase slip translates, via a Fourier transform, into an overall shift of the comb ( $f_o$ ). The  $f_o$  degree of freedom is thus referred to as the translational mode of the comb, while the  $f_{rep}$  degree of freedom is referred to as the breathing mode of the comb (since  $f_{rep}$  controls the spacing between frequency modes)

tector to produce a heterodyne beat at  $f_o$ . A bit of mathematics is required to understand why this happens. The octave spanning pulse bandwidth ensures that we have optical frequencies present in a range from  $\nu_n$  to  $\nu_{2n}$ , while the second harmonic arm converts the  $\nu_n$  light to  $\nu_{2n}$  light via SFG. Filtering out the highest frequencies with the optical filter, and using the frequency comb equation, we can thus write the frequencies present in the two arms as:

$$\nu_{2n} = 2nf_{rep} + f_o \quad (1.12)$$

$$2\nu_n = 2nf_{rep} + 2f_o \quad (1.13)$$

Once these two beams form a heterodyne beat on the photodetector, we can take the difference frequency which is:

$$2\nu_n - \nu_{2n} = (2nf_{rep} + 2f_o) - (2nf_{rep} + f_o) = f_o \quad (1.14)$$

The first demonstration of this method [31] opened the door for an explosion of experiments involving the frequency comb. Precision metrology benefited dramatically from the compact all-in-one nature of the frequency comb (compared to the comb's predecessor: the frequency chain [36]), while new techniques such as broadband cavity-ringdown spectroscopy [37] have been developed based on the comb (see Ch. 4 for more details).

## 1.7 Conclusions

The purpose of this chapter was to introduce the underlying technology used in this thesis: the mode-locked laser. Each of the projects in the following chapter relies on a laser we built in-house (except for one experiment). The reader should now have a good idea of how a mode-locked  $\text{Er}^3$  fiber laser operates. The

next chapter will examine some details of laser stabilization that are crucial to understanding in the subsequent chapters.

## Chapter 2

### Laser Stabilization

Mode-locked fiber lasers can have significant amounts of noise (i.e. phase or intensity noise) on their outputs. This fact is due to several aspects of the fiber laser: the cavities are typically high gain/high loss environments, acoustic fluctuations can cause pressure variations on the fiber at high frequencies, and thermal fluctuations can cause long timescale drifts of the repetition frequency. If these lasers are to find use as a frequency comb for applications such as high-precision spectroscopy or gearing for an optical clock, this situation must be remedied.

The issue of laser noise was addressed very early on in the development of lasers. In fact, the first paper to be published on active feedback control to reduce laser frequency noise was published only 5 years [38] after the first demonstration of the laser by Maiman [12]. The art of laser stabilization is thus quite advanced at the present, with many demonstrations of clever techniques [39, 40]. In principle, the only difference between the various techniques is the way in which they generate the error signal. Once the error signal is created, it can then be filtered and amplified before being used to control an intra-cavity actuator. In this chapter, I will discuss the fundamentals of several laser locking schemes that will be used in later chapters, and discuss several metrics for evaluating system performance.

## 2.1 Feedback Control: The Basics

Feedback control is simply a technique that controls a process by performing a comparison between a system's output and a reference and then adjusting the output in some way. There are two fundamentally different types of feedback: positive and negative. Positive feedback increases the deviation from the normal output. Perhaps the most well-known form of positive feedback is amplification. Negative feedback, on the other hand, minimizes the difference between the output and some reference standard. In this way, the system's output tracks the reference. While it is common in the field of laser stabilization to think that negative feedback is the more useful of the two types, this perception was not always so. In fact, when Harold Black applied for the first patent for negative feedback in 1928 [41], the patent was not granted until more than 9 years later due to the perception that his device would not do anything useful; the patent office reportedly treated it as if it were a perpetual motion machine. Over 80 years later, negative feedback in the form of phase-locked loops (PLLs), key components in FM signal demodulation, is ubiquitous in modern electronics (even available as an integrated circuit). PLLs are also quite common in many labs around JILA where they are implemented in laser stabilization schemes.

Negative feedback in the context of laser stabilization has a rich history with many well developed techniques. The main idea is that if you can somehow derive an error signal (that represents how much the laser's phase or frequency has drifted relative to some reference), then that error signal can be shaped and fed back to an actuator in the laser that controls the laser frequency. In this way, the output frequency of the laser can be stabilized to extremely high levels (coherence times  $\sim$  seconds). While there are a multitude of methods in laser frequency stabilization, these techniques can be broadly lumped into three general categories: phase-

locking to an RF reference, optical reference cavity locking, and spectroscopic transition locking. Phase-locking to an RF reference is experimentally the easiest to implement in terms of required components, but this method does have some drawbacks in ultimate performance. Optical reference cavities yield the tightest locking ( $\sim$ mHz linewidth at optical frequencies), but require a narrow linewidth reference cavity [42]. Spectroscopic transition locking provides long term stability due to the locking of the laser to a natural atomic/molecular resonance. This chapter will examine the first two techniques since they will be used in later chapters.

## 2.2 RF locking via Phase Locked Loops

In this section, I will describe how phase-locked loops can be used to stabilize the pulse train of a mode-locked laser. The fundamental elements required in such a phase-locked loop system are shown in Fig. 2.1. The goal of the PLL is to keep the repetition frequency of the laser at a steady value. This repetition frequency can have lots of noise depending on the construction of the laser. Processes that change the length of the fiber are important as these small length fluctuations cause repetition frequency fluctuations ( $\Delta l/l = \Delta f_{rep}/f_{rep}$ ). Processes such as pressure fluctuations from acoustic waves, thermal fluctuations of the fiber, and mechanical vibrations can all conspire to create a noisy system. In fact, a free-running  $\text{Er}^{+3}$  mode-locked fiber laser can easily have  $> 1$  ps of timing jitter. By locking  $f_{rep}$  to a low-noise reference, we can transfer the stability of the reference to the pulse train and reduce this timing jitter to much lower levels.

To accomplish this, a portion of the output of the oscillator is split off and detected on a photodetector. This signal, which is now an electrical signal, is compared to a reference that produces, in principle, an extremely stable frequency. The mixer device is used for this comparison. As will be discussed later, when

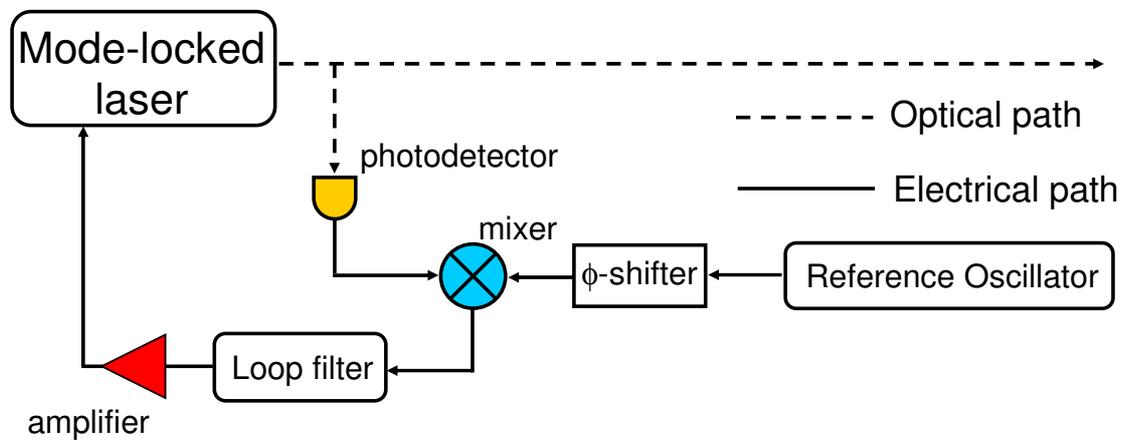


Figure 2.1: Schematic of the fundamental elements of a phase locked loop used for stabilizing the pulse train from a mode-locked laser. The photodetector measures the repetition rate of the laser ( $f_{rep}$ ), which is then stabilized to some RF reference.

the two frequencies are identical the mixer produces a voltage that depends on the phase difference between the two signals. This voltage represents an error signal that can be shaped in the frequency domain by a loop filter and amplified before being fed back to the oscillator. An actuator in the oscillator uses the error signal to correct the oscillator frequency such that the phase difference between the oscillator and the reference is reduced to 0. In the following, I will describe the PLL mathematically and derive some equations that predict the performance for the feedback loop.

The mixer is at the heart of the comparison between the system and the reference. In effect, its output is an analog multiplication of two sinusoidal signals. The signal from the photodetector can be written as  $a(t) = A \sin(\omega t + \phi(t))$ , while the reference signal can be written as  $b(t) = A \sin(\omega t + \phi_2(t))$  (when the reference frequency is equal to the repetition frequency). These two signals are then multiplied in the mixer to yield:

$$V_{mixer}(t) = \frac{K_m AB}{2} \{ \sin(2\omega t + \phi_1(t) + \phi_2(t)) - \sin(\phi_1(t) - \phi_2(t)) \} \quad (2.1)$$

where  $K_m$  is the mixer coefficient. The first term oscillates at twice the fundamental frequency ( $\omega = 2\pi f_{rep}$ ), while the second term is the phase sensitive (DC) term. We are interested in the second term since a phase difference between the two signals produces a voltage. Experimentally, we can easily separate the two signals by low-pass filtering the mixer output to produce just the phase sensitive DC term. One slight detail is missing from this discussion, however. How did we choose that both mixer inputs were sine waves (two cosine waves would work equally as well) instead of one or both being a cosine wave (i.e. shifted by  $90^\circ$ )? Experimentally, it is quite easy to get a time delay difference in the two paths such that one oscillating signal was delayed/advanced by  $90^\circ$  or more. To account

for this, a phase shifter is typically placed in the local oscillator (LO) arm before the mixer (see Fig. 2.1). Using the phase shifter, the experimenter can bring the two signals into phase, or to say this in mathematical terms: make them the same trigonometric function. The end result of all of this is that we have a voltage signal that is proportional to the *phase difference* between our reference signal and our device's signal. Thus, it is only a matter of implementing negative feedback to an actuator in our laser to keep the phase difference (mixer output voltage) at zero.

This next section will take a look at how we shape the error signal through filtering and amplification before feeding back to the actuator. The general classification of this process is called loop filtering, and the construction of the loop filter plays a critical role in achieving a tight phase-locked loop. The loop filter is the only knob in the feedback loop that can be changed to match the actuator's response. Indeed, optimization of the loop filter results in the best achievable lock for a given actuator, thus it is of paramount importance to understand the operation and design considerations of this critical piece.

The simplest feedback loop we can imagine is just to directly take the low-pass filtered output of the mixer and connect it to the actuator. In this scheme, we would hope that an increase in voltage from the mixer tunes the actuator such that it changes the laser in such a way as to reduce the output voltage of the mixer (negative feedback). An equally likely possibility in this scheme is that an increase in mixer voltage tunes the actuator such that it results in an increase in mixer voltage (positive feedback). Thus a slight change in the mixer voltage would immediately cause the system output to run away from any sort of set point. To ensure that we don't end up in a positive feedback system, we could put some electronics (loop filter) in between the mixer output and the input to the actuator that would allow us to flip the sign of the error signal. Thus, we could take our very basic loop filter to any phase locked loop and be assured that we

could arrange the system to be in negative feedback mode. Very quickly, however, we would learn that this simple loop filter has many shortcomings. First of all, the signal from the mixer is most likely too weak to drive the actuator to its full dynamic range. Thus, some sort of amplification of the error signal would be nice.

A more subtle flaw in our loop filter has to do with the fact that we are feeding the full error signal (i.e. all frequencies) back to our actuator. Our actuator won't be able to tune the laser frequency at infinite speed. In fact, there will be a characteristic frequency above which the actuator will move the wrong way (positive feedback). We do not want to provide an error signal at or above this frequency for obvious reasons. Thus, the loop filter needs to have some adjustable filters that allow us to shape the frequency content of the error signal such that it removes all frequencies higher than the cutoff frequency. The loop filter we want now contains all of the essential elements: sign correction, filtering, and amplification. The sign correction and amplification parts are trivial. In the following paragraphs I will examine in detail the filtering process and develop some basic knowledge that every experimenter using a PLL should know.

To optimize the phase locked loop, we need to give the actuator an error signal that is custom shaped to match the actuator's frequency response. In frequency ranges where the actuator responds well, the feedback signal needs to be large so that the actuator can tightly lock to the phase reference. In frequency ranges where the actuator does not respond quickly enough or has a resonance, we need to roll off the feedback signal so that we do not provide positive feedback or drive a resonance. To achieve this sort of control over the feedback signal, we must use some basic operational amplifier circuit designs. The three fundamental types of gain are proportional (flat gain versus frequency), integral (decreasing gain versus frequency), and differential (increasing gain versus frequency) (PID). The circuit diagrams and gain versus frequency plots for each type of gain can be

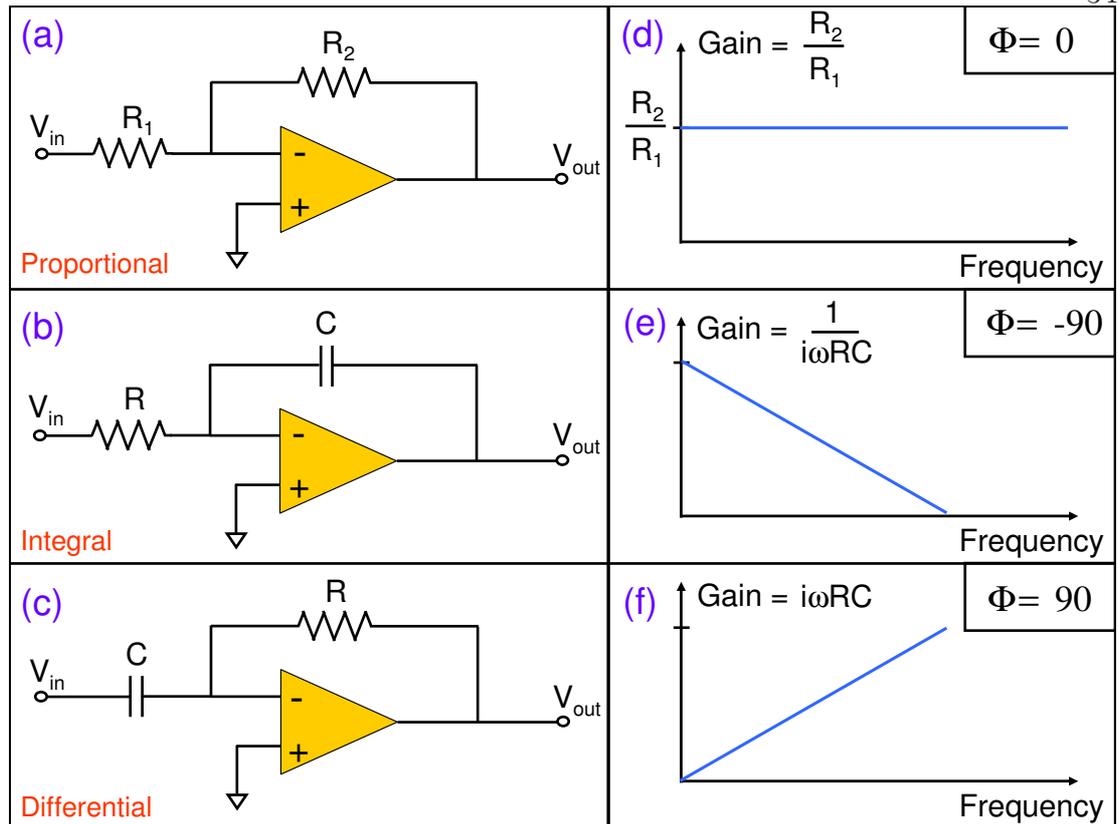


Figure 2.2: Basic circuit diagrams and frequency response of the gain for PID control. (a) and (d) show the circuit diagram and gain versus frequency for proportional gain. The gain is flat across all frequencies, and is just the ratio of the feedback resistor to the input resistor. (b) and (e) show the same plots for integral gain, which exhibits a gain roll-off at a rate of -6 dB/octave. (c) and (f) show derivative gain, which has the same slope magnitude as the integral gain, but with opposite sign.

seen in Fig. 2.2.

The functioning of each of these circuits can be understood using the basic rules of the operation of operational-amplifiers. For instance, the two inputs to the op-amps must be kept at the same potential, or the node at the inverting input must have no current flow. Using this rule, let's first analyze the proportional gain circuit. If  $V_{in}$  goes from zero to some positive voltage, then the current flowing into the node through resistor  $R_1$  is  $I_{in} = V_{in}/R_1$ . To cancel this current, the output of the op-amp has to increase its current output from zero to  $I_{out} = V_{out}/R_2$ .

When  $I_{in} = I_{out}$ , the current flowing into the noninverting input is zero, and the output voltage from the op-amp is  $V_{out} = \frac{R_2}{R_1} V_{in}$ . Thus, proportional gain is independent of frequency, and given by  $G = R_2/R_1$ . For frequencies at which the actuator responds well we should use proportional gain to amplify the feedback signal. However, if we just use proportional gain, we will eventually amplify either a resonance in our actuator or provide positive feedback at high frequencies since proportional gain is the same for all frequencies. Also, proportional gain will actually never allow the system to reach zero error as it becomes vanishingly weak close to zero. Thus, we need some way to roll off the gain such that we can effectively remove frequencies beyond a certain point from our feedback signal and correct this zero error issue of proportional gain. This is where integral gain becomes useful (the circuit diagram for integral gain is shown in Fig. 2.2-(b)). Following the same analysis that we did for proportional gain, the current flow into the inverting input node is held to zero by the op-amp. Again, any current flowing into the input will be detected by the op-amp, and the op-amp will swing its output voltage to compensate for this and make the current flow into the node go back to zero. This time, however, the output sees a feedback capacitor instead of a resistor. Using the previously derived equation for the gain of a feedback mode op-amp,  $G = R_2/R_1$ , we simply need to know the effective resistance of the capacitor (i.e. capacitive reactance). From basic electronics the capacitive reactance is given by  $\chi_c = -\frac{1}{i\omega C}$ . Thus, the gain of this circuit is:

$$G = Z_C/Z_R = -\frac{1}{i\omega RC} \quad (2.2)$$

This equation shows the gain is now frequency dependent. The magnitude of this gain can be calculated by taking the absolute value squared, while the phase of the gain can be calculated by  $\phi = \arctan \frac{Im(G)}{Re(G)}$ . The magnitude of the

gain for the integrator rolls off at a rate of -20 dB/decade (or -6 dB/octave), while the phase of the gain is  $-90^\circ$ . The integrator gain can be thought of as integrating the error signal over some time. Thus, any small error offset of proportional gain is integrated up and corrected by this type of feedback. Finally, the gain for the differentiator (as seen in Fig. 2.2-(c)) is essentially the inverse of the integrator's gain. Since the feedback resistor is divided by the input capacitive reactance, the formula for the gain is:

$$G = Z_R/Z_C = -i\omega RC \quad (2.3)$$

The magnitude of the gain for the differentiator is shown in Fig. 2.2-(f). The slope of the gain response is +20 dB/decade (or +6 dB/octave), and the phase of the differentiator is  $+90^\circ$ . These three op-amp circuits form the basis for all loop filter control. These basic elements can be combined to form complex signal shaping circuits that allow the experimenter to optimize the feedback circuit so that the full potential of the actuator is realized.

### 2.2.1 Phase Locked Loops and Frequency Combs

Phase locked loops can be used to stabilize the  $f_{rep}$  degree of freedom of a mode-locked laser. In fact, the scheme from Fig. 2.1 can be used. A very relevant question though, is what happens when you detect a pulsed laser on the photodetector rather than a cw laser. To answer this question, let's begin by taking a look at the optical frequency domain picture of the frequency comb (Fig. 2.3-(a)). As we have seen before, the energy is contained in the comb teeth at regular intervals of  $f_{rep}$  (note:  $f_o$  doesn't matter here). When the frequency comb is photodetected, the optical comb modes produce beats in the photodiode semiconductor that effectively downshift the optical comb to a radio frequency (rf) comb 2.3-(b). This effect essentially moves the frequency comb into a frequency

range that is commensurate with the speed of modern electronics (MHz to GHz).

Once the rf comb is generated, one comb mode is selected (via band-pass filtering) for locking to an rf reference. The rf comb mode signal is then sent to the mixer, where it is mixed with a stable local oscillator (LO) reference. To achieve the best performance, the mixer needs to have a strong signal from both the LO source and the photodiode. Typically, a large amount of amplification ( $\sim 30$  dB) is needed to bring the photodiode signal up to a usable level for the mixer ( $\sim 0$  dBm for most mixers). As will be discussed in the next chapter, careful photodiode design can increase the rf signal directly from the photodiode by as much as 10 dB. Still, for most phase detection systems, some level of amplification is required.

The phase noise detection process is susceptible to other noise processes that are not related to actual phase errors between the two frequency sources. These sources of *extra* phase error include amplifier flicker noise ( $1/f$ ), amplitude-to-phase conversion, shot noise, and thermal noise [43]. Of these noise processes, shot noise and thermal noise (also known as Johnson noise) are fundamental. The SSB phase noise floor due to thermal (Johnson) noise is given by [44]:

$$L_{\phi}^{thermal}(f) = \frac{kTR}{2V_0^2} = \frac{kT}{2P_{rf}} = \frac{kT}{2R(\eta^2 P_{opt}^2)} \quad (2.4)$$

where  $T$  is the temperature,  $R$  is the resistance,  $\eta$  is the responsivity of the photodetector,  $P_{rf}$  is the electrical power, and  $P_{opt}$  is the optical power incident on the photodetector. The SSB phase noise floor due to photon shot noise is given by [45]:

$$L_{\phi}^{shot}(f) = \frac{ei_{avg}R}{P_{rf}} = \frac{e}{\eta P_{opt}} \quad (2.5)$$

where  $e$  is the electron charge and  $i_{avg}$  is the average photocurrent. A typical InGaAs photodetector has a responsivity ( $\eta$ ) of around 0.9 mA/mW at

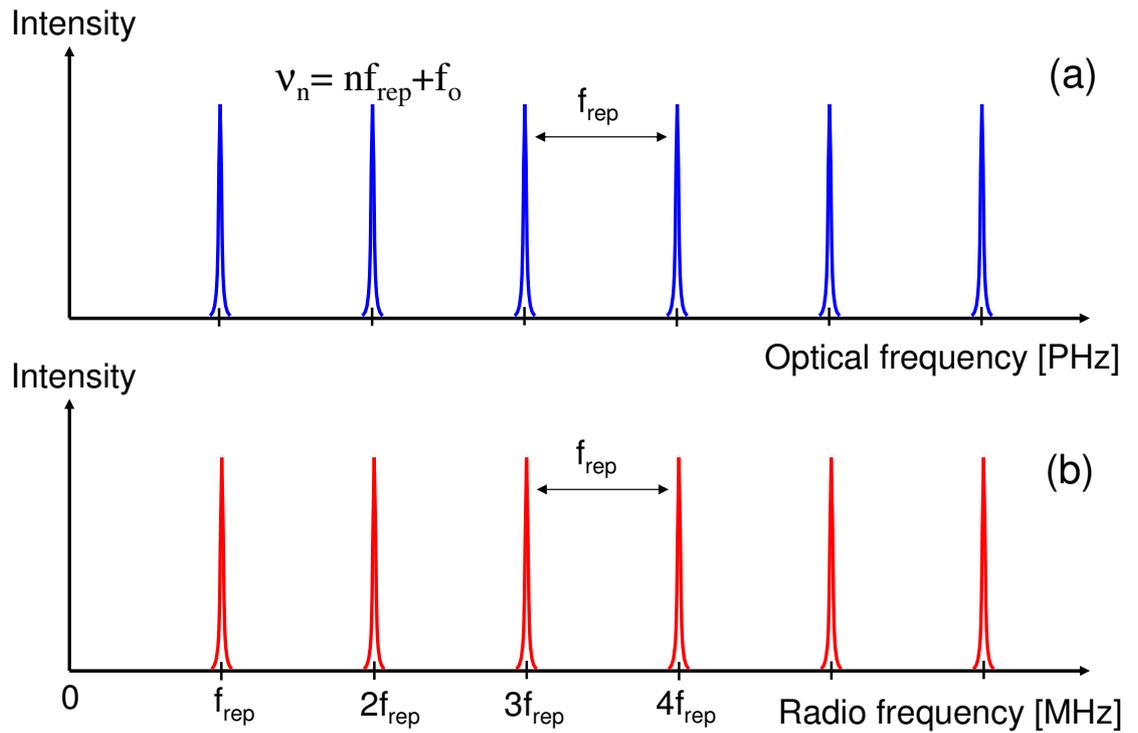


Figure 2.3: Frequency domain picture of a frequency comb. (a) An optical frequency comb with repetition frequency  $f_{rep}$ . (b) The downshifted comb at radio frequencies. Due to beating between optical comb modes on the photodiode, a stable radio frequency comb is generated. The information in this signal can be used for phase locking the frequency comb to an rf reference.

1550 nm. In the equations above, I used the fact that the rf power can be related to the incident optical power on the photodetector by using the responsivity:  $P_{rf} = (\eta^2 R) P_{opt}^2$ . An important point to take away from eqns. 2.4 and 2.5 is that the thermal noise decreases as  $1/P_{opt}^2$ , while the shot noise decreases as  $1/P_{opt}$ . This leads to a crossover point where the thermal noise goes below the shot noise at sufficiently high power. A simple log-log plot of this behavior is shown in Fig. 2.4. Clearly, photon shot noise presents the ultimate noise floor for such a phase detection system.

### 2.3 Optical Frequency Locks: Pound-Drever-Hall Locking

To understand how to best lock a laser to an optical cavity we begin with an analysis of how to best use an optical cavity to derive an error signal. A simple way to implement the cavity as a frequency discriminator is to simply monitor the transmission of the laser power through the cavity. The common side-lock technique does just this by using a photodetector at the output of the cavity to monitor the cavity transmission power. By locking the laser frequency to a point on the side of the transmission resonance peak, one can achieve locking by forcing the laser frequency to always give a set intensity at the output photodetector. This technique suffers from a few drawbacks, however. The most obvious flaw is that amplitude noise from the laser is detected as frequency noise by the photodetector. Thus, via the feedback loop, amplitude noise is written onto the laser as frequency noise. Another drawback of this technique is that if the laser frequency jumps to the other side of the transmission peak, the feedback loop will have the wrong sign of correction and the lock will be broken.

A better method would be to look at the reflected beam and simply lock the laser such that it minimizes the reflection from the cavity (i.e. keep the laser frequency equal to the cavity resonance frequency). At resonance, the reflected

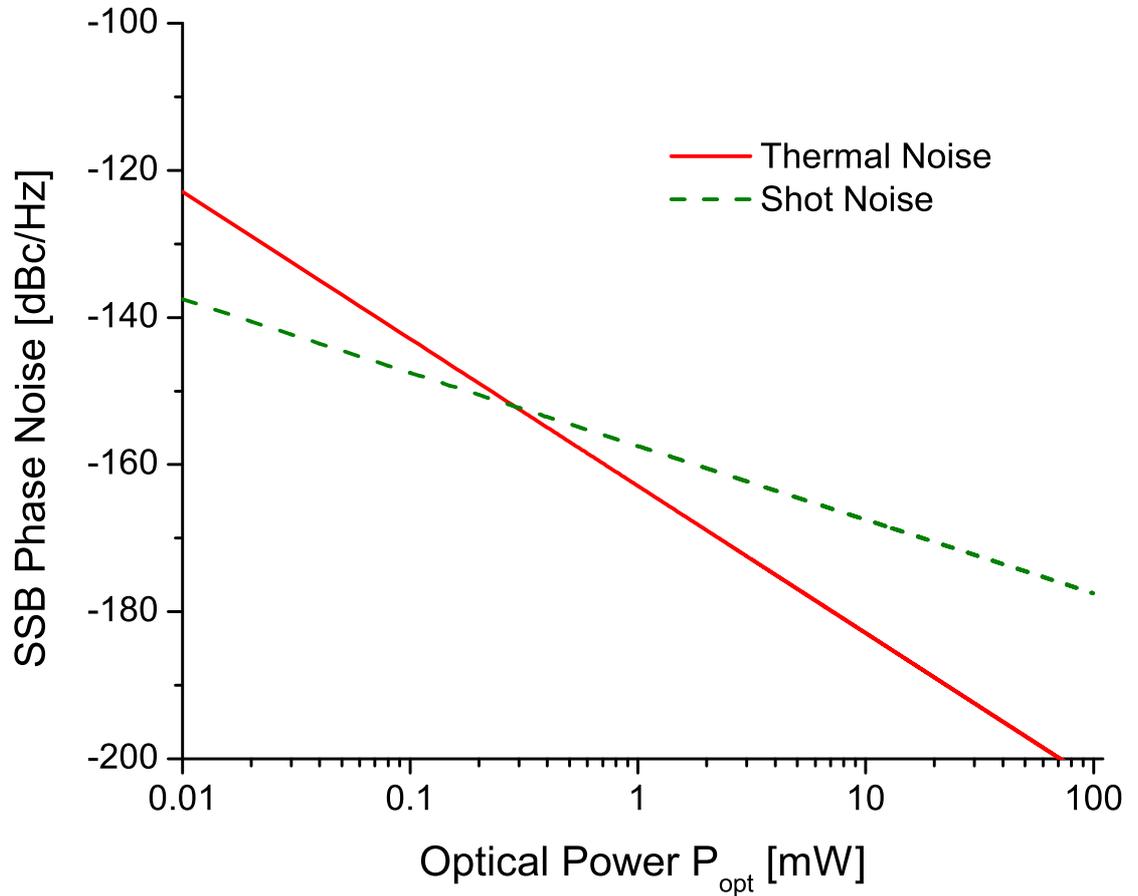


Figure 2.4: Phase Noise floors for thermal and shot noise. This simulation was done with  $\eta = 0.9$ ,  $R = 50\Omega$ , and  $T = 300k$ . Note that the two noise floors cross around 0.5 mW optical power level. Interesting, as the efficiency of the photodetector gets closer to the quantum limit of 100%, the crossing point between these two noise floors moves to lower power. This means that the noise floor for a perfect photodetector would quickly be limited by shot noise.

beam from the cavity and the leakage field from the cavity destructively interfere resulting in the reflected power going to zero. However, when off resonance the immediately reflected beam and the cavity leakage beam don not fully destructively interfere. The typical way to talk about the reflected beam is by using the reflection coefficient  $F(\omega) = E_{ref}/E_{inc}$ , where  $\omega$  is the optical laser frequency. For a Fabry-Perot type cavity composed of two mirrors with reflection amplitude  $r$ , the reflection coefficient as a function laser frequency  $\omega$  is:

$$F(\omega) = \frac{r(e^{i\omega/\Delta\nu_{fsr}} - 1)}{1 - r^2 e^{i\omega/\Delta\nu_{fsr}}} \quad (2.6)$$

Note that if  $\omega$  is an integer multiple of  $\Delta\nu_{fsr}$ , the numerator of eqn. 2.6 vanishes and the reflection coefficient is zero (i.e. there is no reflection and the laser and cavity are said to be impedance matched).

The first way you might think to implement a reflection based locking scheme would be to use the reflected power curve (i.e.  $|F(\omega)|^2$ , see Fig. 2.5) as your error signal. The problem with that, however, is that the curve is symmetric about the resonance. So you would not know from a change in reflected power whether or not the laser frequency had drifted higher or lower. A better option would be to measure the asymmetric phase curve of the reflection coefficient (i.e.  $\tan^{-1}(Im\{F(\omega)\}/Re\{F(\omega)\})$ ), as shown in Fig. 2.6.

In principle, if we could measure this optical phase directly our locking scheme would be complete. However, it is not yet possible (due to electronics speed) to measure an optical phase directly. This leads us to the Pound-Drever-Hall scheme, which is simply a way to indirectly measure the reflection coefficient's phase. The way this scheme accomplishes this is to modulate the incident field's phase (at radio frequencies) to generate sidebands that yield relevant phase information when the sidebands interfere with the carrier beam on a photodetector.

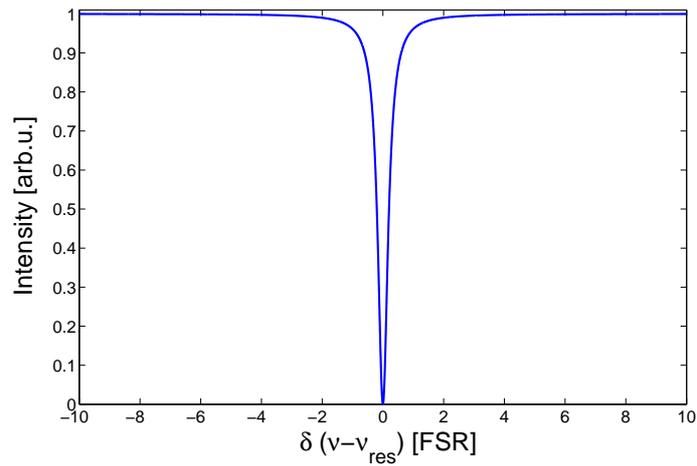


Figure 2.5: Reflected power versus detuning from resonance (in units of cavity FSR) for an optical beam incident on an optical cavity.

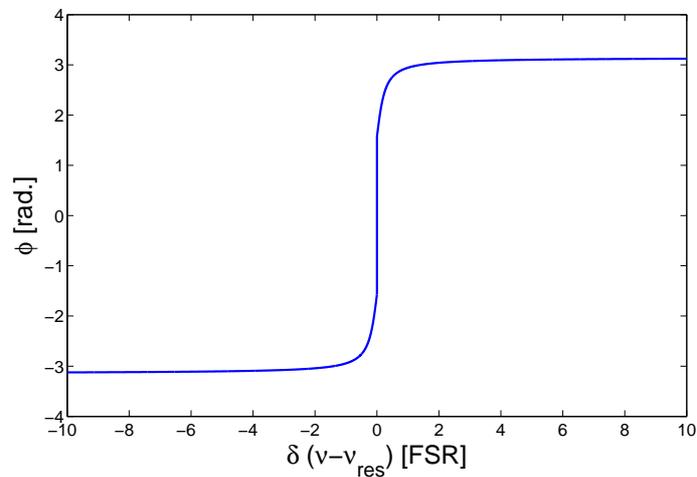


Figure 2.6: Phase of the beam reflected from the optical cavity in a PDH setup versus cavity detuning (in units of the cavity FSR).

In the following section, I will briefly review the mathematics required for a full understanding of how this is accomplished and how the error signal is generated.

The electric field incident on the cavity can be expressed (assuming cw operation) as  $E_{inc} = E_0 e^{i\omega t}$ , where  $\omega$  is the optical carrier frequency. We can use an EOM, Pockels cell, or one of many other options to impart a phase dither on the field incident on the cavity. This phase dither, with a modulation depth  $\beta$ ,

and a frequency  $\Omega$ , moves energy from the carrier frequency to nearby sidebands located at  $\pm \Omega$ . Mathematically, this can be seen by expanding the dithered electric field using Bessel functions:

$$E_{inc} = E_0 e^{i(\omega t + \beta \sin(\Omega t))} \approx J_0(\beta) e^{i\Omega t} + J_1(\beta) e^{i(\omega + \Omega)} - J_1(\beta) e^{i(\omega - \Omega)} \quad (2.7)$$

The first term on the right hand side of eqn. 2.7 is simply the carrier field, the second term is the higher frequency sideband, and the third term is the lower frequency sideband. Two things should be taken away from this equation: (1) the amount of power in the carrier and sidebands depends only on the modulation depth  $\beta$ , and (2) there is a  $180^\circ$  phase shift between the low frequency sideband and the high frequency sideband. For small  $\beta$  (i.e.  $\beta < 1$ ), the carrier and first-order sidebands contain nearly all of the power. This is the regime in which the PDH system operates. The next step is to calculate the error signal based on the fact that these three beams are being reflected off of the cavity and detected at the photodetector. Since the reflection coefficient is a function of frequency, the reflected electric field can be written as  $E_{inc}$  multiplied by the reflection coefficient at the corresponding frequencies:

$$E_{ref} = F(\omega) J_0(\beta) e^{i\Omega t} + F(\omega + \Omega) J_1(\beta) e^{i(\omega + \Omega)} - F(\omega - \Omega) J_1(\beta) e^{i(\omega - \Omega)} \quad (2.8)$$

The reflected signal measured on the photodetector is  $|E_{ref}|^2$ . The terms of  $|E_{ref}|^2$  that oscillate at  $\Omega$  are given by [46]:

$$\epsilon = 2\sqrt{P_{carrier} P_{sb}} [F(\omega) F(\omega + \Omega)^* - F(\omega)^* F(\omega - \Omega)] [\cos(\Omega t) + i \sin(\Omega t)] \quad (2.9)$$

To extract the terms that are oscillating at  $\Omega$  and leave out the myriad other signals present, we first mix the photodetector output with the local oscillator (at frequency  $\Omega$ ), then low pass filter to get the dc error signal. The resulting Pound-Drever-Hall error signal is shown in Fig. 2.7. As can be seen in the figure, the steep

slope of the PDH error signal is caused by the resonance behavior on either side of the central zero point. In fact, for a given modulation depth  $\beta$ , the sensitivity of the PDH interferometer scales as Finesse/FSR. As an example of the error sensitivity of this method, consider this scenario: a 1 mW beam is phase dithered with a modulation depth of  $\beta = 1.08$  (which turns out to be the optimum  $\beta$  for maximizing the sensitivity [47]) before being directed into an optical cavity with Finesse of 500,000 and an FSR of 100 MHz. This system would have a sensitivity of  $13\mu\text{W}/\text{Hz}$ , meaning a 1 Hz change (out of  $10^{15}\text{Hz}$ ) would be easily detected by the electronics! This high level of sensitivity is the reason the PDH lock has become the standard for high precision laser locking [42]. The PDH locking scheme works equally well for mode-locked lasers as is demonstrated in the High-Harmonic-Generation experiment of Ch. 4 of this thesis.

## 2.4 Measuring It

Now that we have introduced several different locking schemes, a natural question to ask is how well is the lock working? In principle, there are many ways to answer this question. The most common method is to analyze the error signal. The simplicity of splitting off a portion of the feedback signal with an rf splitter makes this option very attractive. Another method, commonly employed in optical clock measurements, is to create a heterodyne beat by overlapping the spatial modes of the reference and the source onto a photodiode. Analysis of this beat frequency will reveal the relative frequency stability of the two lasers. In this section I will discuss the fundamentals of noise analysis and illustrate a few experimental examples of a measurement system.

Many tools have been developed to evaluate the stability of a frequency source. In principle, each of these tools answers the same question: how well does the frequency source stay at the desired frequency? The difference between

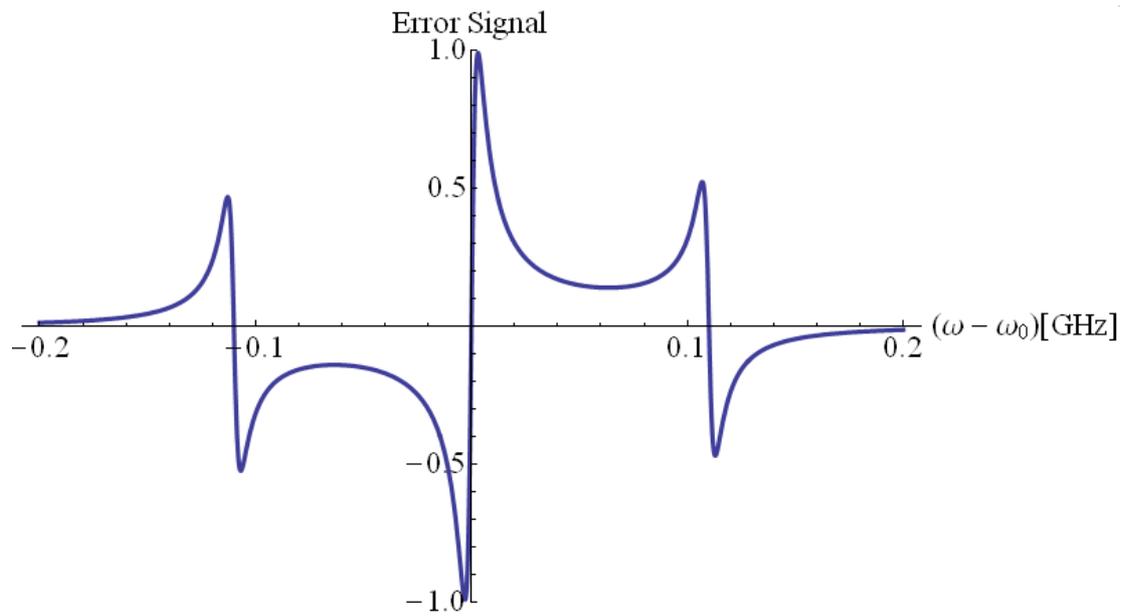


Figure 2.7: Pound-Drever-Hall error signal. The error signal is normalized by dividing by the factor  $\sqrt{2P_c P_s}$ . The modulation frequency  $\Omega$  is 110 MHz, or 37% of the cavity's FSR. A Mathematica notebook for this calculation is given in Appendix A.2.

the various measurement techniques lies in the timescales at which the data is taken. For long time scales (up to hours, days, months, etc.) we have the Allan variance [48], which measures the fractional frequency variation as a function of averaging time. On the other extreme of timescales we have the phase noise measurement technique. This measurement records frequency fluctuations on an extremely fast timescales (limited by the speed of the electronics). Both measurements provide useful characterizations of frequency sources, and their use simply depends on the application. For instance, if a system had large thermal fluctuations that cycled on an hour timescale, the phase noise measurement would not be sensitive to these changes and an Allan variance measurement should be used. However, the phase noise measurement could detect diode current processes that would be much too fast for the Allan variance method.

#### 2.4.1 Phase/Frequency Noise Spectral Density

To begin the discussion of the various techniques, we need to more clearly define what it is that we are measuring. The most simple and complete way of writing a frequency source mathematically is:

$$V_{source}(t) = V_0(t)\cos(2\pi\nu_0t + \phi(t)) \quad (2.10)$$

where  $V_0(t)$  is the amplitude (time dependent when amplitude noise is present),  $\nu_0$  is the natural oscillator frequency, and  $\phi(t)$  is the phase (time dependent when phase/frequency noise is present).

Recall that for the PDH locking scheme, sinusoidal phase modulation of a sine wave leads to sidebands in the frequency domain. A very powerful way to think about noise on a laser is considering the noise as random phase modulation of the carrier frequency [49]. The most intuitive way to quantify this is to simply

take a ratio of the power contained in one of these sidebands ( $P_{ssb}$ ) over the power contained in the carrier ( $P_c$ ), known as the Power Spectral Density (PSD):

$$\frac{P_{ssb}(f)}{P_c} = \left| \frac{E_{ssb}(f)}{E_c} \right|^2 = [J_1(\delta\phi_{peak}(f))]^2 \quad (2.11)$$

$$[J_1(\delta\phi_{peak}(f))]^2 = \frac{\delta\phi_{peak}(f)^2}{4} = \frac{\delta\phi_{rms}(f)^2}{2} \quad (2.12)$$

The only remaining detail to address is the issue of over which bandwidth we measure the power in the sideband. The standard definition is 1 Hz. Thus, the quantity  $P_{ssb}(f)$  is the power contained in a 1 Hz frequency window located at fourier frequency  $f$  away from the carrier. We are now in a position to define the Single-Sideband Phase Noise Spectral Density as:

$$L(f) = \frac{1}{2} \frac{\delta\phi_{rms}(f)^2}{b} \quad (2.13)$$

where  $b$  is the detection bandwidth (1 Hz). The units of  $L(f)$  are thus  $\text{rad.}^2/\text{Hz}$ . Experimentally, the quantity that is often measured is twice  $L(f)$  due to that fact that an FFT machine lumps the positive and negative fourier frequencies together. This quantity is known as the double-sided phase noise spectral density  $S_\phi(f)$ . The SSB-Phase noise is often expressed in dBc/Hz (where dBc refers to dB below the carrier of the sideband in a 1 Hz window). Mathematically, the SSB-Phase noise can be written in terms of the double sided PSD:

$$L_\phi(f) = 10 \log_{10} \left( \frac{1}{2} S_\phi(f) \right) \quad (2.14)$$

So we simply need to measure  $S_\phi(f)$  to get the phase noise, but how do we do that? A typical way to do this is to use a doubly-balanced mixer. As mentioned in the last section, the mixer output is a voltage that depends on the relative phase of its two inputs. A good way to think of the phase noise then is to imagine one of the inputs as a perfect sine wave with no phase errors (i.e. reference). The other

input, however, has some amount of phase error associated with it and thus jitters back and forth. The mixer voltage then oscillates according to the phase differences detected between the reference and the noisy input. For stabilized lasers the phase excursions are small and we can employ the small angle approximation:

$$V(t) = V_0 \sin(\phi(t)) \approx V_0 \delta\phi(t) \quad (2.15)$$

The mixer output,  $V(t)$  is then fed to an FFT-machine that calculates the Fourier transform of the time domain voltage record. The FFT machine records the signal in units of  $V/\sqrt{Hz}$ .

$$V(f) = V_0 \delta\phi_{rms}(f) \Rightarrow \delta\phi_{rms}(f) = \frac{V(f)}{V_0} \left[ \frac{rad.}{\sqrt{Hz}} \right] \quad (2.16)$$

The quantity  $V_0$  is the phase to voltage coefficient of the mixer, which can be measured by recording the voltage change as the phase difference between the two inputs is tuned between 0 and  $\pi$ . Thus, by recording  $V(f)$ , we can construct  $S_\phi(f)$  or  $L_\phi(f)$ .

It is useful here to make a connection to another commonly quoted metric of laser noise, the Frequency Noise Spectral Density (FSD). The basic difference is that we record frequency fluctuations as a function of Fourier frequency rather than phase fluctuations. These two types of fluctuations are related by the equation:

$$2\pi\delta\nu(t)t = \delta\phi(t) \quad (2.17)$$

Taking a time derivative of this equation, performing a Fourier transform, and solving for  $\delta\nu(t)$  yields:

$$\delta\nu_{rms}(f) = f\delta\phi_{rms}(f) \quad (2.18)$$

Thus, the frequency noise spectral density (given by the square of eqn. 2.18) is simply the phase noise spectral density multiplied by  $f^2$ . The units of frequency

noise spectral density are  $\text{Hz}^2/\text{Hz}$ . A commonly quoted quantity for frequency fluctuations is the residual frequency fluctuations defined as:

$$\Delta\nu = \sqrt{\int_{f_l}^{f_h} f^2 \phi_{rms}(f)^2 df} \quad (2.19)$$

### 2.4.2 Timing Jitter

The two noise spectral densities mentioned above each have their place in the lab. Which one you choose to measure is determined partially by the system that you are measuring, the tools available, and the goals of the experiment itself. However, there is yet another noise spectral density we can measure that is suitable in certain experiments known as Timing Jitter Spectral Density (TSD). TSD is essentially a scaled version of PSD (scaled by the carrier frequency). The advantage to quoting noise in terms of TSD is that the final number for timing jitter is independent of the carrier frequency. A phase slip of 1 radian at high frequency represents a much different amount of stability than a phase slip of 1 radian at low frequency. Thus, PSD as a noise metric must always come with information about the carrier frequency. However, TSD represents only the timing fluctuations of the signal which are independent of carrier frequency. This type of noise metric is especially useful for experiments involving synchronization of pulses (see Ch. 3).

Returning to the equation 2.17, we can derive the TSD by letting the time variable  $t$  become itself time dependent:

$$\delta T_{rms}(t) = \frac{\delta \phi_{rms}(t)}{2\pi\nu_0} \quad (2.20)$$

Since experimentally we take the Fourier transform of the mixer output, the quantity we are really measuring is:

$$\delta T_{rms}(f) = \frac{\delta\phi_{rms}(f)}{2\pi\nu_0} \quad (2.21)$$

Using equation 2.16 we can write the TSD in terms of the voltage fluctuations from mixer:

$$\delta T_{rms}(f) = \frac{V(f)}{2\pi\nu_0 V_0} \left[ \frac{fs}{\sqrt{Hz}} \right] \quad (2.22)$$

We can also define the integrated timing jitter as:

$$\Delta T_{rms} = \sqrt{\int_{f_i}^{f_h} [\delta T(f)]^2 df} \quad (2.23)$$

It is important to note that assumption that the modulation depth of the perturbations is small ( $\phi(f) \ll 1$ ) is critical to all of the previous analysis. If this assumption is not true, then we would have to include higher order Bessel terms in the equation for the power spectral density. For reasonably stabilized lasers, however, the peak phase excursions at a given Fourier frequency are generally much less than 1 radian (in a 1 Hz window).

### 2.4.3 In-loop error vs. Out-of-loop error

To wrap up this discussion of error analysis, we need to examine one last detail about the way we measure noise. The most straightforward method is to simply monitor the error signal that is being used for feedback. One must take care, however, in calculating the noise spectral density from this sort of error signal. This so-called *in-loop* error signal(see Fig. 2.8-(a)) can yield a system noise that is artificially low. The reason for this is that with high gain the servo loop can suppress the fluctuations to the noise floor and then lower than the noise floor by writing the negative of the noise floor onto the laser. Thus, the detected phase noise will appear to be lower than the actual noise floor set by the detection

process! A remedy for this unfortunate scenario is provided by the *out-of-loop* measurement. In this setup (see Fig. 2.8-(b)), a separate measurement system is used for the error analysis. This ensures that we are truly measuring the frequency stability of the laser versus the reference since the loop cannot artificially suppress a noise floor it has no information about.

This chapter has introduced the tools and concepts needed to stabilize lasers. In particular, the PLL is an essential tool for stabilizing the timing of pulse trains from a mode-locked laser and will be used extensively in the next chapter. This language of frequency/phase noise is perhaps a bit confusing at first, but with application is much easier to understand. The next chapter will put into action some of these concepts and thus hopefully add a practical dimension of understanding to the concepts covered in this chapter. The reader is encouraged to use this chapter as a reference for any frequency stability questions that may arise during the next few chapters.

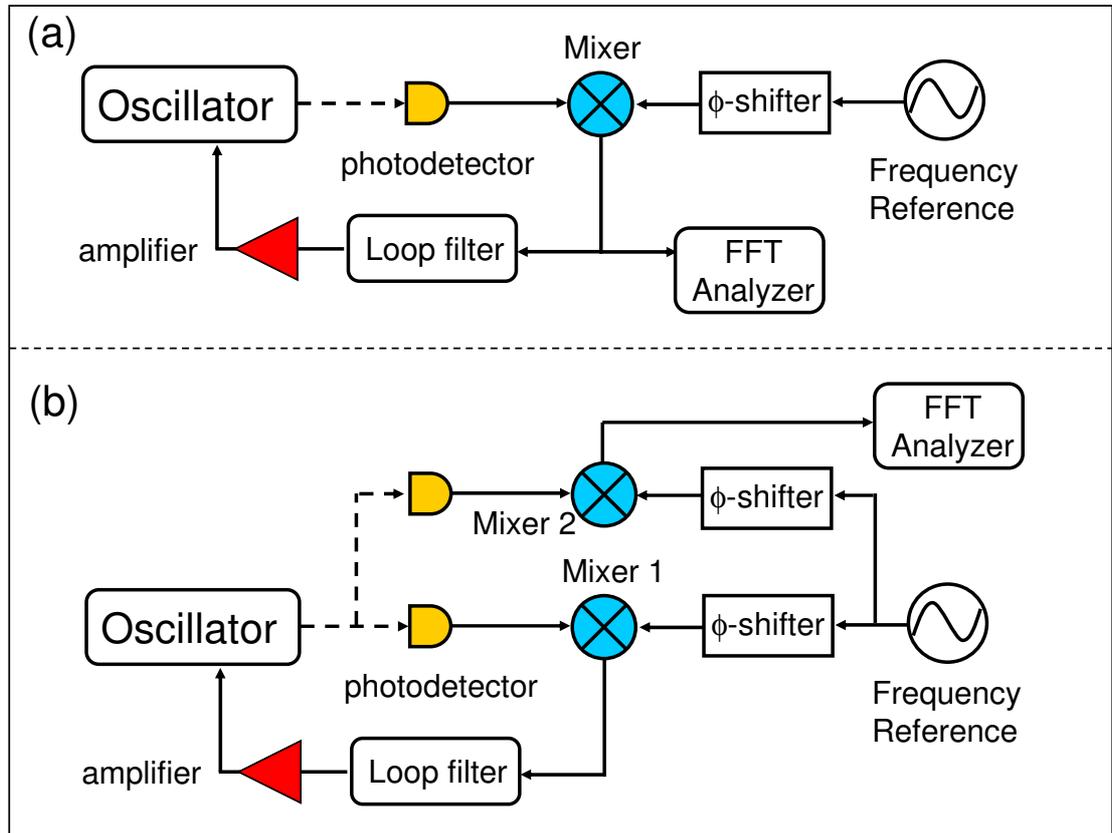


Figure 2.8: (a) In loop error signal measurement. The error signal from the mixer is used for both servoing the laser frequency and for noise analysis. (b) Out of loop error signal measurement. In this setup, the error signal is derived from one set of photodetectors/mixers/amplifiers while the actual noise analysis measurement is done with a completely different set of these components.

## Chapter 3

### Tight Timing Control of a Mode Locked Erbium Fiber Laser

Controlling the repetition frequency and the offset frequency of a mode-locked laser has received much attention since the advent of the frequency comb [50, 51, 52]. Ostensibly this is due to the fact that the precision measurement community (which developed the first frequency combs) was interested in making a precise frequency ruler for clock measurements. There are other applications, however, that require tight control of the mode-locked laser output. In particular, several of these projects benefit from a mode-locked laser with tight timing control of its pulses (i.e.  $f_{rep}$  control).

One such application [53] involves synthetic aperturing in the field of radio astronomy. This technique uses arrays of small telescopes separated by some distance to achieve the same resolution as a telescope with a physical size equal to the entire collection of smaller telescopes (see Fig. 3.1). The mathematical machinery for this sort of experiment is based on the Fourier transform and requires the full amplitude and phase of the signal in order to perform the synthetic aperture transformation. While the data collection for the signal amplitude is straightforward, the phase data requires some level of synchronization between all of the telescopes involved. An optical fiber with an ultra-stable mode-locked laser could present a nice solution for this system. In this scheme, a master laser would send its pulse train through an optical fiber to all of the telescopes. The telescopes

would then lock their local oscillator to the repetition frequency of the master laser (i.e. a microwave frequency that is generated through photodetection). This common phase locking between all telescopes involved results in a phase coherent data scan, allowing for the synthetic aperture processing to work.

Another application involving tight timing synchronization of mode-locked lasers can be found in linear accelerator and x-ray physics [55, 56, 7]. The emergence of linear accelerator based x-ray free electron lasers (XFEL) as tools for pump-probe experiments has led to much interest in tight timing synchronization between laser pulses. In this technology, a conventional laser pulse is used to pump the sample, while the time-delayed x-ray probe is used for interrogating the sample dynamics. Timing jitter between the pump pulse and the probe pulse clearly degrades the temporal resolution of the system [7]. A major source of timing jitter in these experiments is path length fluctuations between the two pulses. Typically, there are large distances involved (i.e.  $\sim 3$  km) in this type of pump-probe due to the nature of linear accelerators (see Fig. 3.2). Thus, path length fluctuations over a few kilometers must be reduced to as low as possible to achieve high time resolution data.

In this chapter I will discuss two experiments that were aimed at transferring a stable optical pulse train over long distances ( $\sim$  kilometers) via an optical fiber and synchronizing a second laser to the transmitted pulse train. The first experiment focuses on creating a laser that can be locked to a stable reference. To achieve extremely precise control of this laser, we introduce a novel intra-cavity actuator with an extremely high actuation bandwidth [6]. The second experiment takes this actuated laser and synchronizes it to a pulse train that is sent from another laser over a kilometer scale optical fiber [8]. Several different optical fibers (installed and spooled fiber) were tested and a system of active stabilization was implemented.

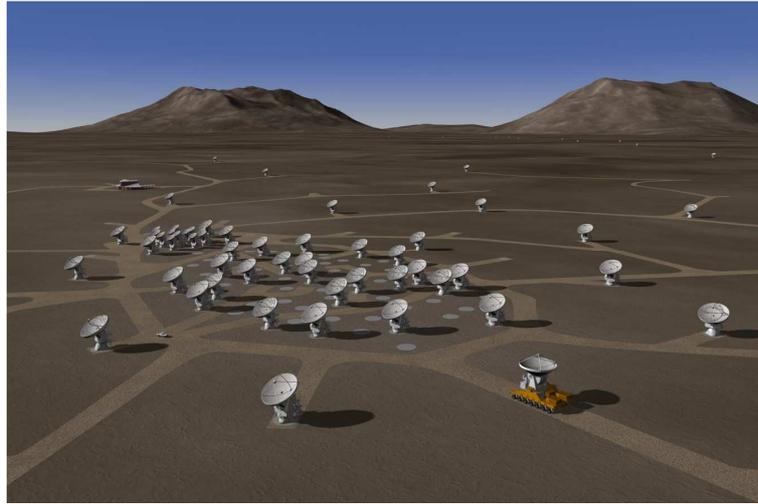


Figure 3.1: Artist's rendering of the Atacama Large Millimeter Array (ALMA) telescope [54] in Chile. Phase coherent data collection, achieved by phase-locking the individual receivers to a master reference, allows for synthesis of a larger aperture than the physical aperture of each telescope. When complete in 2012, this telescope will be configurable to have an aperture of 150 m to 14 km. The end result is a telescope with spatial resolution 10 times better than that of the Very-Large Array and the Hubble Space Telescope.

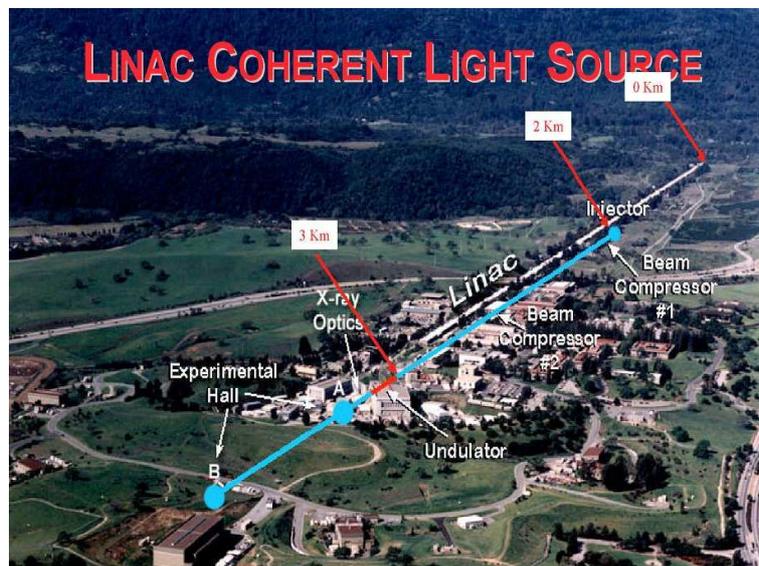


Figure 3.2: Aerial view of the the Linac Coherent Light Source (LCLS) [57] at the Stanford Linear Accelerator. Pump pulses from a conventional laser system must be synchronized to the probe pulses of the XFEL to achieve good timing resolution. The LCLS provides a unique tool to scientists studying a broad range of topics from nanoscience to biology.

### 3.1 Using an EOM for fast feedback control

Early work with frequency combs generated from fiber lasers [5] found a large amount of high-frequency noise on both the pulse repetition rate and the offset frequency. These fluctuations of the repetition frequency are likely due to the high-gain/high-loss condition in the fiber laser cavity, which means the pulses receive a strong spontaneous-emission perturbation on every pass through the erbium-doped fiber [58]. Therefore, applications involving timing synchronization require a broad bandwidth feedback loop to stabilize the repetition rate. Moreover, in frequency metrology it is often necessary to minimize the linewidth of the individual comb components [59]. Accordingly, a high-bandwidth actuator ( $> 100$  kHz) capable of correcting these fluctuations is a key component for femtosecond frequency combs produced by fiber lasers to become useful scientific tools. In the following I will discuss a project involving the use of an electro-optic modulator (EOM) inside a APM-NPR mode-locked laser cavity. This modulator acts as a large bandwidth frequency/phase servo transducer, perfect for stabilizing a noisy mode locked fiber laser.

Due to the long upper-state lifetime ( $\sim 10$  ms) of erbium-doped fiber, direct amplitude modulation of the 980 nm pump diodes is limited in speed. An alternative, non-mechanical actuator is an intracavity EOM [60]. While EOMs can be used to modulate many aspects of light such as amplitude, polarization (through the birefringence of the crystal) and phase, we were only interested in using it to modulate the phase. To achieve such a phase modulation, we aligned the input electric field polarization to be completely along the optical axis of the  $\text{LiNbO}_3$  crystal that was in the direction parallel to the applied electric field (see Fig. 3.3). In this configuration, we took advantage of the Pockel's effect (linear electro-optic effect) of the Lithium Niobate crystal and thus the index of refraction of the crys-

tal along that optical axis was a function of applied electric field. One can think of an index change of the intracavity EOM as a small change in the cavity length, thus producing a small change in the repetition frequency.

While the EOM does introduce some chromatic dispersion into the fiber laser cavity, the change in the overall cavity dispersion is small and can be easily compensated. The combination of a high-bandwidth actuator consisting of an intracavity EOM and a low-bandwidth, long-range actuator consisting of an intracavity piezoelectric- (PZT-) actuated mirror allow for tight stabilization of the fiber laser repetition frequency to a reference over a large dynamic range. There also exists the possibility of controlling not only the repetition frequency but also the offset frequency via an intracavity EOM. In this scenario, the phase index would be set by a DC voltage across the EOM and the group index could be tuned via an AC voltage modulation, at a harmonic of the repetition frequency, on the EOM crystal [61]. The EOM modulation would impart a positive, negative, or zero frequency shift to the pulse, depending on which part of the modulation waveform the pulse encountered. This frequency shift coupled with the dispersive fiber in the cavity would allow for group index control. The fiber laser used in this experiment was based on a standard ring cavity design [62] with a polarizing beam splitter as the output coupler (Fig. 3.4).

The EOM we employed was a 2 cm long, 5 mm thick piece of  $\text{LiNbO}_3$ , with an insertion loss of 2%. The group velocity dispersion (GVD) at  $1.55 \mu\text{m}$  for  $\text{LiNbO}_3$  is  $+100 \text{ fs}^2/\text{mm}$ , thus leading to a group delay dispersion (GDD) of  $+0.002 \text{ ps}^2$ ). The EOM dispersion was compensated by adding a 9 cm long fused silica rod (GVD =  $-22 \text{ fs}^2/\text{mm}$ ) inside the fiber laser cavity. Overall, the laser cavity had a net anomalous group-delay dispersion ( $-0.008 \text{ ps}^2$ ), thus operating in the soliton regime. This laser proved to be quite difficult to mode-lock. Most likely this was due to the many free space components in the cavity which each

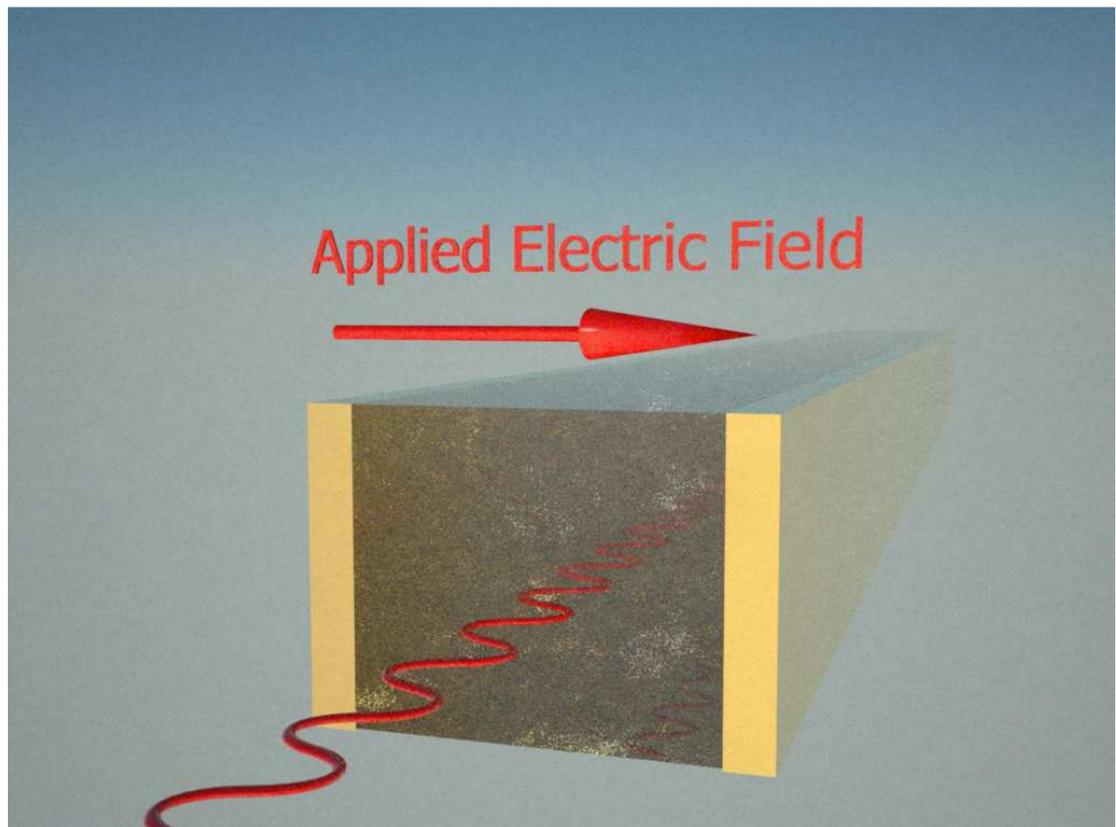


Figure 3.3: Electro-optic modulator. Electrodes form a parallel plate capacitor around the EOM. The input polarization is matched to the direction of the applied electric field of the electrodes so that the linear electro-optic effect only effects the phase of the input light.

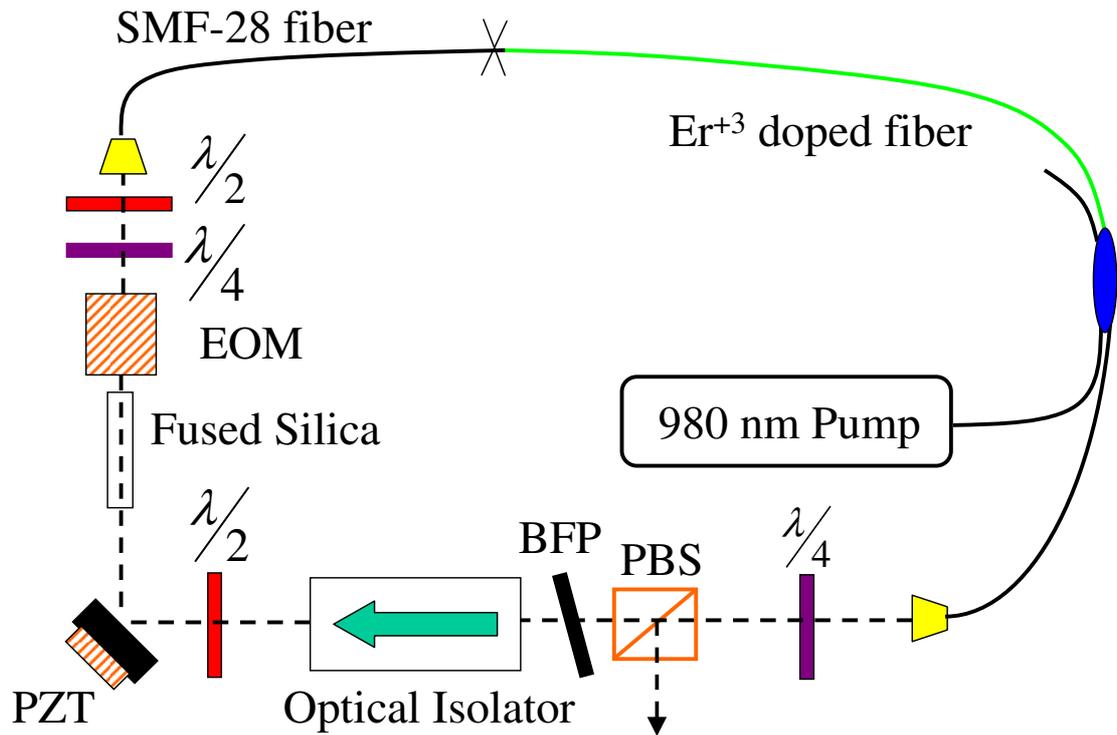


Figure 3.4: Schematic of the fiber laser with intracavity EOM and intracavity PZT. The output coupler is the polarizing beam splitter (PBS). A half-wave plate after the optical isolator allows for polarization adjustment into the e-wave axis of the EOM. BFP, birefringent tuning plate; SMF, single-mode fiber.

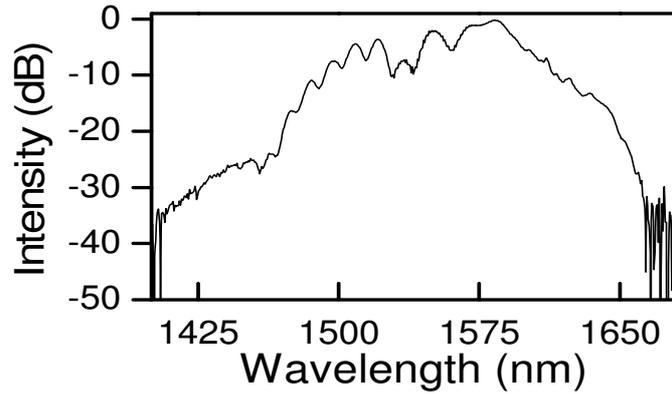


Figure 3.5: Optical spectrum of the EOM fiber laser pulses.

presented some loss. It is likely that the finesse of this cavity was very low ( $\sim 1-10$ ). Nevertheless, we were able to achieve mode-locking and in fact found that the laser produced relatively clean pulses with pulsewidths around 170 fs (measured via autocorrelation and assuming a hyperbolic secant pulse shape). The spectrum for this laser is shown in Fig. 3.5.

### 3.2 Characterizing the performance of the EOM

We performed various tests on the EOM to understand its action on the frequency comb. In terms of a servo actuator, two key tests are the dynamic range of the EOM and its servo bandwidth. To measure the dynamic range, we applied various voltage steps across the two EOM electrodes and measured the responses of the repetition frequency. To precisely monitor these changes, we stabilized the laser to an RF reference using a low-bandwidth 150 Hz, low-gain PZT lock. This procedure allowed the DC value of the repetition frequency to be stable while the fast changes in the repetition frequency induced by the EOM were monitored through the in-loop servo error signal, which is not affected by the slow PZT servo (Fig. 3.6). Measuring the change in the phase of the repetition frequency signal with time after applying the voltage step allowed us to calculate the shift

of the repetition frequency. The EOM driver provided a maximum voltage step of 500 V to the EOM. As seen in the inset of Fig. 3.6, within this voltage range the frequency change is linear and the maximum frequency change is  $\sim 1$  kHz out of the 80th harmonic of the fundamental repetition rate of 93 MHz. This maximum variation is equivalent to a total cavity length change of  $0.28\mu\text{m}$ . The theoretical cavity length change is given by  $\Delta l = nl = r_{22}Vn^3/2d$ , where  $r_{22}$  is the only nonzero component of the electro-optic tensor for  $\text{LiNbO}_2$ ,  $V$  is the voltage applied across the EOM,  $n$  is the index of refraction of  $\text{LiNbO}_2$ ,  $l$  is the EOM length (20 mm) and  $d$  is the distance between the electrodes on the EOM. With an applied voltage of 500 V this equation gives a cavity length change of  $0.37\mu\text{m}$ , in reasonable agreement with that extrapolated from the step response measurement.

The EOM actuation bandwidth is determined via the measurement of its transfer function as shown in Fig. 3.7. Again, the laser was weakly stabilized to a reference via a low-bandwidth 150 Hz, low-gain PZT lock. We monitored the error signal from the phase-locked loop in the Fourier frequency domain using an FFT machine while a swept sine wave was applied to the EOM electrodes. For the input signal frequency below 150 Hz, the transfer function is suppressed artificially due to the low bandwidth PZT lock that was being applied. From the plot we see that the -3 dB roll-off point of the EOM response is around 230 kHz, while the phase lag reaches  $90^\circ$  at 200 kHz. It is important to note that this speed of actuation is extremely hard to reach via a moving actuator (i.e. a PZT actuated mirror).

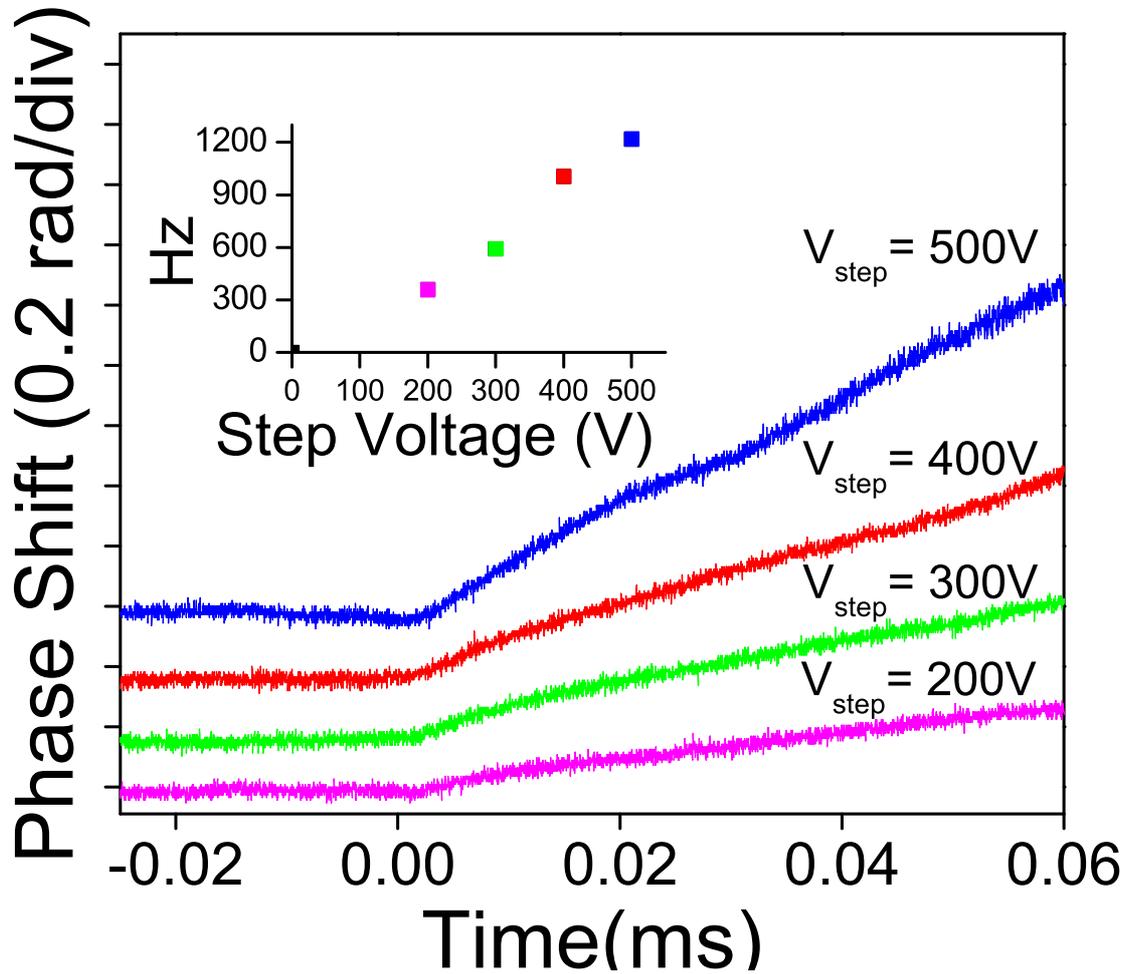


Figure 3.6: Step response of the EOM measured at increasing input voltages. The ordinate axis shows the phase change of the 80th harmonic of the repetition frequency signal. The voltage step turns on at 0 ms on the plot and stays on for 10 ms. The repetition frequency change of the laser reaches a limit of around 1 kHz (at the 80th harmonic) at a maximum of 500 V across the EOM. The inset shows the change of the 80th harmonic of the repetition frequency of the laser with increasing voltage across the EOM.

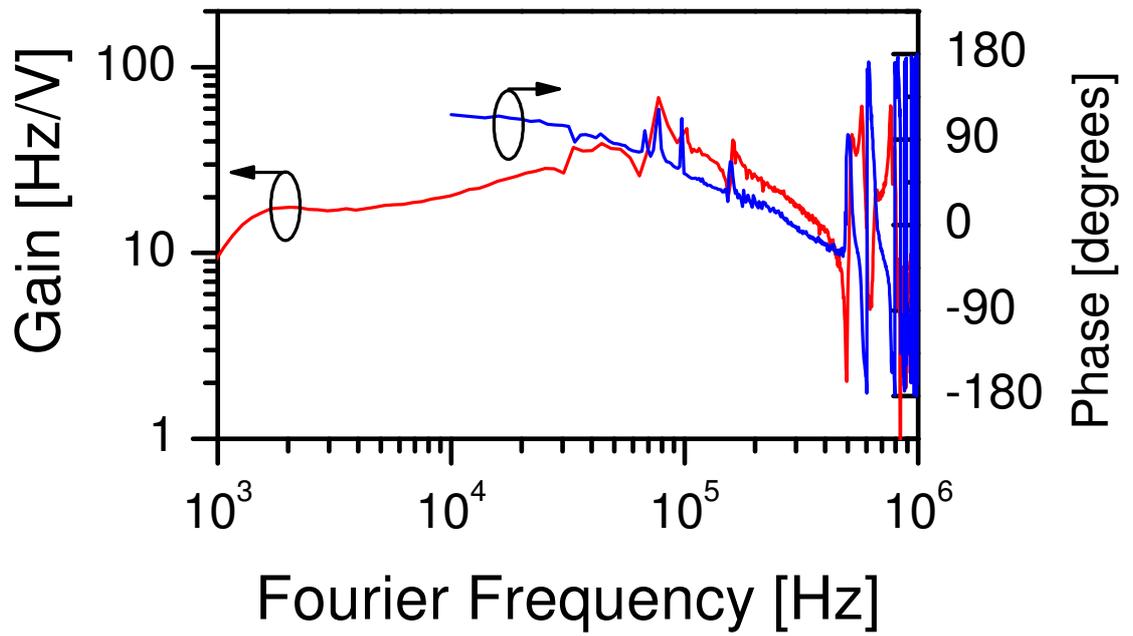


Figure 3.7: Transfer function of the EOM. The -3 dB roll-off frequency is approximately 230 kHz, while the phase lag reaches  $90^\circ$  at approximately 200 kHz.

### 3.3 Synchronizing the EOM laser to a free-running mode-locked fiber laser

To investigate the practical application of the intracavity EOM we used it to lock the repetition frequency of the fiber laser to a second, independent mode-locked fiber laser with a fundamental repetition frequency of around 31 MHz. We built this second laser with a long-travel translation stage that allowed length tuning of the free-space section of the laser cavity, thus allowing for large relative changes in  $f_{rep}$ . Separate photodetectors were used to detect the 80th repetition frequency harmonic of the 93 MHz laser and the 240th harmonic of the 31 MHz laser, which were then phase-sensitively compared. The error signal was then filtered and fed back to the actuators in the slave (EOM) laser cavity. Spectral analysis of the locking was accomplished by use of a fast Fourier-transform spectrum analyzer. The in-loop jitter spectral densities for the free-running case and locking with the PZT and the EOM are shown in Fig. 3.8-(a).

The EOM and PZT together reduce the integrated jitter (over a bandwidth from 1 Hz to 100 kHz) from approximately 1800 to 10 fs [Fig. 3.8-(b)]. The measurement noise floor was determined by feeding an identical signal into both arms of the mixer, with appropriate phase shift and amplitude adjustment. The integrated jitter of the noise floor from 1 Hz to 100 kHz is less than 5 fs. Figure 3.8-(a) shows that the locking is limited by the noise floor, which is dominated by the RF amplifiers, below 60 Hz. At frequencies above 100 kHz the noise spectrum is also limited by the RF amplifiers. A realistic estimate of the jitter at a larger bandwidth can be performed by rolling off the noise floor to match the extrapolated free-running roll-off at high frequencies, due to the lack of servo actions there. This method yields 20 fs of integrated jitter for a 10 MHz upper limit and 21 fs of jitter for a 50 MHz (Nyquist frequency) upper limit.

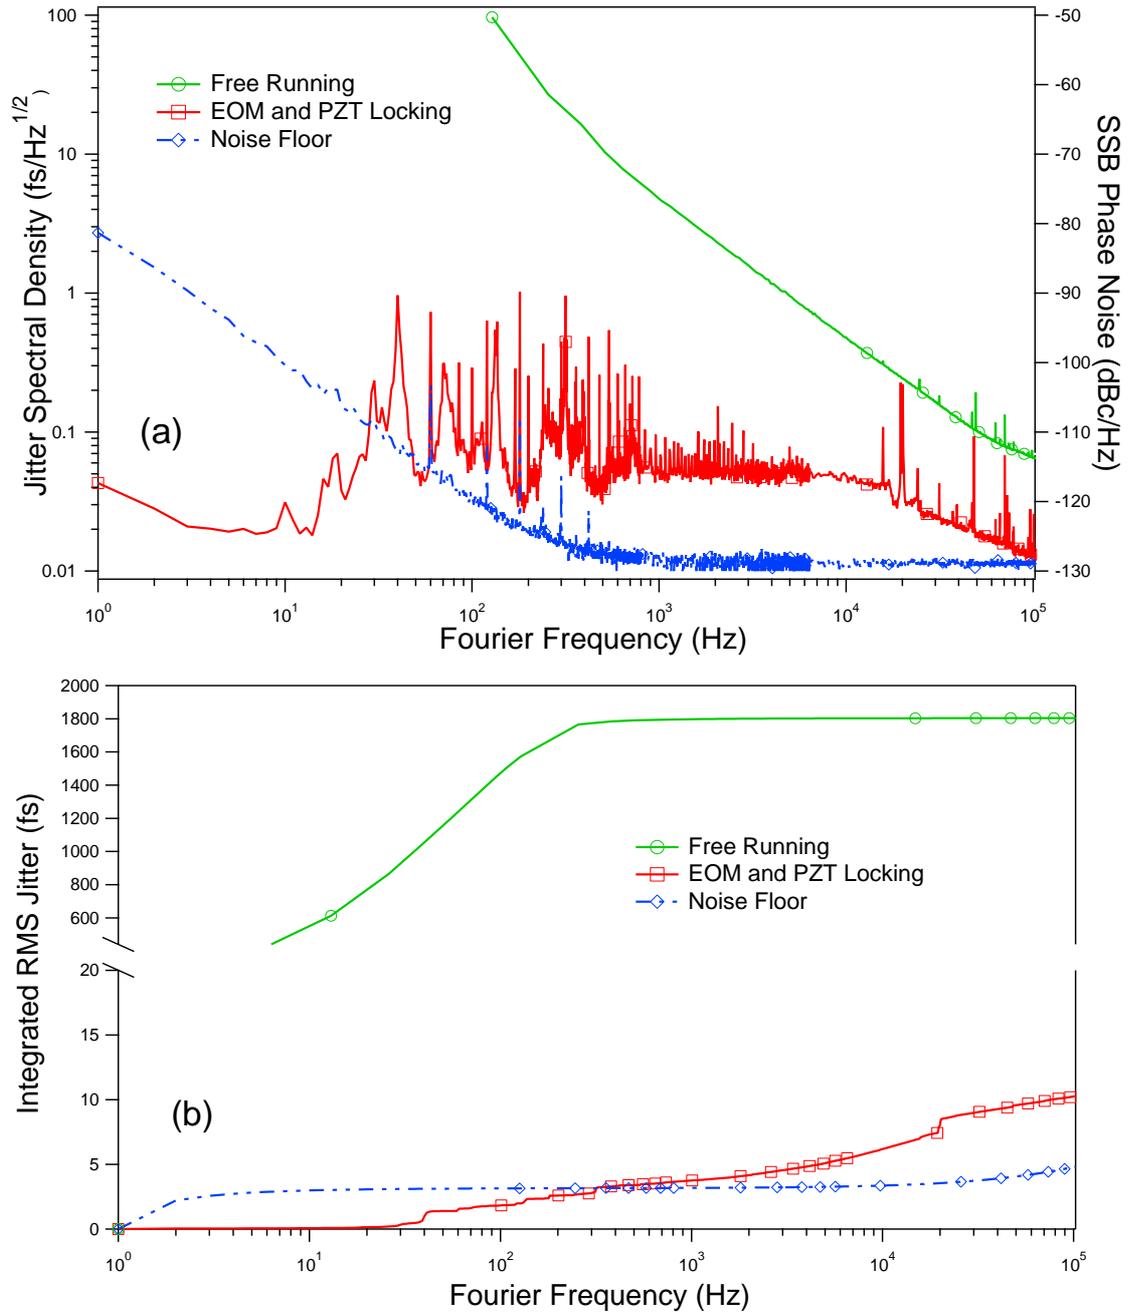


Figure 3.8: (a) Jitter spectral density (left axis) and single-sideband phase noise (right axis) for locking the 80th harmonic of the repetition frequency 7.5 GHz of the laser to the second fiber laser. (b) The free-running rms jitter is around 1800 fs. With the intracavity EOM used in conjunction with the PZT, the integrated rms jitter is reduced to 10 fs. The noise floor of the measurement is limited by the thermal noise, which has an integrated rms jitter over 1 Hz to 100 kHz of less than 5 fs.

The EOM loop bandwidth is ultimately limited by resonances above 500 kHz. These resonances are most likely due to piezo-like electromechanical resonances in the EOM crystal. To minimize this sort of phenomenon, we mounted the EOM in a Teflon casing with vibration-absorbing material surrounding the crystal. However, at high driving frequencies the mechanical resonances are not completely suppressed. These issues could perhaps be resolved by replacing the free-space EOM with an inline fiber EOM.

### 3.4 Conclusion for EOM Locking

In this experiment we showed that synchronization of a mode-locked fiber laser to a reference via an intracavity EOM can lead to very tight locking. A record-low level of 10 fs of integrated jitter over a bandwidth of 1 Hz to 100 kHz was achieved. This is a reduction in the timing jitter of 2 orders of magnitude between the free-running laser and the EOM stabilized laser. In the next section, I will discuss an experiment to synchronize this fiber laser to another, remotely located, fiber laser.

### 3.5 Remote Synchronization

Some applications of synchronized lasers, such as coherent pulse synthesis [63] and precise pump-probe experiments, [64] use collocated lasers with direct optical links of only a few meters. However, for many applications, such as those mentioned in the introduction to this chapter, it is necessary to have remotely located lasers separated by distances ranging from tens of meters to several kilometers [65]. Transmitting optical pulses over kilometer-scale distances places strict limits on the optical system; the most feasible way is using telecommunications grade optical fibers, which have low loss 0.25 dB/km at wavelengths around 1550 nm. Erbium-doped fiber lasers, which emit light centered at 1550 nm, are thus em-

ployed to utilize this transmission medium. Optical fibers, however, are sensitive to environmental perturbations such as acoustic vibrations, thermal fluctuations, and mechanical stresses. The resultant optical path length fluctuations introduce timing jitter on the pulse train, which must be canceled if a remotely located laser is to be synchronized to the incoming pulse train. In the following experiment we demonstrate that active cancellation of this fiber transmission path noise and a large feedback bandwidth allow for tight synchronization of two fiber lasers over kilometer-scale transmission distances. We present an out-of-loop time-domain analysis of the timing jitter via optical cross correlation, which allows sensitive measurement of the jitter and verifies the frequency-domain in-loop measurements.

The two fiber lasers used in this experiment are both erbium-doped, ring cavity design lasers with a nonlinear polarization rotation scheme as the mode-locking mechanism [62]. The free-running reference laser, which had a repetition frequency of 31 MHz and operated in the normal dispersion regime (i.e. dispersion managed soliton regime), was enclosed in an acrylic box that was mounted on an optical breadboard. The fiber in the laser cavity was environmentally isolated by securing the fiber to leadbacked foam inside the enclosure. The slave laser, which had a repetition frequency of 93 MHz and operated in the anomalous dispersion regime (i.e. soliton regime), had two crucial intracavity actuators, an electro-optic modulator (EOM), and a piezoelectric transducer (PZT). As was demonstrated in the last section, the intracavity EOM had a servo bandwidth of over 200 kHz, allowing for local synchronization with a timing jitter of 10 fs. The intracavity PZT has a long dynamic range of 14  $\mu\text{m}$ , which allowed for locking times of greater than 12 hours. The next element of synchronization is the fiber link between the two lasers. We used two different fiber links to test our method of actively canceling the group delay noise introduced on the pulse train by the fiber: an installed 6.9 km fiber in the Boulder area (known as BRAN [66]) and a 4.5

km spool of dispersion-shifted fiber (DSF). The timing jitter introduced by either of these two fiber links is canceled (to the same degree for either link) by a PZT fiber stretcher that has a large dynamic range [67, 68]. We used these fiber links to simulate the timing jitter that the pulse train would experience in a real-world implementation of this system. In this demonstration we use the configuration shown in Fig. 3.9, where we cancel the round-trip jitter of the link instead of the one way jitter. However, for transmission to a physically different location, a portion of the light at the remote end must be retro-reflected and detected at the local end to derive the timing jitter information of the fiber link [69]. A truly remote system incorporating retro-reflection with a fiber link half as long as ours is expected to perform at the same level as the measurements that we report here.

These two elements-fiber transfer and synchronization-are combined to achieve the synchronization of the two lasers over a kilometer-scale fiber link. Fig. 3.9 shows a high-level diagram of the experiment. Active noise cancellation is used to deliver a highly stable pulse train over the fiber link. Once this is accomplished, the intracavity actuators in the slave laser synchronize its repetition frequency to that of the incoming pulse train.

We characterize the residual timing jitter of the synchronization by using a crossed-beam, background-free, optical cross correlation of the two lasers pulse trains. The two pulse trains are focused onto a Lithium Iodate ( $\text{LiIO}_3$ ) crystal (type-I phase matching), which generates sum frequency light (SFG) when the two pulses overlap in time and space. To achieve temporal overlap, we use two phase-locked loops that operate at two different timing resolutions [50]. A coarse-timing loop operates at the fundamental frequency of 93 MHz, while a higher-resolution loop operates at 7.6 GHz (80th harmonic of the fundamental). A phase shifter in the fundamental frequency loop allows coarse-timing adjustments such that temporal overlap between the two pulse trains can be found. Since the slave

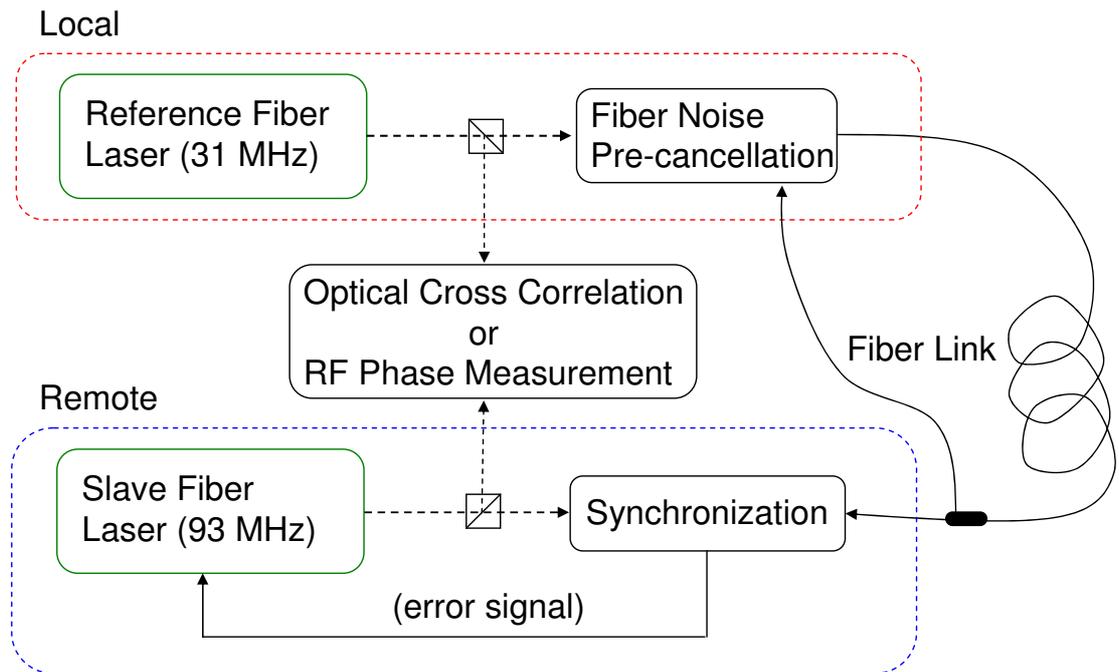


Figure 3.9: Schematic of the transfer and local synchronization setup. The synchronization loop compares the incoming pulse train, which is stabilized via a fiber noise cancellation loop, with the slave lasers pulse train; the resulting error signal is fed back to the intracavity EOM and PZT of the slave laser. The cross correlation is performed locally.

laser's repetition frequency is three times that of the reference laser, only every third pulse from the slave laser overlaps a pulse from the reference laser. Once an SFG signal is observed on a photomultiplier tube (PMT), we measure the total cross-correlation width to calibrate the data. We then transfer control from the fundamental frequency loop to the 7.6 GHz loop. A phase shifter in the high-harmonic loop allows us to finely tune the time overlap of the pulses to position the SFG signal at the steepest point of the cross-correlation slope to obtain the most sensitive measurement of the pulse timing jitter, which is proportional to the amplitude fluctuations of the SFG signal. The fluctuations are monitored through a 50 MHz low-pass filter to determine the timing jitter within an integration bandwidth up to the Nyquist frequency.

It is important to note that the cross correlation is performed on the same optical table that holds the two lasers used in the experiment. This configuration allows a direct comparison of the two lasers, which reveals the timing jitter due to both the transmission path and the slave lasers locking ability. This measurement verifies our synchronization capability over the fiber network. For a successful implementation of true remote synchronization, the only change is to detect the retro-reflected light through the transmission path, as discussed above. To characterize the performance of the servo loops for synchronization over a fiber link, we first analyze the Fourier frequency spectrum of the error signal of the phase-locked loops. The first step is to analyze the timing jitter between the two lasers when synchronized locally (i.e., without a kilometer-scale fiber link connecting the two). The residual in-loop timing jitter between the two pulse trains is measured by way of the residual phase noise fluctuation between the two repetition frequency signals in the Fourier frequency spectrum, which can be converted into a timing jitter spectral density. Fig. 3.10-(a) shows the jitter spectral density out to 100 kHz and a corresponding integrated jitter of 10 fs.

Next, we characterized the jitter of the fiber-link transfer of the pulse train from the reference laser. The error signal in this case is derived from mixing the local rf signal with the rf signal from a photodetector that detects the transmitted light; this error signal is then fed back to the PZT fiber stretcher, which has a resonance around 18 kHz. This wide bandwidth actuator yields an improvement over previous measurements [68]; we achieved approximately 16 fs of timing jitter for either the BRAN or DSF fiber, integrated over 1 Hz to 100 kHz, as can be seen in Fig. 3.10-(b). Importantly, we measure an out-of-loop error signal (as discussed in Ch. 2) by using independent photodetectors for the combined transfer and synchronization. This is shown in Fig. 3.10-(c), with a total jitter of 19 fs, integrated over 1 Hz to 100 kHz.

Next, we improved the jitter measurement by using an out-of-loop optical cross correlation between the two pulse trains. The goal here is to overlap pulses from the two lasers in time and space on a nonlinear crystal. If we overlap the pulses at their steepest slopes (see Fig. 3.11, then any timing jitter between the two pulses will result in an amplitude change. These amplitude changes are recorded on a photo-multiplier tube (PMT), low pass filtered, and then digitized on an oscilloscope. The voltage width of the timing jitter signal can be converted into timing jitter by noting how much time delay is required to move the PMT signal from trough to peak. This approach is advantageous, since it provides a highly sensitive detection of timing jitter without electronic noise contributions, and it also provides an independent assessment of the system performance out of the servo loop. The local synchronization result measured by the optical cross correlation is shown in Fig. 3.12-(a).

The conversion coefficient of the timing jitter (phase noise) to the amplitude fluctuations of the SFG signal is pre-calibrated, and thus, we can determine the rms timing jitter directly from the amplitude fluctuations of the SFG signal.

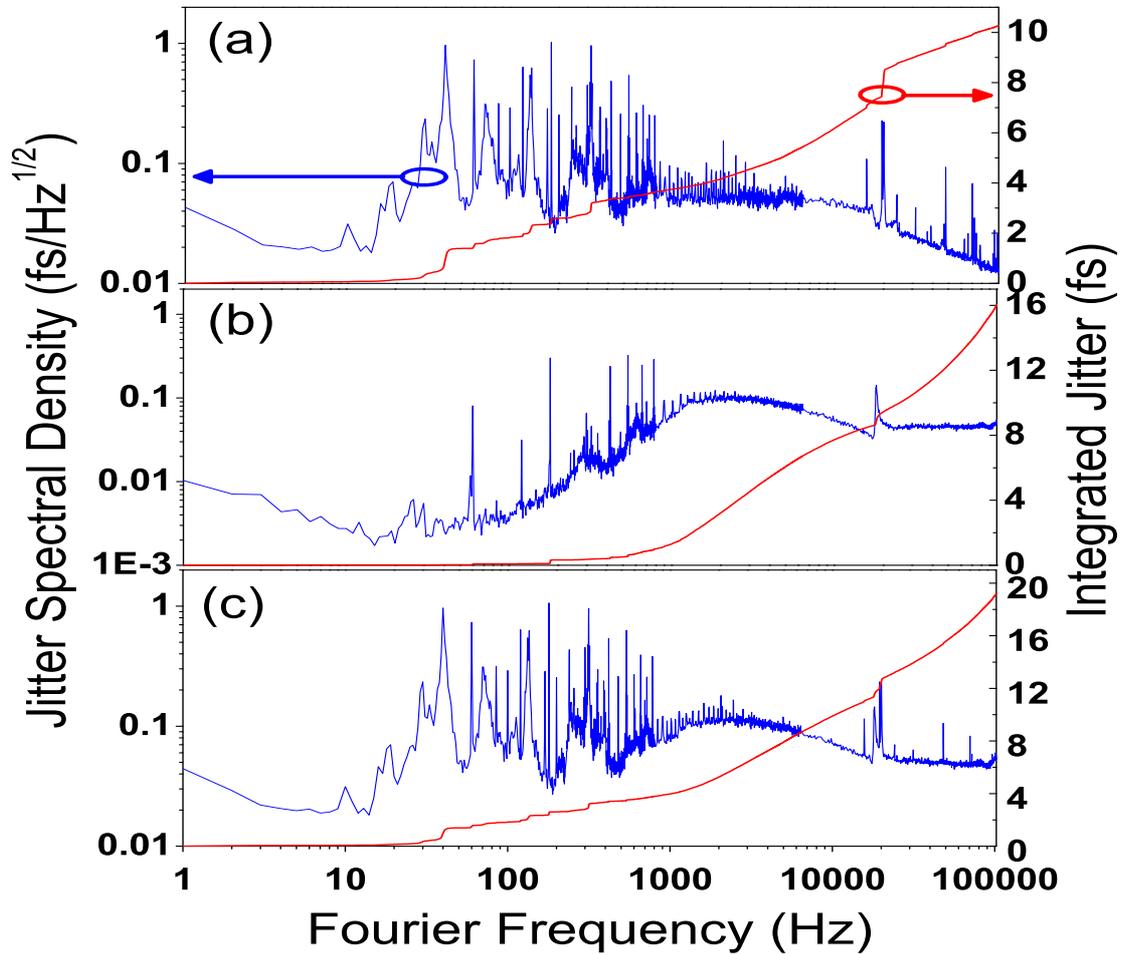


Figure 3.10: In-loop error signal for (a) local synchronization, (b) fiber noise cancellation, and out-of-loop error signal for (c) long-distance synchronization. All plots are shown over a 1 Hz to 100 kHz integration bandwidth. The fiber noise cancellation loop has 16 fs of timing jitter, the local synchronization has 10 fs of timing jitter, and the long-distance synchronization has a timing jitter of 19 fs.

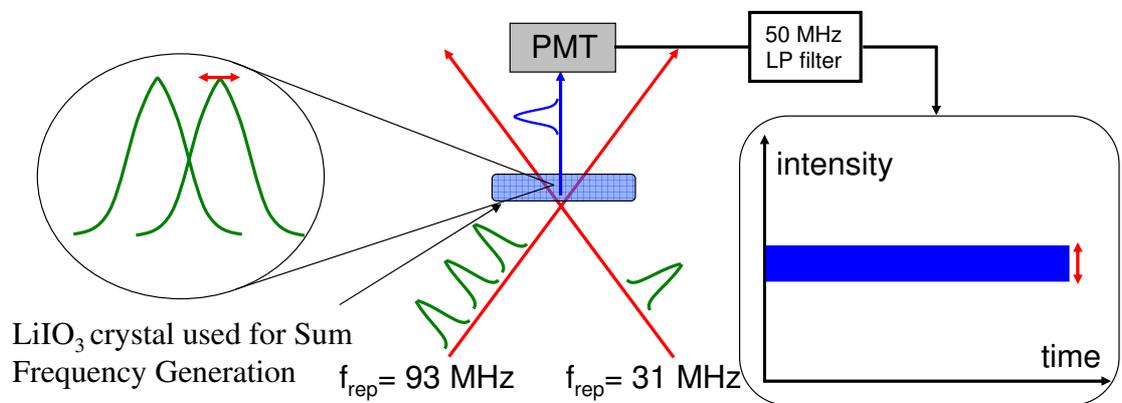


Figure 3.11: Schematic of the cross correlation layout. The pulses from the two lasers overlap in space on the nonlinear crystal and in time at their steepest points. A timing jitter between the two pulses results in an amplitude change on the photomultiplier tube. This signal is then recorded by an oscilloscope.

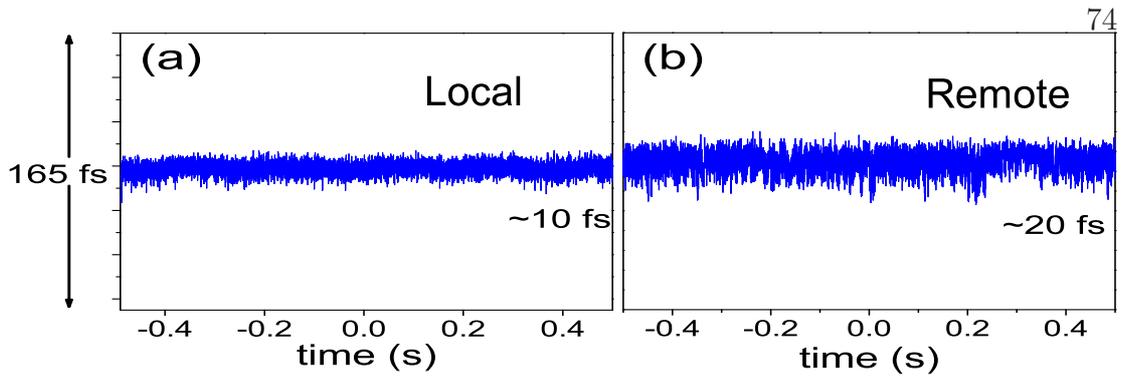


Figure 3.12: Optical cross-correlation measurement of the timing jitter for (a) local synchronization, and (b) long-distance synchronization. Both traces show the rms timing jitter within the Nyquist bandwidth (50 MHz).

This measurement is taken in a bandwidth equal to the Nyquist frequency of 50 MHz, which ensures that it is an accurate representation of all of the noise on the repetition frequency synchronization. We note that the direct jitter measurement from the optical cross correlation agrees with that from the indirect approach of integrating the residual phase noise of the error signal in the feedback loop over a 100 kHz bandwidth, indicating that there is no significant noise contribution beyond 100 kHz. Finally, we characterize the timing jitter of the synchronization over the fiber link (DSF) by using the optical cross correlator. The cross correlation between the two lasers reveals 19 fs of timing jitter observed over the Nyquist frequency bandwidth as can be seen in Fig. 3.12-(b), in agreement with the result shown in Fig. 3.10-(c). The total timing jitter is essentially the root-square sum of the residual jitters from the local synchronization and the fiber transfer.

### 3.6 Cross-correlator as a timing jitter detector and ultimate limits

In experiments in which the reference laser was multiple pulsing, the cross-correlation data showed much higher timing jitter, even though the in-loop error signal analysis indicated tight synchronization. This observation highlights the importance of using optical cross correlation and the danger of relying on the in-

loop error signal. For cases in which the laser had several extra pulses (i.e., more than four pulses), the timing jitter under the tight lock condition was so severe that the cross-correlation signal explored its full range of 165 fs. By selecting an appropriate polarization state of the laser, however, we were able to eliminate the multiple pulsing and achieve the best timing jitter results.

The timing jitter suppression is currently limited by the bandwidth of the actuators in the system: the PZT fiber stretcher, the intracavity EOM, and the intracavity PZT. However, the fundamental limits imposed by Johnson (thermal) noise on the rf amplifiers and shot noise on the photocurrent are quickly being approached. Both of these effects contribute white phase noise that scales inversely with the rf signals power. The synchronization over the fiber link employs the weakest optical power owing to losses in the fiber after several kilometers of transmission, which leads to  $P_{signal} = -48$  dBm for the rf signal from the photodetector, while the photodetector in the local system has sufficient incident light to generate -30 dBm in the 7.6 GHz carrier. For this carrier frequency at an rf power level of -48 dBm, Johnson and shot noise yield phase noise floors of -129 and -132 dBc/Hz. Integrating these noise floors over the slave lasers actuator bandwidth of 200 kHz yields timing jitters of 4.7 and 3.3 fs, respectively. Employing fast photodetectors that can receive larger optical powers will lower these fundamental noise limits. In summary, we have demonstrated synchronization through a 7 km fiber link of two femtosecond fiber lasers at a timing jitter level of 19 fs over the Nyquist bandwidth. This measurement was performed via an optical cross correlation, which provides the most sensitive measurement of timing jitter. We achieved this result by combining the fiber transfer of a reference laser with the synchronization of a slave laser using a fast intracavity actuator.

### 3.7 Putting it into context

The remote synchronization experiment resulted in 19 fs of timing jitter from 1 Hz to the Nyquist frequency. A natural question to ask at this point is how good is that? One can answer this question two ways. The first way is to consider a mental picture of the pulses. For the active stabilization of the fiber link, we achieved 10 fs of integrated jitter. After propagating through several kilometers of fiber, the optical pulses are stretched to pulsewidths of many picoseconds. If you could look at the pulse with your eyes, you would see a 10,000 fs long electric field jittering back and forth with an rms value of 10 fs (so you would be looking for a 0.2% change in the pulse position). The other way to view this level of stability is in regard to what other experiments have reported. The integrated timing jitter for the system shown in Fig. 3.1 was reported to be at a level of 34 fs over a bandwidth of 3 kHz to 3 MHz [53], while the integrated jitter for the LCLS experiment (see Fig. 3.2) was reported to be at 200 fs over a 1 Hz to 40 MHz bandwidth [7]. Thus, our system of high bandwidth actuation via an intra-cavity EOM in conjunction with actively stabilized fiber links has lead to quite remarkable levels of long distance synchronization.

## Chapter 4

### Applications of mode-locked fiber lasers

In this chapter I will discuss two atomic physics experiments involving mode-locked fiber lasers and passive enhancement cavities. The first experiment deals with sensitively and simultaneously detecting the presence of a wide range of molecules in trace amounts. The applications for this are many: ranging from early medical disease detection [70, 71] to security screening for combustible materials. The laser used in this experiment was a home built Erbium doped mode-locked fiber laser, with several actuators for controlling the optical frequency. A chirped-pulse amplifier (CPA) was also constructed so that the pulses could be launched into a highly nonlinear fiber (HNLF) with enough intensity to undergo spectral broadening. The second experiment I will discuss examined using a passive enhancement cavity in conjunction with a noble gas to generate harmonics of the incident frequency comb. The main goal in this experiment is to shift a 1  $\mu\text{m}$  frequency comb to shorter wavelength for high resolution spectroscopy in the UV range. This experiment used an Ytterbium doped fiber laser (developed by IMRA America) in conjunction with a multiple-clad amplifier fiber to deliver 90 fs pulses ( $\sim 136$  MHz repetition frequency) with an average power of 10 W. These applications of rare-earth doped fiber lasers (Erbium and Ytterbium) give us a glimpse of the potential of this class of laser.

## 4.1 Coupling a Frequency Comb to a High Finesse Cavity

The common theme of this chapter is coupling a frequency comb to an optical cavity. This section will review the basic ideas of this process. An optical cavity is a device that consists of at least two mirrors (Fabry-Perot) arranged in a configuration such that the photons that enter the cavity are reflected around the cavity and end up back at the position in which they entered. If the cavity is carefully constructed, the electric field of the entering photons and the field of the recycled photons can interfere constructively, leading to an enhancement of the electric field. Clearly, a high finesse cavity in which the photons are recycled many times before leaking out, offers the possibility of achieving a large electric field enhancement. Another consequence of a high finesse cavity is the enhanced light-matter interaction that can take place due to the many roundtrips of the photons. The reader should already begin to see the power of such a system.

As seen in Fig. 4.1-(a), the resonant modes of an optical cavity are equally spaced by the free-spectral-range (FSR). For a simple Fabry-Perot cavity of length  $L$  the FSR is equal to  $c/2L$ . In the frequency domain this means that only light at certain frequencies (the resonances) will be transmitted into the cavity. In Fig. 4.1-(b), the frequency spectrum of a frequency comb is shown. If the two degrees of freedom of the comb,  $f_{rep}$  and  $f_o$ , are controlled to match the transmission spectrum of the cavity, then the light of the frequency comb will be effectively coupled to the cavity. This condition means that a coherent accumulation of pulses will occur in the cavity. However, dispersion in the cavity mirrors causes the FSR to be wavelength dependent (see Fig. 4.1-(c)). For a large bandwidth such as that produced by a frequency comb, a wavelength dependent FSR can cause the coupling bandwidth to be reduced. Therefore it is necessary to use carefully designed mirrors that achieve low chromatic dispersion over a large bandwidth

(preferably as large as the frequency comb bandwidth).

It is useful to understand the coherent accumulation effect in both the frequency and the time domain. The frequency domain, as explained above, is quite straightforward in the sense that you simply want to match the frequency spectra of the laser and the cavity. What does this mean in the time domain? Clearly, you need to match the repetition frequency so that the entering and recycled pulses overlap in time (i.e. envelope overlap). This is not enough, however, to achieve constructive interference. We also need the actual phase of the electric fields to match (see Fig. 4.2 and 4.3). To achieve this we need to be able to control the carrier envelope offset phase so that the recycled and entering electric fields have exactly the same electric field shape at the point at which they interfere.

## 4.2 Molecular Detection Experiment

Since its introduction in 1988 by O’Keefe and Deacon [72] Cavity Ring-down Spectroscopy (CRDS) has become a powerful tool for trace detection of molecules. Detection sensitivities of up to a few parts-per-trillion by volume have been reported [73]. At the heart of a cavity ringdown system is a high finesse cavity with dielectric mirrors that exhibit extremely high reflectivity. The high reflectivity allows photons in the cavity to complete many round trips before they leak out of the cavity. In fact, with common parameters a photon can travel an effective length of several kilometers ( $\sim 10 \mu\text{s}$  storage time), greatly enhancing the light-matter interaction. The nature of the measurement also means that it is insensitive to intensity noise on the laser source. A basic layout for cavity ringdown spectroscopy is shown in Fig. 4.4.

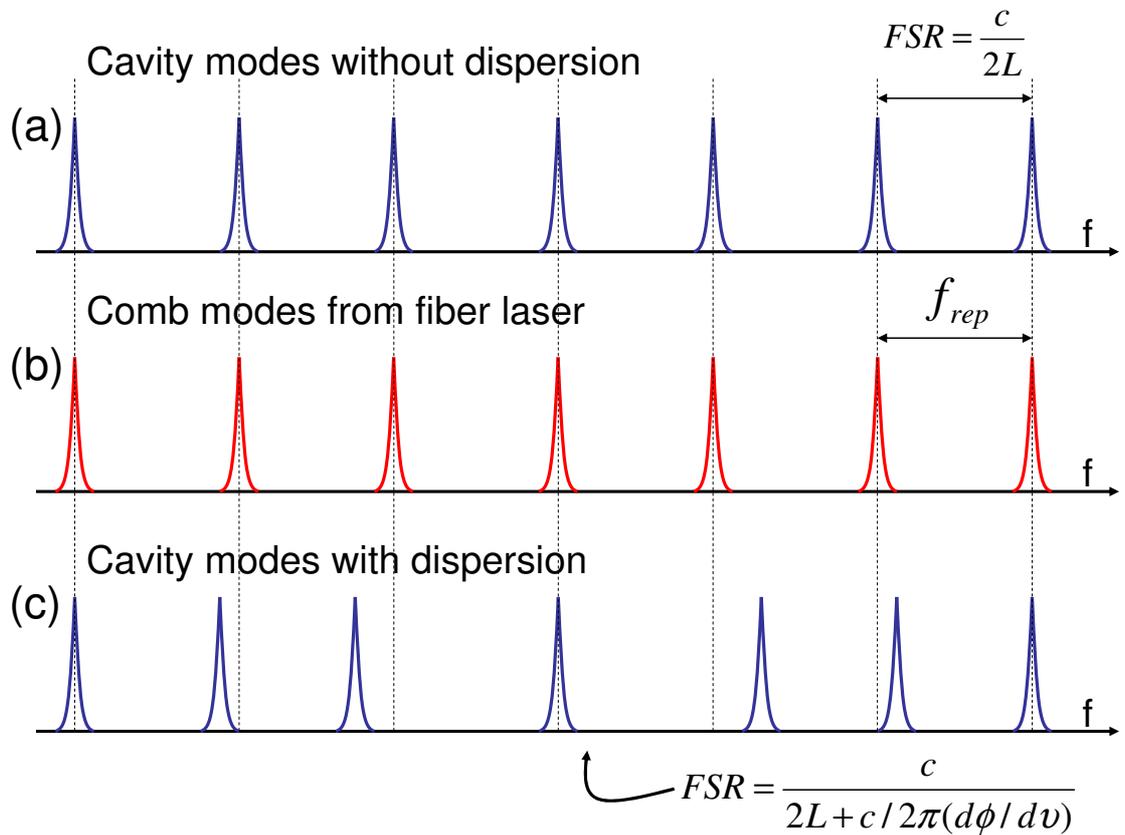


Figure 4.1: (a) Passive cavity frequency modes spaced by the FSR of the cavity. (b) Frequency comb modes spaced by the repetition rate,  $f_{rep}$ , of the mode-locked laser. (c) Passive cavity frequency modes with cavity dispersion. The equation for FSR then becomes dependent on the spectral phase ( $d\phi/d\nu$ ) as shown. The effect of this dispersion is to shift the resonance frequencies of the cavity by non-uniform amounts.

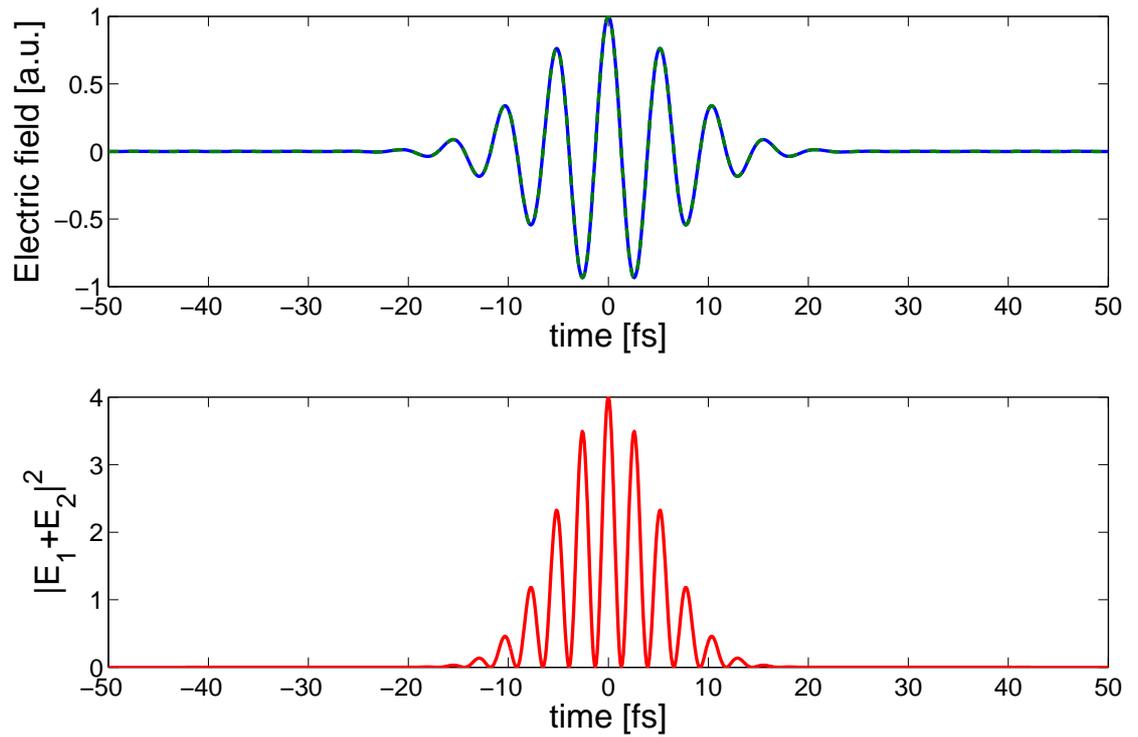


Figure 4.2: Interference of two pulses in a passive optical cavity at the input coupler. The incident pulse (blue line) has the same phase as the round trip pulse (dotted green line), which leads to constructive interference (shown in the lower pane).

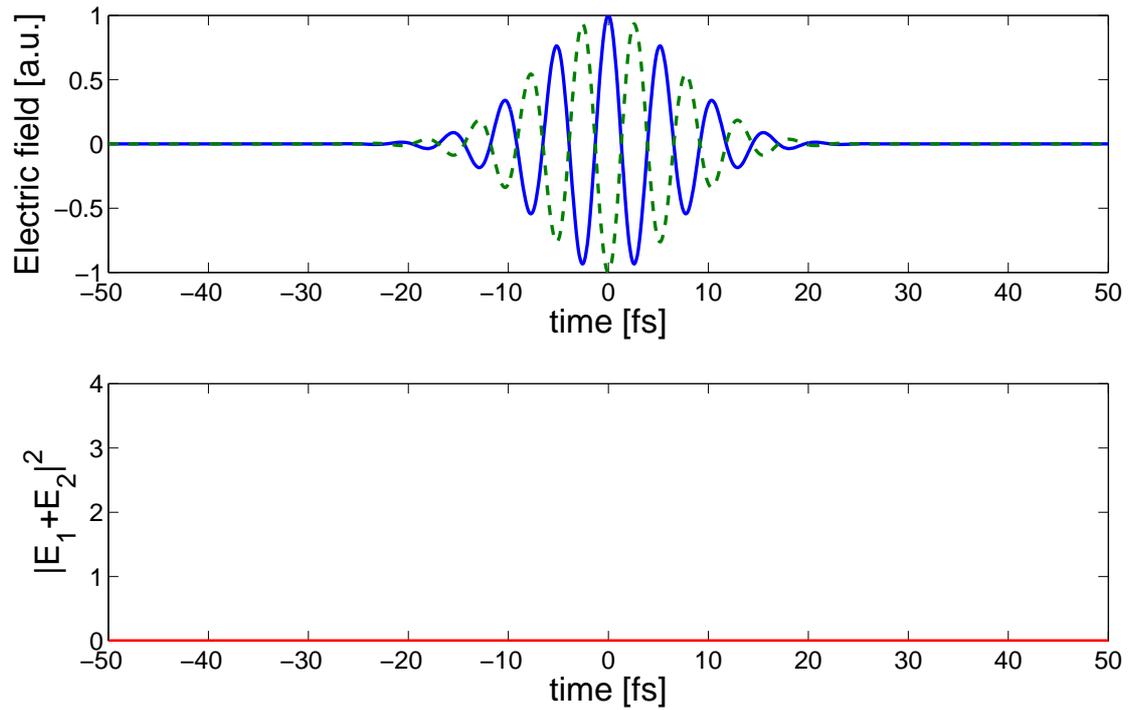


Figure 4.3: Interference of two pulses in a passive optical cavity at the input coupler. The incident pulse (blue line) has a phase difference of  $\pi$  relative to the round trip pulse (dotted green line), which leads to destructive interference (shown in the lower pane).

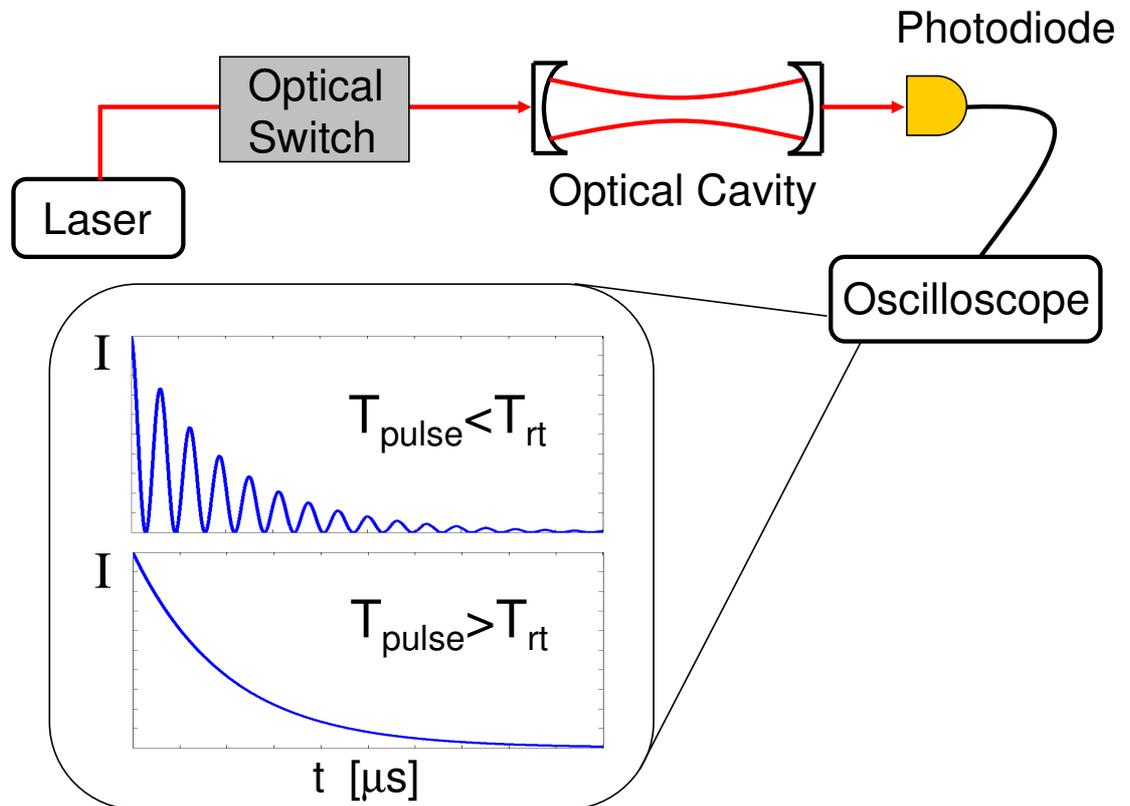


Figure 4.4: Cavity ringdown spectroscopy setup. The optical switch shuts off the light to the cavity, then the ringdown signal is measured as a function of time. If the pulse width is less than the round-trip time, the output shows an exponentially decaying oscillatory structure due to the pulse sampling the output mirror every round trip. If the pulse width is longer than the round trip time, the oscillations are blurred and the signal is just the exponential decay.

### 4.2.1 Brief explanation of CRDS

In a standard CRDS experiment, laser light is coupled to a cavity and then shut off with an optical switch. The ringdown signal is simply the time record of the intensity at the output mirror. If the pulse is shorter in time than the roundtrip time of the cavity  $T_{rt}$ , the detector at the output will see intensity spikes every time the pulse bounces off of the output mirror. Due to mirror losses ( $L_0 = 2(1 - R)$  for a Fabry-Perot cavity), though, the overall intensity detected at the detector will exponentially decay as a function of time. If the pulse is longer in time than the cavity, the ringdown signal will decay exponentially but the fast *carrier* signal will be washed out. Mathematically we can define the ringdown time as the time it takes the field to reach a level of  $1/e$  of the maximum transmitted field. For the empty-cavity case this ringdown time is:

$$\tau_0 = \frac{T_{rt}}{L_0}. \quad (4.1)$$

However, if we introduce some absorbers (i.e. atoms or molecules) into the cavity we must add a new loss term in our equations for loss through absorption. This can be done by adding a term in the total loss as follows:

$$L_{tot} = L_0 + \alpha l_c \quad (4.2)$$

where  $l_c$  is the cavity length and  $\alpha$  is the absorption coefficient at a particular wavelength. This, in turn, leads to a new ringdown equation for the case where absorbers are present in the cavity:

$$\tau = \frac{T_{rt}}{L}. \quad (4.3)$$

The total loss equation can be solved for the absorbance (i.e.  $\alpha l_s$ ) to yield the absorbance as a function of measured ringdown times:

$$\alpha l_s = (1 - R) \frac{\tau_0 - \tau}{\tau} \quad (4.4)$$

Thus, experimentally you only need to measure the empty cavity ringdown time ( $\tau_0$ ) and the absorber ringdown time ( $\tau$ ) to determine the absorbance for your system. Note that the ultimate limit of the sensitivity is given by how well you can time resolve the ringdown signal and the reflectivity of the mirrors. This is quite reasonable since a higher reflectivity means the light-matter interaction length is longer and a better time resolution means we can determine the 1/e time more accurately.

#### 4.2.2 Broadband CRDS

In this experiment we performed a version of CRDS known as Broadband Cavity Ringdown Spectroscopy (BCRDS) [37]. In this setup, a mode-locked laser is coupled to a cavity (as was discussed in the first section of this chapter). This configuration allows for the simultaneous detection of cavity ringdown events at many different wavelengths. Thus, the use of the mode locked laser helps overcome one of the main limitations of conventional cw-CRDS. Obviously, the key component in such a system is the laser itself. For the purposes of this experiment, the laser had to meet two important requirements: emit light over a large spectral bandwidth and be able to control both degrees of freedom of the comb. To achieve the first goal, we employed a stretched pulse cavity design (i.e. net normal dispersion) operating with a repetition frequency of 100 MHz. This type of laser can emit light over a larger spectral bandwidth than a soliton laser (net anomalous dispersion) due to the absence of the Kelly sidebands [74], which clamp the oscillator spectrum. To address the issue of frequency control in this laser, we added three intra-cavity actuators. For long-range frequency tuning (required

to keep the laser on the cavity resonance as the cavity is evacuated) we used a picomotor with a dynamic range of  $\approx 1/100^{th}$  of the cavity length, yielding a frequency tuning of 1.5 MHz. To achieve faster control of  $f_{rep}$ , two piezo electric transducers were employed: a long range (17  $\mu\text{m}$  dynamic range) PZT with a relatively slow response and an inline fiber PZT for fast modulation ( $\approx 20$  kHz). The carrier-enveloped-offset frequency was controlled via pump current. The spectrum directly from the oscillator was measured to be around 100 nm wide. While this is quite large for an Erbium doped system, we employed an amplifier and a section of Highly Nonlinear Fiber (HNLF) to broaden the spectrum even further such that we were limited by the bandwidth of the mirrors (see Fig. 4.5).

As can be seen in Fig. 4.5-(b), the dispersion of the mirrors effectively cuts off the comb and cavity coupling at the points 1450nm on the short wavelength side and 1650 nm on the long wavelength side. Thus, we have a detection bandwidth of 200 nm, limited by the mirror dispersion.

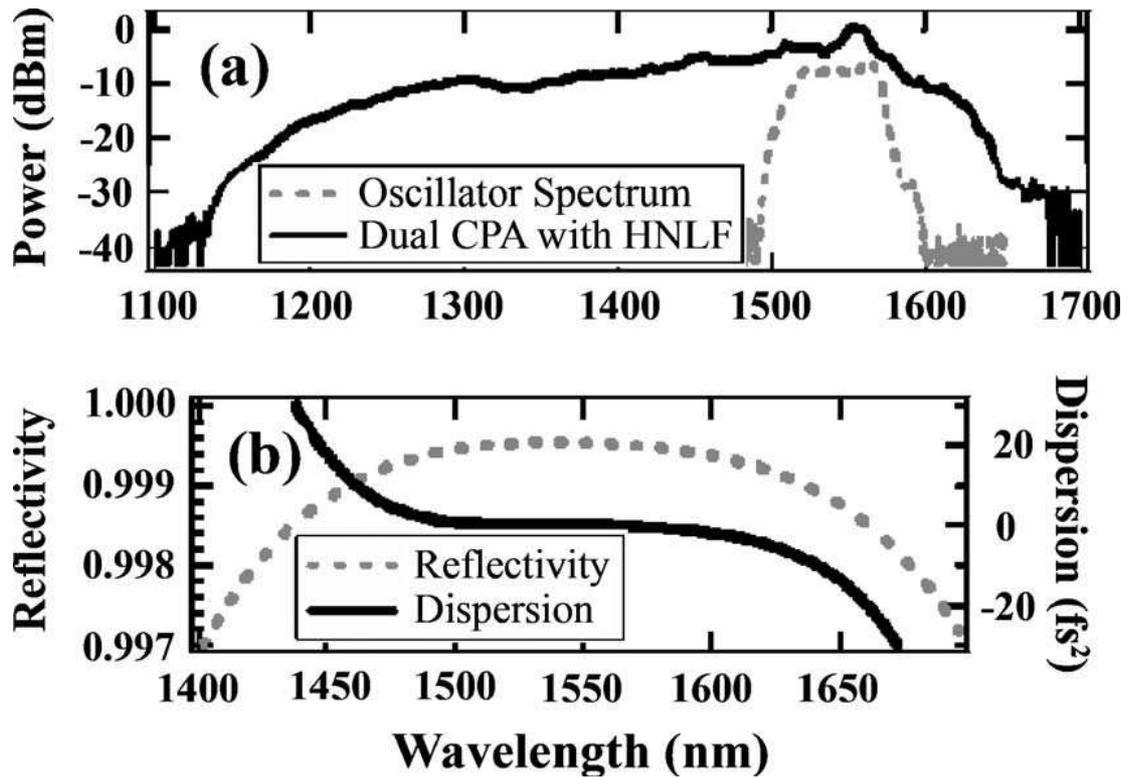


Figure 4.5: (a) Spectral power density in (dBm/nm) for the fiber laser before and after amplification and spectral broadening. (b) Dispersion and reflectivity curves of the cavity mirrors show that low dispersion  $< 15 \text{ fs}^2$  and high reflectivity ( $R > 0.999$ ) are maintained between 1450 and 1650 nm.

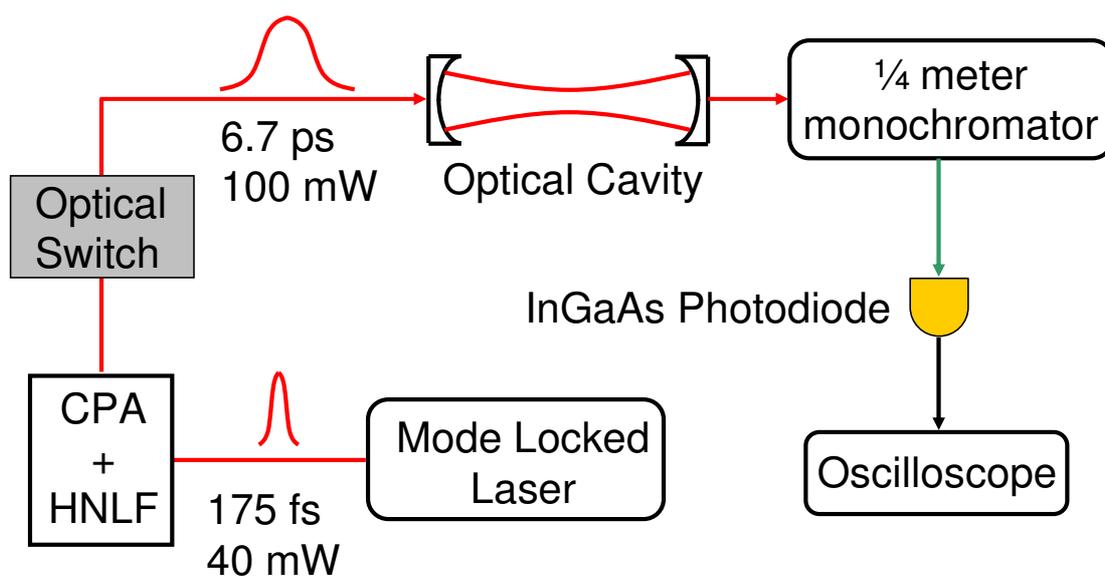


Figure 4.6: Broadband CRDS setup. The mode locked laser is coupled to the cavity using mode matching optics (not shown). The optical switch turns off the light and the monochromator directs individual wavelengths to the photodiode for recording of the ringdown signal. In this experiment, the grating inside the monochromator was scanned and the ringdown event was measured at each individual wavelength.

The final experimental layout is shown in Fig. 4.6. The monochromator allowed for ringdown measurements at each wavelength, with a frequency resolution of 25 GHz. This frequency resolution has since been improved using various other technologies included a virtually imaged phased array (VIPA) [75]. We used the setup to measure overtone spectra of three different molecules: Carbon Monoxide (CO), Methane (C<sub>2</sub>H<sub>2</sub>), and Ammonia (NH<sub>3</sub>). The measured absorption spectra for each of these gases is given in Fig. 4.7. By averaging the ringdown signals for 1 second, we achieved a signal-to-noise ratio of 1000, which yielded an absorption sensitivity of  $2 \times 10^{-8} \text{ cm}^{-1} \text{ Hz}^{-1/2}$ . With this sensitivity we were able to detect CO at 6 parts per million, NH<sub>3</sub> at 20 parts per billion volume (ppbv), and C<sub>2</sub>H<sub>2</sub> at 8 ppbv.

### 4.3 High Harmonic Generation: Moving the frequency comb to the UV and beyond

In this section I will discuss coupling a mode-locked fiber laser to a passive enhancement cavity for high-field physics purposes (i.e. intensities  $\sim 10^{12}$ - $10^{16} \text{ W/cm}^2$ ). This type of experiment allows for observation of some very interesting phenomena: High-order Harmonic Generation (HHG) of frequency combs [76, 77], attosecond pulse generation [78, 79], and even imaging molecular orbitals [80]. The following experiment focused on producing HHG for its possible application to spectroscopy. In particular, the experiment aimed at extending the spectral coverage of frequency combs to vacuum ultra-violet (VUV) and extreme ultra-violet (XUV) wavelengths.

Electron energy level spacings for light atoms such as Hydrogen and Helium are typically on the order of many eV. Accordingly, to do spectroscopy on these atoms one needs a light source in the VUV and XUV range. Also, to calibrate the light source for high resolution spectroscopy one needs a frequency link back to the

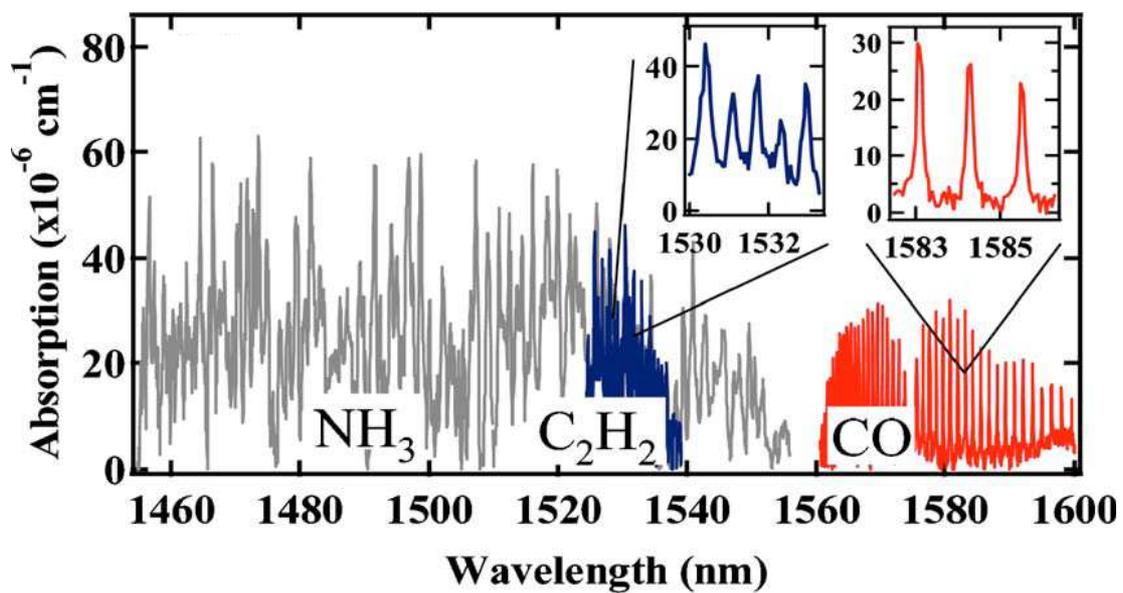


Figure 4.7: Absorption spectra for 2 Torr  $\text{CO}$ , 10 mTorr  $\text{NH}_3$ , and 1.5 mTorr  $\text{C}_2\text{H}_2$ , showing 150 nm of spectral information.

standard definition of the second (Cesium standard). Thus, a UV frequency comb presents a great way to do calibrated spectroscopy of the deeply bound electron quantum states of light atoms. One could in principle generate such a UV comb by using a standard nonlinear crystal ( $\chi_2 \neq 0$ ) to perform sum-frequency-generation (SFG) which shifts the comb modes to half the wavelength of the fundamental comb. However, to reach even the low-energy side of the VUV region one would need to start at a wavelength of 400 nm (the beginning of the UV region). Clearly, this approach is not feasible. A better way to move the comb to short wavelengths is to use a medium in which high-order harmonics are generated [81].

The HHG process can be understood as a 3 step process: (1) a strong laser field drives a deeply bound electron from an atom to the continuum via tunneling enhancement, (2) the free electron is accelerated in the laser field picking up kinetic energy, and (3) the electron recombines with the ion emitting a single photon with a much higher energy than the energy of the fundamental photons. This radiation, which is in the same direction as the driving field, can have a maximum energy (cutoff energy) equal to the ionization potential of the electron ( $\sim 10$  eV) plus the kinetic energy imparted by the laser field [82, 83, 78]. This second term, known as the pondermotive energy, depends on the driving laser intensity and thus can be quite large ( $\sim 60$  eV). However, much of the time the free electrons do not follow a path that leads to recombination with the nucleus due to collisions and other effects. For most systems this is a major inhibitor of HHG photon flux. However, a high finesse optical cavity presents a great way to recycle the fundamental photons, thereby increasing the overall efficiency of the system. This fact, along with the power enhancement aspect make passive enhancement cavities a natural fit for combining HHG and the frequency comb.

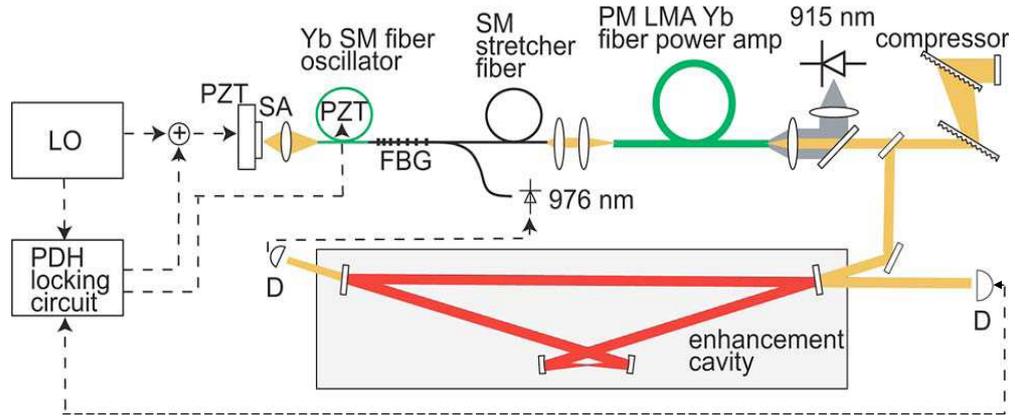


Figure 4.8: Experimental setup: LO, local oscillator; SA, saturable absorber; PZT, piezo actuator; FBG, fiber Bragg grating; D, photodetector; PM, polarization maintaining; SM, single mode.

#### 4.3.1 The Ytterbium Fiber Laser and Amplifier

The light source employed in this experiment (see Fig. 4.8) was a Yb-doped similariton [84, 85] laser in conjunction with a double-clad high power fiber amplifier. The similariton laser produces parabolic shaped pulses that are *self similar* solutions to the Ginzburg-Landau equation. The most useful aspect of the similariton pulses is that they exhibit optical wave breaking only at nonlinear phase shift levels that are 1 to 2 orders of magnitude higher than dispersion-managed(DM)-solitons. Thus, the the pulse energy that can be achieved before the onset of wave breaking can be on the order of 100 nJ/pulse. Thus, a similariton laser is a good choice for a system involving the study of high field physics.

Another crucial element of this laser system was the amplifier section, which used double clad Yb-core fiber. Double clad fiber consists of a Yb doped core section surrounded by two cladding layers. Multimode pump light is guided by the two cladding structure, while the singlemode laser light is contained in the core section. This allows for the use of high power multimode diodes ( $\sim 100$  W of cw power at 915 nm) as optical pumps for the amplifier gain. The femtosecond pulses

out of the laser cavity have to first be stretched so that the large energy amplification does not induce wave breaking through nonlinearity. In our system, the approximately 100 fs pulses were stretched to 70 ps before entering the amplifier. After the amplification, the pulses were recompressed using high-quality transmission gratings leading to a net average output power of 13.1 W, with pulses as short as 75 fs.

### 4.3.2 Measurement of the intra-cavity intensity

The pulses from the amplifier were then sent to the enhancement cavity (see Fig. 4.8). The cavity used in this experiment was a 4 mirror bow-tie configuration. The input coupler had a reflectivity of 99.94%, while the three high reflectors had reflectivities of 99.988%. The input coupler was selected purposely to have a transmission that was higher than the net losses of the other mirrors (i.e. not impedance matched). A lower IC reflectivity allows more of the frequency comb to be coupled into the cavity due to the lower dispersion. To match the amplified laser pulse spectrum to the cavity spectrum we used a birefringent filter between the oscillator and the amplifier, resulting in a 90% spectral overlap between the input spectrum and output spectrum of the cavity (see Fig. 4.9-(a)).

To keep the laser frequencies matched to the cavity resonances, we employed a Pound-Drever-Hall lock (see Chapter 2) by modulating the intra-cavity PZT with a low amplitude, 1 MHz sine wave. The feedback loop was then closed using the PZT (which had a 70 kHz actuator bandwidth) and a slower 1 kHz intra-cavity fiber stretcher. The offset frequency,  $f_o$ , was controlled using via temperature tuning a section of the fiber oscillator. Using these actuators, we were able to actively lock the laser frequencies to the cavity resonances for a time span of several hours. The net result of the cavity was a power enhancement of  $\approx 230$ , calculated by measuring the output power through one of the high reflectors.

This means on average a photon inside the cavity will be combined with 230 other photons, leading to a dramatic power enhancement (2.3 kW average power for 10 W input). Assuming a temporal gaussian envelope, the peak intensity at the intra-cavity focus was  $3 \times 10^{14} \text{W/cm}^2$ .

To verify this level of peak intensity, we performed ionization experiments using noble gases inside our enhancement cavity. A gas jet was mounted such that gas could flow through the intra-cavity focus position, while electrodes (at 10 V/mm bias voltage) were mounted orthogonal to the jet to allow for the collection of ions. We measured the plasma current through the electrodes for Xenon (ionization energy  $\approx 12$  eV), Krypton (ionization energy  $\approx 14$  eV), and Argon (ionization energy  $\approx 16$  eV). The results of this measurement are shown in Fig. 4.9-(b). As expected from the ordering of the ionization energies, Xenon reached the highest level of plasma current, followed by Krypton, then Argon. This data and the ionization thresholds are comparable to previously published data [86] for ionization of these noble gases.

This experiment marked the first time peak pulse intensities on the order of  $10^{14} \text{W/cm}^2$  at MHz level repetition rate were produced. Using cavity techniques and an ionization measurement we confirmed the high peak intra-cavity intensity. Further increases in peak intensity are being pursued [87, 88], but not without difficulty. Perhaps the most daunting challenge is finding a way to increase the damage threshold of the high finesse mirrors used in the enhancement cavity. We experienced many problems with mirror damage (resulting in a significant loss in cavity finesse) when running the experiment at high power for extended periods of time. The problems of mirror damage can be alleviated somewhat by ultrasonically cleaning all the components before putting them in the vacuum chamber. Perhaps another technique to reduce the peak intensities on the mirrors is to design a cavity that has large spatial mode sizes at the mirror positions.

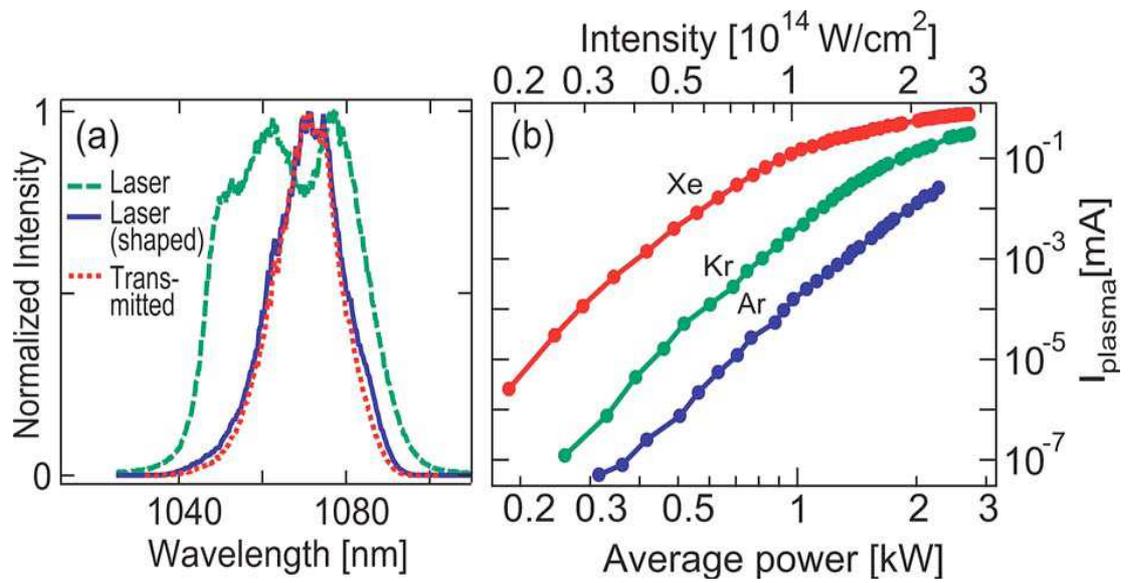


Figure 4.9: (a) Optical spectrum transmitted through the cavity (dotted, 17.2 nm FWHM) and emitted by the laser system with (solid, 19.3 nm FWHM) and without (dashed, 40 nm FWHM) filtering before the amplifier. (b) Current through plasma for various noble gases at 10 V/mm bias as a function of laser power and peak intensity. The gas pressure was 750 mTorr.

## Chapter 5

### Introduction to Waveguide Arrays

The last chapter demonstrated the many wonderful experiments one can perform using a mode-locked fiber laser. These experiments were successful, in large part, due to our precise control of the fiber laser repetition frequency ( $f_{rep}$  and  $f_o$ ). The next half of this thesis will address a way to achieve another important goal that we have for mode-locked fiber lasers: a robust mode-locking mechanism. The technology used in the lasers of the previous chapters (APM-NPR) is robust in a long timescale operation sense, but not in an environmental sense. To make this clearer, consider the following example: one of the NPR mode-locked lasers in our lab has been running virtually non-stop for my entire graduate career (6 years), but if I were to move the laser box the slightest amount the mode-locking would surely stop. To mode-lock the laser again could take a few minutes, a few hours, or even a few days. This is a frustrating point for those interested in using the fiber laser as a tool outside of a lab environment. Even for those working in laboratories a more robust system could prevent lots of time being wasted trying to find the right polarization for mode-locking. In the following chapters, we will introduce and characterize a new type of effective saturable absorber known as waveguide arrays for use as a robust mode-locking mechanism.

The response time of the NPR saturable absorber technique is essentially instantaneous, thus the pulse duration is only fundamentally limited by the gain

bandwidth of the laser. However, there are several downsides to NPR mode locking including the environmental sensitivity of the laser and the problem of finding the initial mode locking state by a random walk in polarization space. The problem of finding the initial mode locking state is alleviated somewhat in the Kerr lens mode locking of the Ti:Sapphire laser. In that laser, the experimenter's goal is to mis-align the cavity for cw operation, while making the pulsed operation more gain favorable. Thus, there is a *recipe* for mode-locking and less time is spent wandering around in a random polarization walk. However, the Ti:Sapphire laser also has sensitivity issues and can drop out of mode locking with relatively modest perturbations. Using a saturable absorber based on a totally different mechanism (waveguide arrays), it appears that we can combine the good aspects of a KLM laser with a robust fiber system that is polarization insensitive [89].

### 5.1 What are Waveguide Arrays?

Perhaps the most well known waveguide device is an optical fiber. Light can be coupled into the optical fiber and guided with extremely low loss ( $<0.2$  dB/km) via total internal reflection. If you bring the cores of two optical fiber close together (i.e. to within a distance  $\sim$  wavelength of guided light), the light contained in the fibers will slosh back and forth between the two cores over some characteristic distance (known as the coupling length). This power coupling is due to the evanescent wave that exists outside of the core region. This technique has been used to great practical success in optical fiber splitters, where the input light power is split into two output ports. The ratio of the power between the two outputs is simply controlled via the distance between cores or the length of the interaction between the two cores.

The amount of coupling in this interaction is dependent on the propagation constants ( $\beta_{1,2}$ ) of the two waveguides. If  $\beta_1 = \beta_2$  (i.e. the phase velocities are

equal), the light field in one waveguide will completely couple over to the neighboring waveguide (a scenario known as phase synchronism). If  $\beta_1 \neq \beta_2$ , the light field will only couple partially between the two waveguides. A phenomenological model of this is easily constructed with two coupled differential equations:

$$\frac{dA_1(z)}{dz} = i\beta_1 A_1 + i\kappa A_2 \quad (5.1)$$

$$\frac{dA_2(z)}{dz} = i\beta_2 A_2 + i\kappa A_1 \quad (5.2)$$

where  $A_{1(2)}$  is the field amplitude in waveguide 1(2) as a function of propagation distance  $z$ , and  $\kappa$  is the coupling constant between the two waveguides. Note that  $\kappa \neq 0$  only if the waveguides are separated by a distance of a few wavelengths. If we begin with the initial conditions that we launch all of the power into the 1<sup>st</sup> waveguide, then  $A_1(z=0)=1$  and  $A_2(z=0)=0$ . The resulting power versus propagation plots are shown in Fig. 5.1, with part (a) showing the case for equal propagation constants, and (b) showing the unequal case.

The scenario we have just analyzed is the case for the smallest *array* of waveguides possible. While the math becomes more involved as we go from two waveguides to 10's of waveguides, the basic ideas stay the same: if the propagation constants are all equal then power flows freely between the waveguides, if the propagation constants are not equal this interaction is quenched. The next question then is, how can we selectively change the propagation constant? The answer to this question lies in the fact that index of refraction is, in general, intensity dependent due to the nonlinear coefficient  $n_2$  and is given by:

$$n(I) = n_0 + n_2 I \quad (5.3)$$

where  $n_0$  is the usual refractive index and  $I$  is the intensity in the medium.

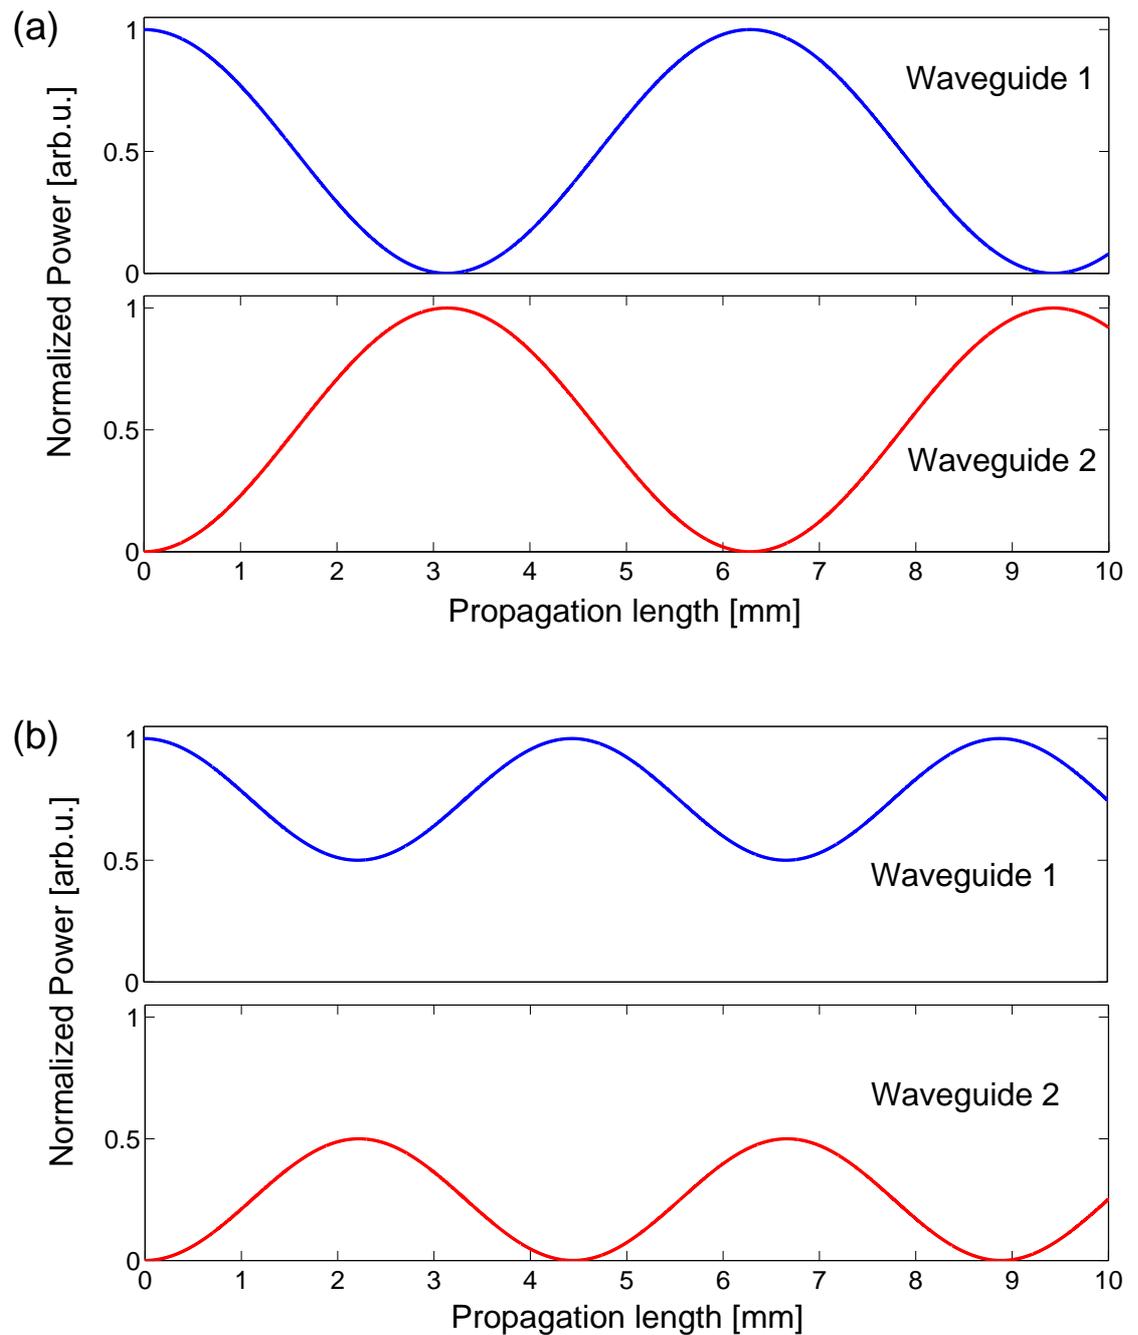


Figure 5.1: Power versus propagation distance for two cases. In both cases, light is launched initially into waveguide 1 (i.e.  $A_1(z=0)=1$  and  $A_2(z=0)=0$ ). (a) Power distribution for equal propagation constants  $\beta_1=\beta_2$ . (b) Power distribution for two waveguides with unequal propagation constants. In this simulation,  $\beta_1$  was 10% greater than  $\beta_2$ .

This means that the propagation constant is intensity-dependent, and  $\beta(I) = \frac{\omega n(I)}{c}$ . Thus, if we launch a low intensity field into a single waveguide, we can expect a broad field distribution at the output of the array due to the coupling between adjacent waveguides. However, if we launch a high intensity field into the same waveguide, the field will change its propagation constant relative to the neighboring waveguides effectively shutting off the power coupling. The field distribution at the output in this scenario will be localized around the input waveguide.

These two processes (the spreading and localizing of the power distribution) are discrete forms of the more well-known continuous phenomenon of diffraction and self-focusing. In the literature, the low intensity scenario described above is referred to as *discrete diffraction*, while the high intensity scenario is called *discrete self-focusing* [90]. The reader may already be seeing a connection between this discussion and the mode locking mechanism of the Ti:Sapphire laser system. In the next section I will expand on this idea and introduce a medium (AlGaAs) in which to build this effective saturable absorber.

## 5.2 Waveguide Arrays and Mode-Locking

The principle of operation of a waveguide array as a mode locking mechanism is essentially a discrete version of the KLM mode-locking of Ti:Sapphire lasers. When several identical waveguides are separated spatially by a distance that is on the order of the wavelength of light, evanescent-wave coupling occurs. This results in the input field experiencing discrete diffraction. However, when the light in the waveguide reaches a high intensity, the nonlinear index becomes significant and effectively shuts off the coupling to neighboring waveguides. By coupling the center waveguide to a laser cavity, we have an intensity discriminator. The low intensity (cw) fields will experience high loss due to the outer waveguides being uncoupled to the laser cavity, while the high intensity (pulsed)

fields will experience less loss due to the self-focusing effect and the coupled center waveguide. Because this process relies on the nonlinear index of refraction of the medium, the effective saturable absorber is as fast as the NPR saturable absorber (i.e.  $\approx$  a few optical cycles).

The waveguide arrays we will be discussing are rectangular waveguides composed of 3 layers of Aluminum-Gallium-Arsenide (a III-IV semiconductor) grown by molecular-beam epitaxy (MBE) on top of a Gallium-Arsenide substrate. The vertical confinement is provided by a core layer of  $\text{Al}_{0.18}\text{Ga}_{0.82}\text{As}$  with an index of refraction of  $n_{\text{cladding}} = 3.343$ , and an upper and lower cladding layer of  $\text{Al}_{0.24}\text{Ga}_{0.76}\text{As}$  with an index of  $n_{\text{cladding}} = 3.312$  (see Fig. 5.2). The lateral confinement is provided by etching ridges into the top cladding layer, which effectively writes a periodic index profile in the lateral dimension [91]. This sort of device can be constructed in a different media such as silica (i.e. optical fiber). However, AlGaAs has a major advantage in that its nonlinear coefficient ( $n_2$ ) is roughly three orders of magnitude higher than that of silica. This means that self-phase modulation and thus discrete self-focusing can occur at relatively low peak intensities ( $\approx 500 \text{ W/cm}^2$ ).

Also of importance is the ability to tune the semiconductor band gap via aluminum alloy concentration. Previous work on these devices demonstrated [92] that the alloy concentrations mentioned above minimized multi-photon absorption, a crucial point for application in a mode-locked laser cavity (a discussion of this can be found in Chapter 6).

### 5.3 Testing the Saturable Absorber Properties

From a mode-locking perspective, the waveguide array seems to be a good candidate for a fast saturable absorber due to the spatial effects [93, 90, 94, 95]. In fact, the physics that determines the spatial profile of the transmitted beam

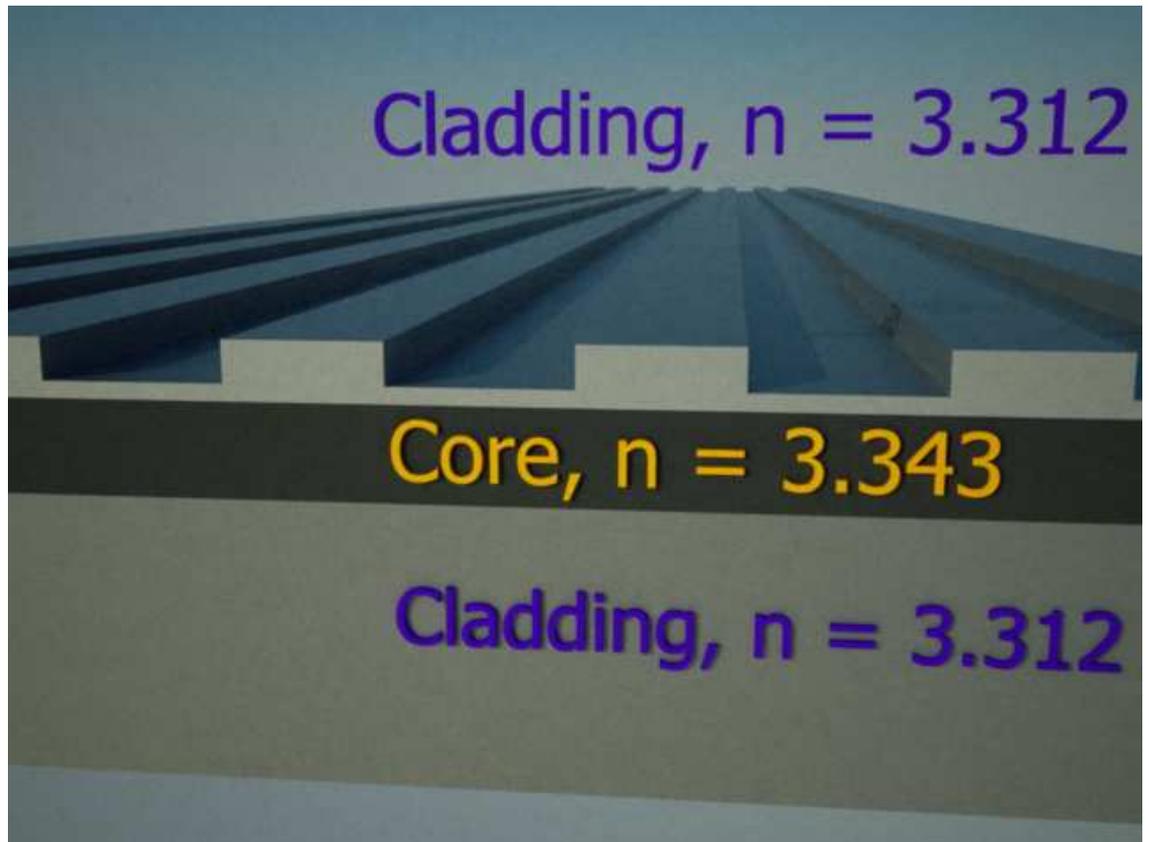


Figure 5.2: The waveguide array. Two vertical cladding layers provide index guiding in the vertical dimension, while the ridges in the top cladding layer provide guiding in the horizontal direction. The guided mode has an ellipticity of 2.7 and an effective mode area of  $19 \mu\text{m}^2$ .

is quite rich and suitable to theoretical treatment via the discrete nonlinear Schroedinger equation. This nonlinear equation can be tailored to include many physical processes that are present in the waveguide array such as discrete diffraction, normal dispersion, and self-phase modulation. Many novel spatial phenomena have been demonstrated using a waveguide array including discrete spatial solitons [93, 90], discrete modulational instability [94], and optical discrete surface solitons [95].

Despite the fact that most of the early experiments on waveguide arrays used pulses to achieve the necessary peak powers, temporal effects were initially largely ignored. Clearly, for application as a saturable absorber in a mode-locked laser the temporal effects on the pulse by the waveguide array must be understood. This situation led us to perform the first precise time domain measurements of pulse shaping in waveguide arrays [11]. In the following experiment, we carefully examined how the pulse shape in each waveguide depends on peak power of the input pulse using intensity autocorrelation. The output of the central waveguide showed significant shortening for high peak power due to attenuation of its lower power wings, as predicted in 2005 [96]. Simulations of the governing coupled-mode equations corroborate the observed experimental pulse-shaping results.

### 5.3.1 Theory of coupled modes

Coupled-mode theory provides an analytic reduction of the governing equations describing the propagation of electromagnetic energy in waveguides and waveguide arrays [97]. The theory assumes that the electromagnetic field is localized transversely in each waveguide and that the exchange of energy between the waveguides can be accurately modeled by an evanescent, linear coupling. When intense electromagnetic fields induce a self-phase modulation effect, coupled-mode theory can be modified to include the nonlinear index of refraction [93]. The result-

ing nonlinear coupled-mode theory agrees well with experiment [90, 98, 99, 100]. To leading-order, the nearest-neighbor coupling of electromagnetic energy in the waveguide array is included in the discretely coupled nonlinear Schrodinger equations (DNLSE):

$$i\frac{\partial A_n}{\partial z} - \frac{\beta''}{2}\frac{\partial^2 A_n}{\partial t^2} + \gamma|A_n|^2 A_n + c(A_{n+1} + A_{n-1}) = 0. \quad (5.4)$$

where  $A_n$  represents the normalized electric field amplitude in the  $n^{\text{th}}$  waveguide ( $n = -N, \dots, -1, 0, 1, \dots, N$  and there are  $2N+1$  waveguides). For the simulations in later sections we assumed a linear coupling coefficient of  $c = 0.82 \text{ mm}^{-1}$  and the nonlinear self-phase modulation parameter to be  $\gamma = 3.6 \text{ m}^{-1}\text{W}^{-1}$ . The parameter  $\beta'' = 1.25 \text{ ps}^2/\text{m}$  is the experimentally measured chromatic dispersion in the waveguide array. The simulations of eqn. (5.4) of the next sections are performed with 41 ( $N = 20$ ) waveguides [100] for various launch powers that match experimental conditions. A pseudo-spectral method was implemented that spectrally transforms the time-domain solution and uses a fourth-order Runge-Kutta for propagation in the waveguide.

### 5.3.2 Experimental System and Measurements

To generate the input pulses, we used a conventional mode-locked, Erbium doped fiber laser with a repetition rate of 25 MHz (operating at 1550nm) and a chirped-pulse amplifier/compressor system (see Fig. 5.3). Using dispersion compensating fiber (DCF), the normally chirped pulses from the fiber laser were further broadened to several picoseconds to avoid nonlinearities in the amplifier. These stretched pulses were then coupled to a bi-directionally pumped Erbium amplifier [101], which increased the pulse energy by a factor of 7, while maintaining the original pulse shape. The output of the amplifier was temporally com-

pressed/stretched in free-space by a diffraction grating compressor [102]. This compressor allowed us to achieve autocorrelation widths of several ps down to 200 fs. For this experiment, the compressor was adjusted to produce 600 fs pulses (FWHM as measured by autocorrelation), which are normally chirped and 3.8 times the Fourier transform limit. The output pulse energy was 3.5 nJ.

The pulses were coupled into the waveguide array using standard microscope objectives (40x) mounted on 3-axis stages. The input field was mode matched to the waveguide with a coupling efficiency  $> 50\%$ , corresponding to a peak power of 1.5 kW. The waveguide array had a 10  $\mu\text{m}$  center-to-center spacing between waveguides, with 1.5  $\mu\text{m}$  tall ridges and 4  $\mu\text{m}$  wide waveguides.

To verify that discrete self-focusing was occurring and to estimate the coupling coefficient between adjacent waveguides, the output power distribution of the array was measured as a function of input power (see Fig. 5.4). The energy localizes in the center waveguides for high power due to discrete self-focusing in the waveguide array [90]. At low power, the input light easily couples to neighboring waveguides and thus yields a nearly uniform power in each waveguide at the output end.

To measure the temporal reshaping effects of the waveguide array, background free autocorrelations were performed on the output of each waveguide. The autocorrelation measurements were performed in the crossed-beam geometry with a thin BBO crystal used for Type-1 second harmonic generation (SHG). A translation stage provided a scanning delay, while a 16-bit digitizer recorded the SHG signal detected by a photomultiplier tube. To increase the signal-to-noise of the measurement we also employed lock-in detection by chopping the beam just before it entered the autocorrelator. The data traces were continuously scanned and averaged. For reference, an autocorrelation of the input pulse was also recorded.

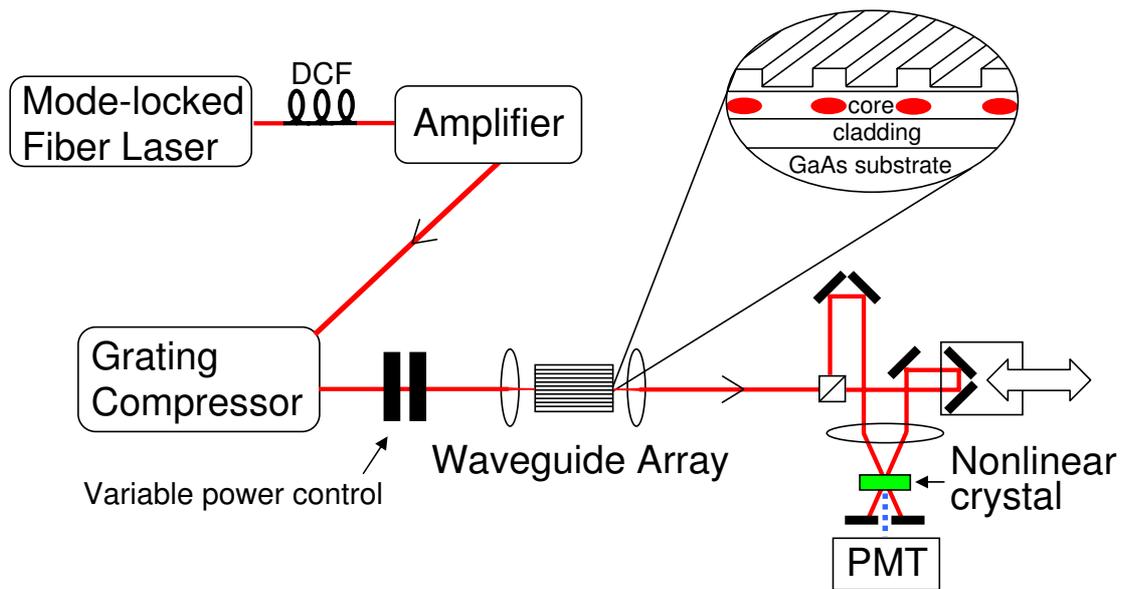


Figure 5.3: Experimental setup. The output of the fiber laser is broadened by dispersion compensating fiber (DCF) to avoid nonlinearities in the amplifier. The grating compressor is tuned to produce a 600 fs pulse. The variable power control consists of a half-wave plate and a polarizer. A temporal intensity autocorrelation of the output pulses is recorded using a photomultiplier tube (PMT).

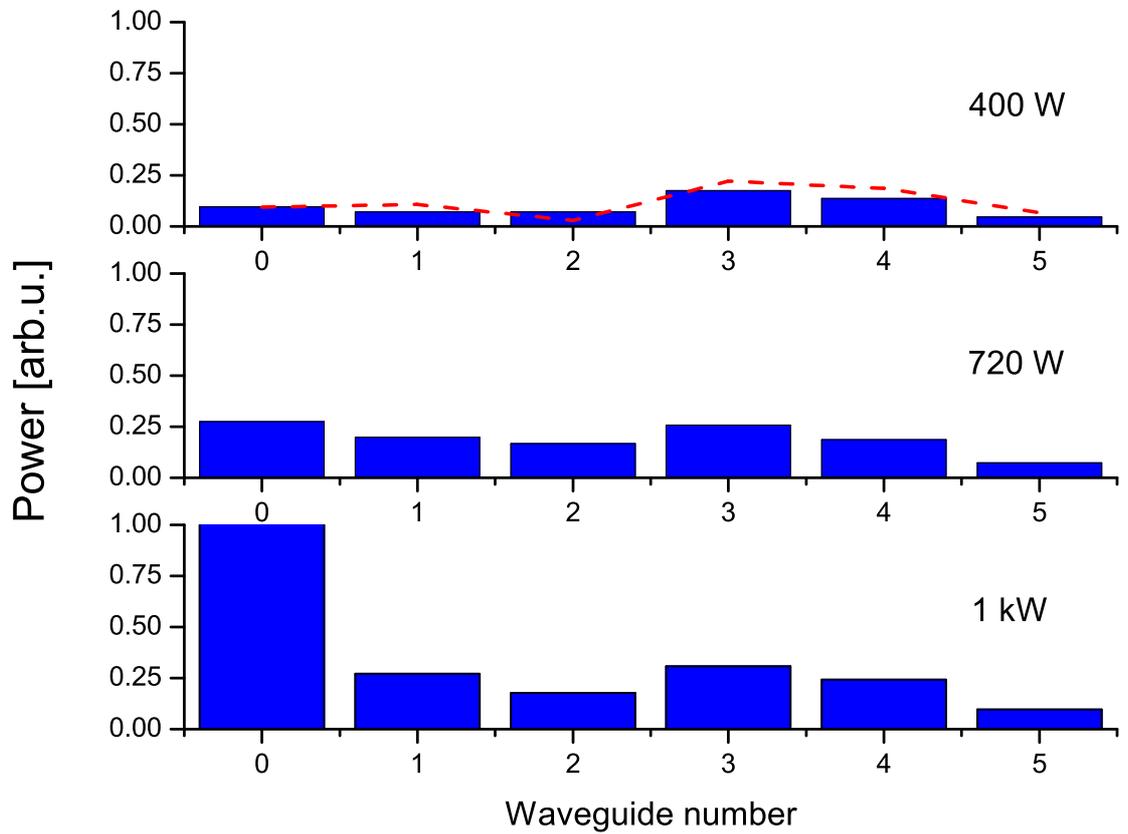


Figure 5.4: Measured power distribution of the output of the waveguide array. At low peak power the energy is delocalized and has a Bessel function dependence (dashed red line) on waveguide number. As the peak power of the launched pulse is increased, the power distribution shifts from the Bessel distribution to a localized distribution. The waveguide modes located symmetrically about the central waveguide had a symmetrical power distribution (not shown).

Fig. 5.5 shows the pulse reshaping effects of the waveguide array at each of the input powers, with experimental results on the left and a numerical simulation (provided by our theory collaborator: Nathan Kutz) of eqn. (5.4) on the right. At a peak power of 400 W, the output pulses from the central and outer waveguides were essentially identical to the input pulse (Fig 5.5-(a) and (d)). In this regime the  $\gamma$  term of eqn. 5.4 is negligible. The weak pulse launched into the center waveguide evanescently couples to neighboring waveguides. Thus, at the output multiple copies of the input pulse can be observed in each waveguide. As the input peak power is increased to 720 W ((b) and (e)), the pulse reshaping of the central waveguide begins to emerge. At the highest input peak power (1.5 kW) the  $\gamma$  term in eqn. (5.4) becomes non-negligible and the peak of the pulse decouples from neighboring waveguides. Meanwhile, the low intensity wings of the pulse are coupled to the nearest neighbor waveguides. The result is a shortened pulse in the center waveguide with the wings removed in agreement with the predicted nonlinear pulse shortening [96, 89]. Fig. 5.5-(c),(f) shows the output of the waveguide array at high power. The triple peaked autocorrelation of the outer waveguides in Fig. 5.5-(c),(f) is evidence of a double peaked pulse shape. The experimental results agree well with the numerical simulation at each power level.

Taking a closer look at the central waveguide pulse shape as a function of input power (Fig. 5.6) shows the reshaping increases strongly at high peak power. In Fig. 5.6, a 600 fs pulse was launched into the central waveguide and the output autocorrelation of the central waveguide was measured as a function of input power. The output pulse for the highest power case shows a pulse width of less than half that of the input pulse.

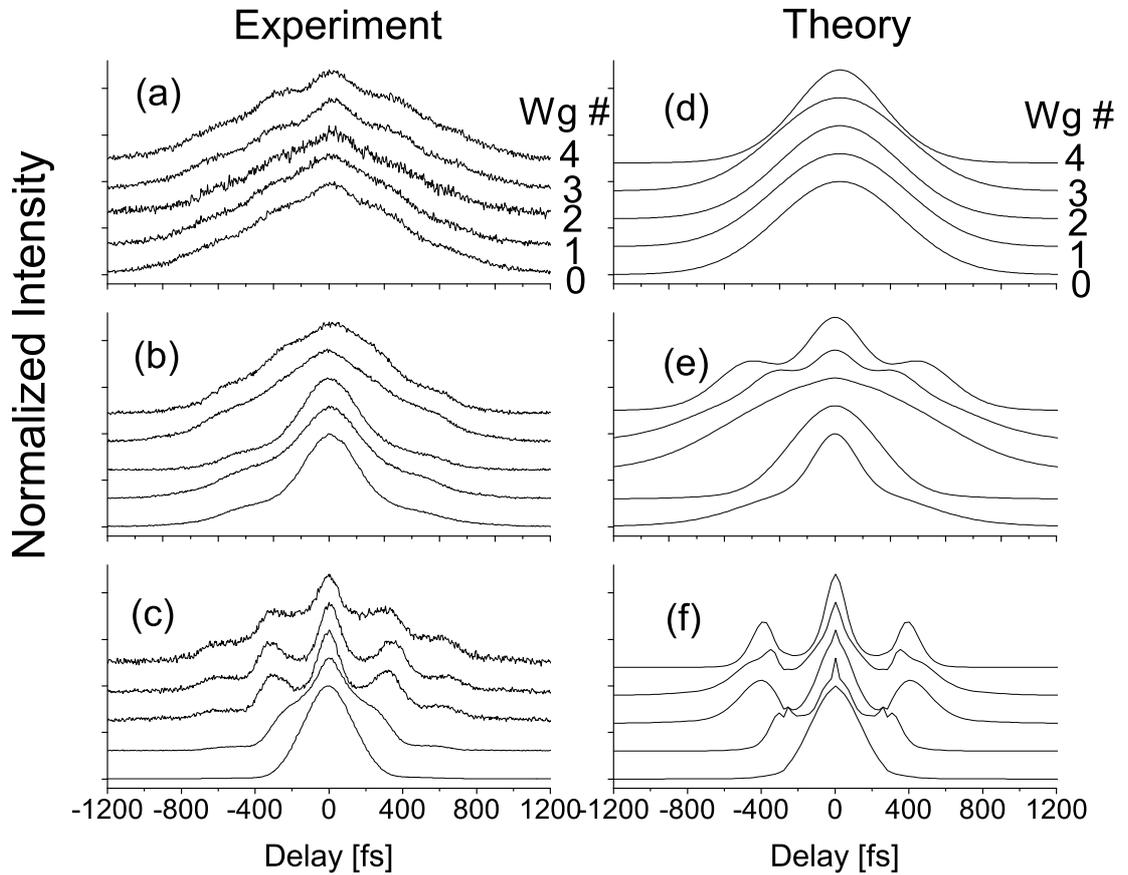


Figure 5.5: Autocorrelation signal versus waveguide number, with experimental results on the left ((a)-(c)) and theoretical simulations on the right ((d)-(f)). The three power levels shown correspond to those in Fig. 5.4, with (a) and (d) at 400 W, (b) and (e) at 720 W, and (c) and (f) at 1.5 kW. Pulse shortening in the center waveguide is observed in the 720 W and 1.5 kW cases. The autocorrelations are offset vertically for clarity, with the central waveguide being the lowest and the outer waveguides plotted sequentially higher on the vertical scale.

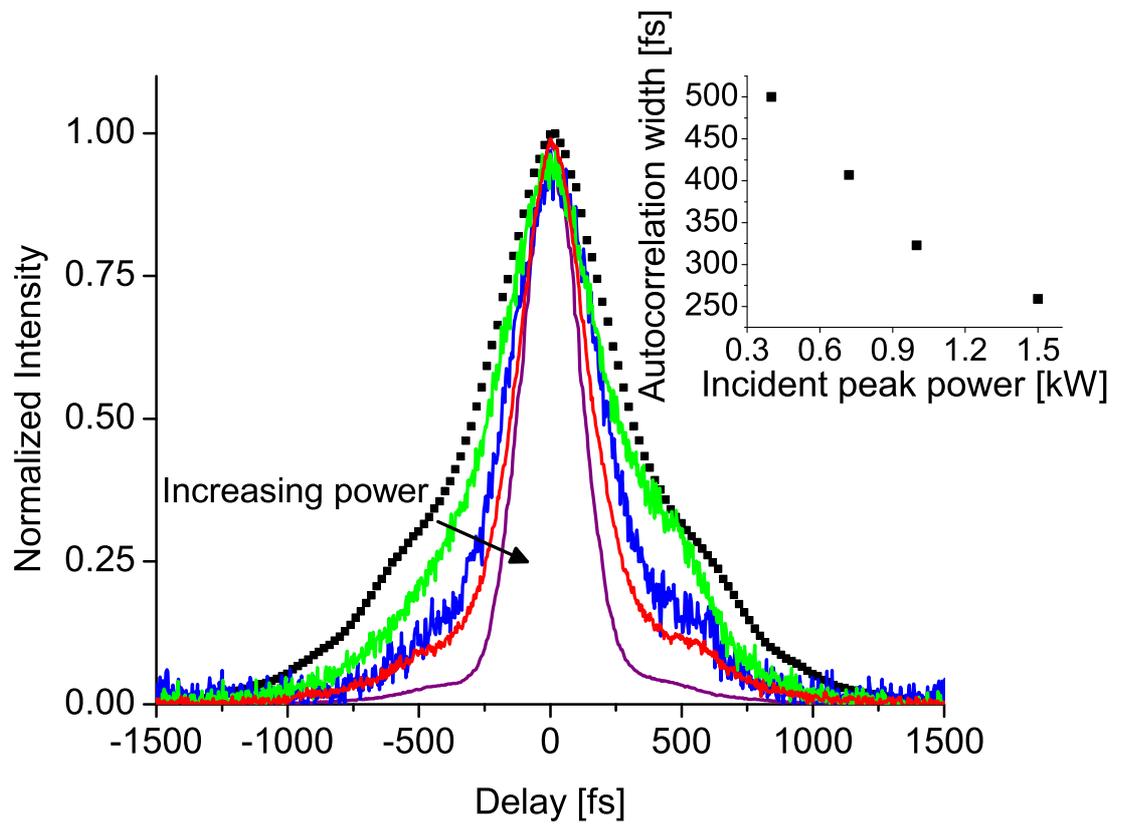


Figure 5.6: (color online) Output autocorrelation of central waveguide for input powers of 400 W, 720 W, 1 kW, and 1.5 kW. The inset shows the autocorrelation FWHM as a function of input power. The dotted trace is an autocorrelation of the input pulse.

### 5.3.3 Dispersion Measurement of the Waveguide Array

To confirm that the pulse shaping we observed was due solely to nonlinear effects in the waveguide array we had to check the effects of chromatic dispersion in the waveguide array. However, the fabricators of this device had little interest in knowing the chromatic dispersion and thus getting a number for this quantity was difficult. Thus, we decided to measure it ourselves by building a white-light interferometer [103, 104]. The white-light interferometer uses a He-Ne laser to calibrate the time axis of the data for a Fourier transform which allows for high precision measurements of the dispersion in the material. As summarized in Fig. 5.7, we found that the overall GVD of the waveguide was around  $+1500$   $\text{fs}^2/\text{mm}$  (i.e. normal).

This measurement helped make a very strong case for the validity of the nonlinear chopping effect we observed. First of all, and perhaps the strongest point, the input pulse was chirped with the same sign as the dispersion of the waveguide material. Thus, if anything, one would expect the pulse to broaden in time. Secondly, given the length of the waveguide array and the spectral width of the input pulse, dispersion should only change the pulse length by around 60 fs, well below the change observed ( $>300$  fs). And finally, a purely dispersive compression [96] would be independent of the peak power in the waveguide. Thus, the white light interferometry measurement confirmed that the pulse shaping was due to the self-focusing/discrete-diffraction effects in the waveguide array.

## 5.4 Conclusions

This experiment was the first observation of nonlinear pulse shortening in a waveguide array. From this measurement, the idea of using a waveguide array as a saturable absorber seems to be possible. This phenomenon could also have

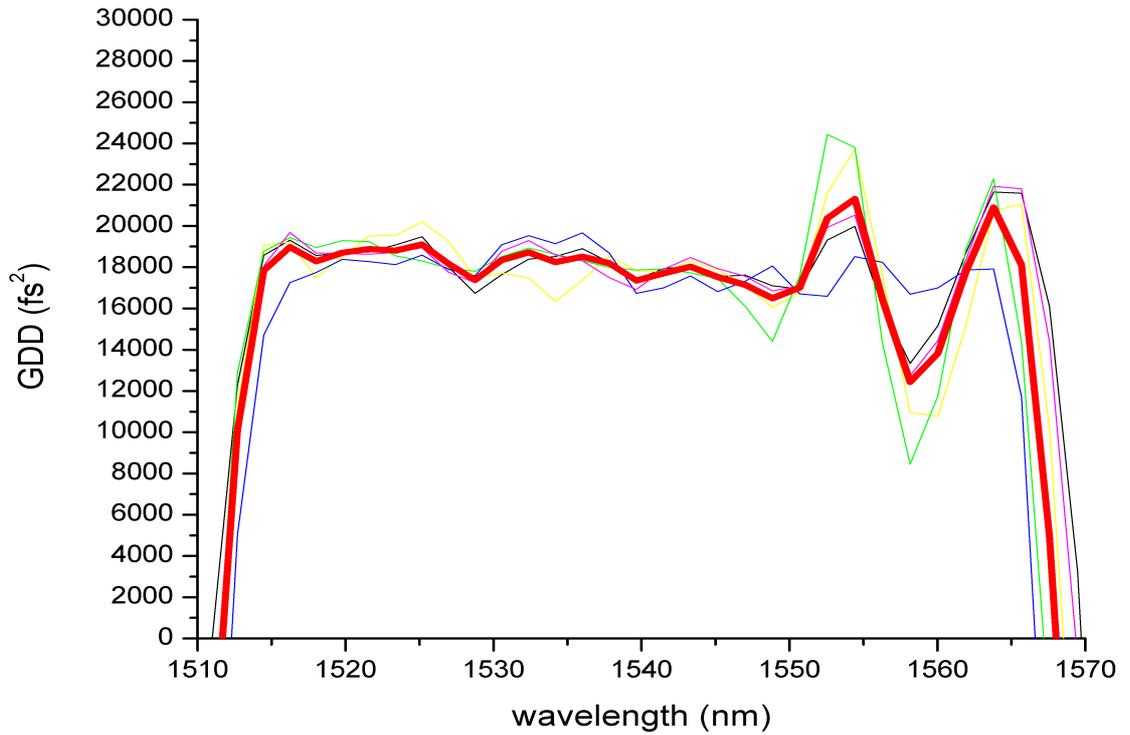


Figure 5.7: White light interferometer measurement of the waveguide array device showing the group-delay dispersion versus wavelength for several independent measurements (the thick red line is the average of all the measurements). The slope of this plot yields the third-order dispersion (TOD). We found that the waveguide array had a GVD of around  $1500 \text{ fs}^2/\text{mm}$  at  $+1550 \text{ nm}$ . This overall dispersion is dominated by chromatic dispersion, with the waveguide dispersion coming in at  $\sim +100 \text{ fs}^2/\text{mm}$ .

a wide range of applications involving pulse reshaping for long distance telecommunications. In this scenario, the fidelity of the pulses broadened by dispersion in the fiber could potentially be regained via the chopping mechanism as only the highest intensity parts (i.e. peaks) of the pulses would be transmitted. In the next chapter, I will discuss two measurements we performed on the waveguide array to examine multi-photon absorption in the AlGaAs and to investigate the pulse shaping in terms of the full electric field.

## Chapter 6

### Multi-Photon Absorption and Spectral Phase Clamping

In this chapter, I will discuss two important and perhaps surprising phenomenon that we observed in waveguide arrays. The first topic is multi-photon absorption in the AlGaAs material [92]. It is important to fully map out the intensities at which this effect becomes significant as multi-photon absorption represents an inverse saturable absorber (i.e. losses increases with intensity). A clean measurement of the effect of multi-photon absorption via power measurement through an individual waveguide is difficult due to the coupling between waveguides. Thus, a full power distribution measurement at the output is required. In the experiment described below we measure the power distribution at the output of the waveguide array for various input powers at various input chirps. We find that the discrete spatial soliton distribution becomes clamped at high peak intensity, with virtually no change in the output power distribution once the intensity is increased above the threshold value.

The second experiment in this chapter involves a measurement of the full electric field (amplitude and phase) of the pulses in the central waveguide via Frequency-Resolved Optical Gating (FROG). Surprisingly, the waveguide appears to set the spectral phase of the pulses to a specific value, regardless of the input chirp. This aspect of the waveguide arrays could lead to application as a novel pulse compression system.

## 6.1 Discrete Spatial Solitons

A discrete spatial soliton (DSS) is a spatial power distribution that remains unchanged over propagation. This type of soliton has been intensely researched over the last 20 years. Beginning with their theoretical prediction in 1988 [93], and the first experimental observation in 1998 [90] in AlGaAs waveguide arrays, it has been shown that discrete spatial solitons have some fundamental differences with their continuum counterparts [105, 99]. In particular, discrete solitons can be formed that are insensitive to angular perturbations on the input light field. Less stable solitons will shift their power distribution laterally across the array as the soliton propagates if there is a slight perturbation on the input angle. This shifting is not observed for continuum solitons as they possess rotational and translational invariance. The novel dynamics of such discrete solitons has generated interest for soliton steering, which could have wide ranging applications in optical switching and optical computing. Another fundamental difference between the continuum solitons and discrete solitons is the ability of the discrete soliton to be either bright or dark. This fact is due to the ability of the waveguide array exhibit either normal or anomalous diffraction [105]. Mathematically, we can describe the DSS using the discrete nonlinear Schrödinger equation (DNLSE) from Ch. 5.

Solutions of the DNLSE for the waveguide array show that, at high power, discrete spatial solitons form [93]. In the following experiment, we examine the evolution of the DSS power distribution as the power incident into the waveguide is increased to levels that induce heavy multi-photon absorption.

## 6.2 Multi-photon absorption theory

Apart from the features mentioned in the last chapter, AlGaAs allows for band gap engineering via changing the alloy concentration. At room temperature

( $T = 300\text{K}$ ), the band-gap energy for  $\text{Al}_x\text{Ga}_{1-x}\text{As}$  is  $E_g(x) = 1.424 + 1.247x$  [eV] (where  $x < 0.45$ ) [106]. For alloy concentrations of  $x > 0.45$ , the band gap becomes indirect. By changing the alloy concentration we have the ability to move the band gap energy around so that multi-photon absorption is minimized at a given wavelength [92]. In particular, an alloy concentration of  $x=0.18$  yields a half band gap energy (1.65 eV) for which two photon absorption (at 1506 nm) is just out of the range of the telecommunications C band (1530 to 1565 nm). Although it has been shown that this kind of band-gap engineering can eliminate the effects of two-photon absorption (2PA) and minimize the effects of three-photon absorption (3PA), at very high intensities multi-photon absorption invariably becomes non-negligible. Theoretically, we can describe the total absorption coefficient  $\alpha$  as the sum of the absorption coefficients from the various multi-photon absorptions:

$$\alpha = \alpha_1 + \alpha_2 I + \alpha_3 I^2 \quad (6.1)$$

where  $\alpha_1$  is the linear absorption coefficient ( $\sim 0.1 \text{ cm}^{-1}$ ),  $\alpha_2$  is the two-photon absorption coefficient ( $\sim 0.04 \text{ cm/GW}$ ), and  $\alpha_3$  is the three-photon absorption coefficient ( $\sim 0.05 \text{ cm}^3/\text{GW}^2$ ). Using the Erbium fiber laser and amplifier system from the previous chapter, we were able to reach intensities of  $24 \text{ GW/cm}^2$  in the waveguides. At this intensity level, multi-photon absorption in the waveguide was visible. In the the next section, I will examine the effect of 3PA on discrete spatial solitons and show that 3PA effectively clamps the output power distribution of the discrete spatial soliton; with further increases in input power having almost no effect on the output distribution.

While this phenomenon is not yet fully understood, we are working on simulations to aid in understanding. Physically, light in the central waveguide self-focuses tighter at higher power but also undergoes more nonlinear absorp-

tion. As the average power is increased and more light is coupled into the central waveguide, the outer waveguides also receive more energy due to the evanescent coupling. This experimental observation is explained by the fact that we cannot change the relative propagation constants enough to completely shut off nearest-neighbor coupling. Thus, the outer waveguides still grow in energy (and experience no 3 photon absorption) while the power in the central waveguide increases at a reduced rate due to the 3 photon absorption. The net effect of this is that the relative power distribution does not change beyond a certain intensity level.

### 6.3 Experimental Setup and Results of the Three Photon Absorption Measurement

The pulses that were launched into the waveguide in this experiment were generated using the same mode-locked, Erbium doped fiber laser and amplifier that was used in the last chapter. The pulses were coupled into the waveguide array using standard microscope objectives (40x) mounted on 3-axis stages. The input field was mode matched to the waveguide with a coupling efficiency  $> 50\%$  (excluding Fresnel losses), and we again used a waveguide array with  $10 \mu\text{m}$  center-to-center spacing between waveguides, with  $1.5 \mu\text{m}$  tall ridges and  $4 \mu\text{m}$  wide waveguides. The objective lens focused the input beam to around  $20 \mu\text{m}^2$ , which is roughly equal to the area of one waveguide mode. This tight focus, together with the highest peak power from the compressor systems, yielded a peak intensity of  $24 \text{ GW}/\text{cm}^2$ .

Initial results were obtained by tuning the compressor such that the fluorescence from the waveguide was at a maximum and measuring the resulting power distribution. A 3-dimensional plot of this power distribution measurement is shown in Fig. 6.2-(a). Adding a three-photon absorption term to the DNLSSE resulted in good theoretical matching with experiment (see Fig. 6.2-(b)). After

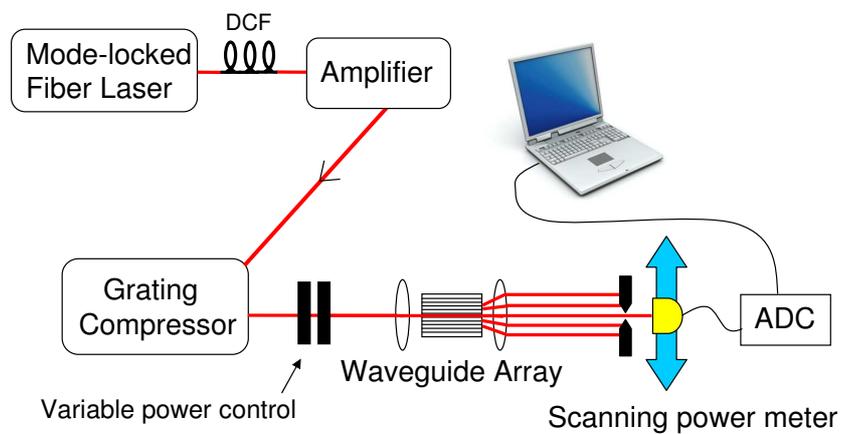


Figure 6.1: Experimental Setup for the three photon absorption measurement. The power at the output of the waveguide is apertured (to restrict measurement to individual waveguides) and measured on a photodiode. The photodiode is mounted on a translation stage oriented perpendicular to the direction of light propagation. This signal is then digitized and stored on a computer.

this measurement, we set out to get a full picture of the multi-photon absorption as a function of peak intensity at three different input chirps.

We began by tuning the compressor to deliver a rather broad pulse with normal chirp. An autocorrelation of this pulse is shown in Fig. 6.3-(a). The structure on the pulses from the compressor is mainly due to self-phase modulation in the erbium doped amplifier. This pulse was coupled to one waveguide in the array, and the resulting power distribution at the output was measured (see Fig. 6.3-(b)). Self-focusing overcomes discrete diffraction at roughly  $7 \text{ GW/cm}^2$ . Beyond this point, the central waveguide dominates the power distribution. From Fig. 6.3-(b), however, we cannot definitely say that the power distribution of the discrete soliton is clamped to a fixed distribution. It does appear that the distribution is headed that way, but our intensity levels are not high enough to make a strong case.

To achieve higher intensity levels, we adjusted the compressor so that it produced the shortest output pulse (i.e.  $\text{GDD} \approx 0 \text{ fs}^2$ ). The shorter pulse allows (see Fig. 6.4-(a)) us to reach an intensity of  $24 \text{ GW/cm}^2$ . An identical power distribution scan was then performed for this shorter input pulse, and the results can be seen in Fig. 6.4-(b). As will be demonstrated later, the higher intensity access allows us to see the distribution clamping in effect.

We then tuned the compressor so that it delivered anomalously chirped pulses to the waveguide array. This input pulse is shown in Fig. 6.5-(a). The power distribution for this pulse is shown in Fig. 6.5-(b). The data appears to show a slightly higher level of clamping than the normally chirped pulse. While this pulse is roughly as long temporally as the pulse in Fig. 6.3-(a), the sign of the chirp is opposite. Since the waveguide array presents normal dispersion, this input pulse must recompress slightly due to chromatic dispersion. Thus, the anomalously chirped pulse will remain shorter inside the waveguide array and

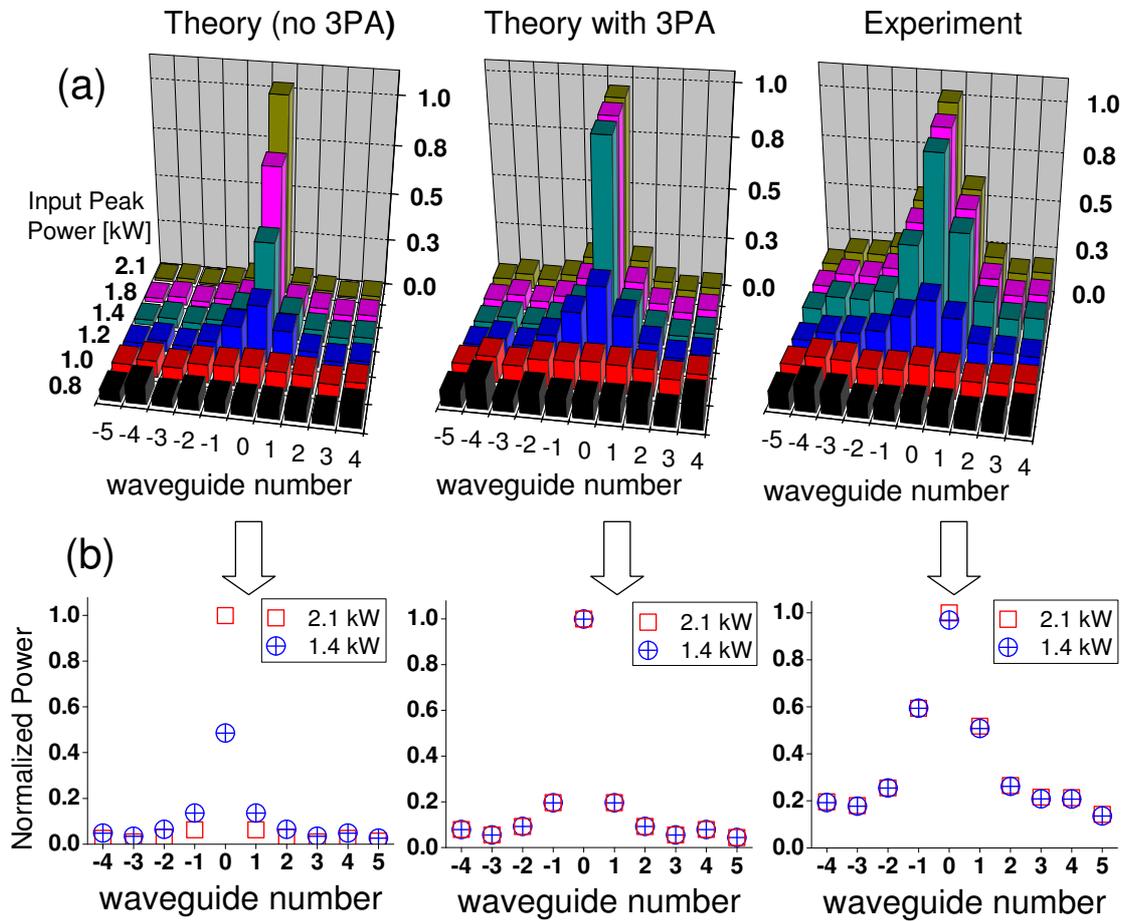


Figure 6.2: (a) 3d plot of the theoretical power distribution (left), theoretical distribution with 3PA (middle), and experimental distribution (right). (b) 2d plot of the distributions for two high power levels. This measurement was the first indication that power clamping was occurring in the waveguide array.

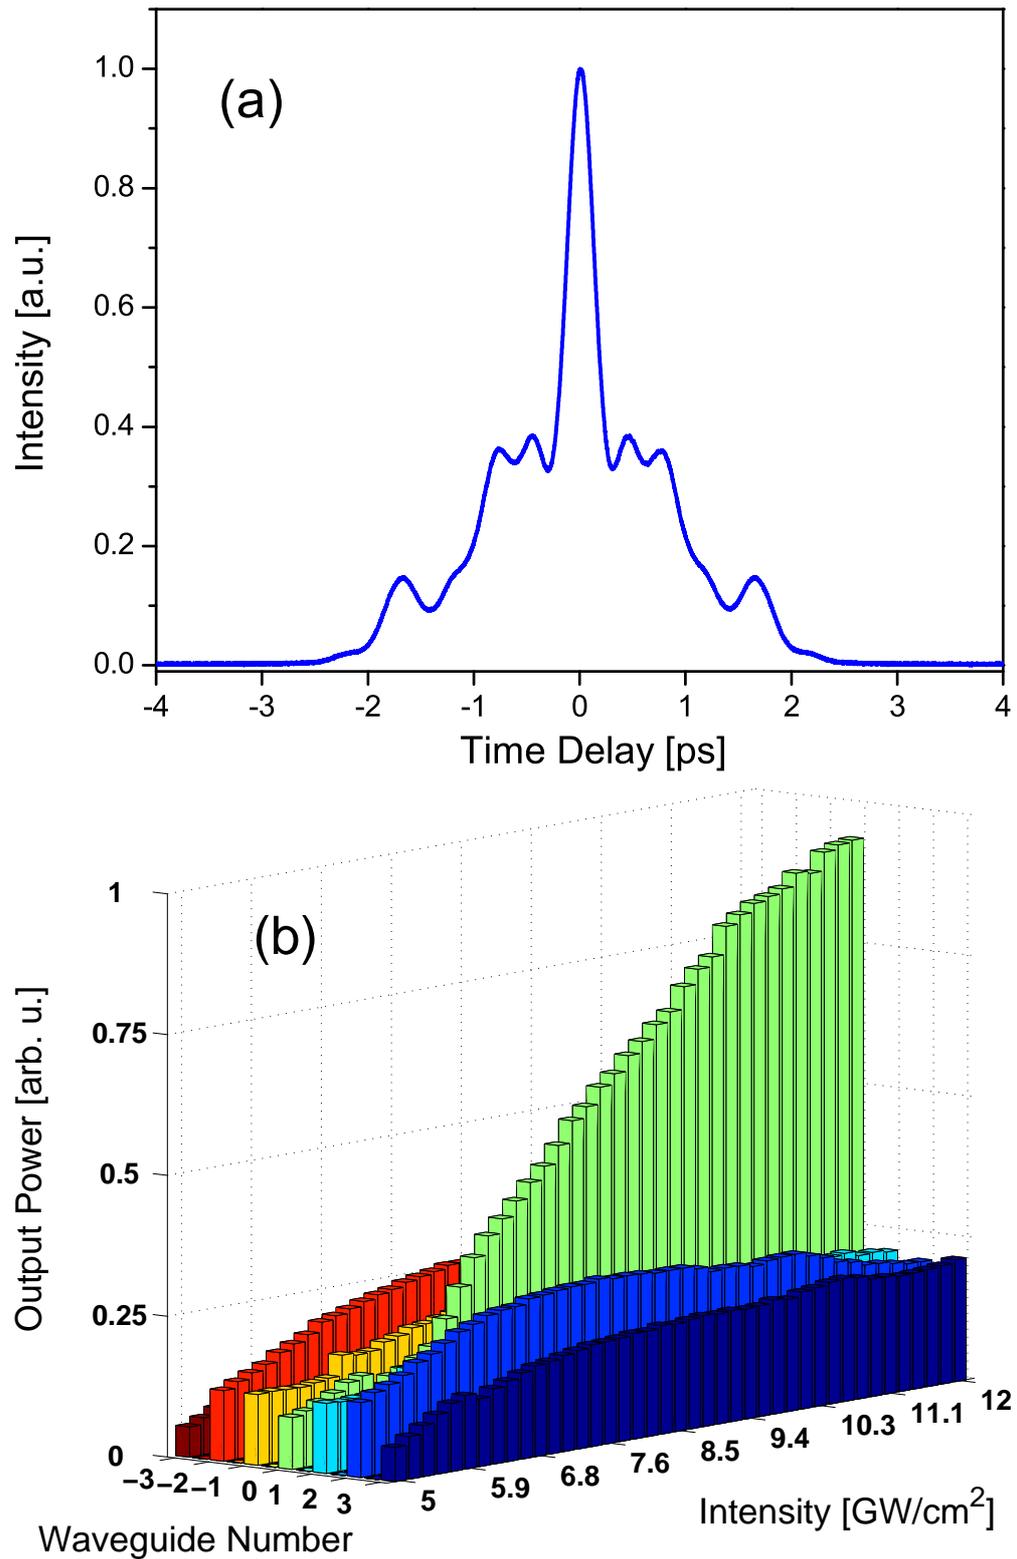


Figure 6.3: (a) Autocorrelation of the normally chirped input pulse, and (b) the resulting 3D power distribution as a function of peak intensity. At low input power, the power distribution is spread out and the outer waveguides actually have more power than the central waveguide. However, as the input power is increased, discrete diffraction gives way to self-focusing and the central waveguide eventually dominates the power distribution.

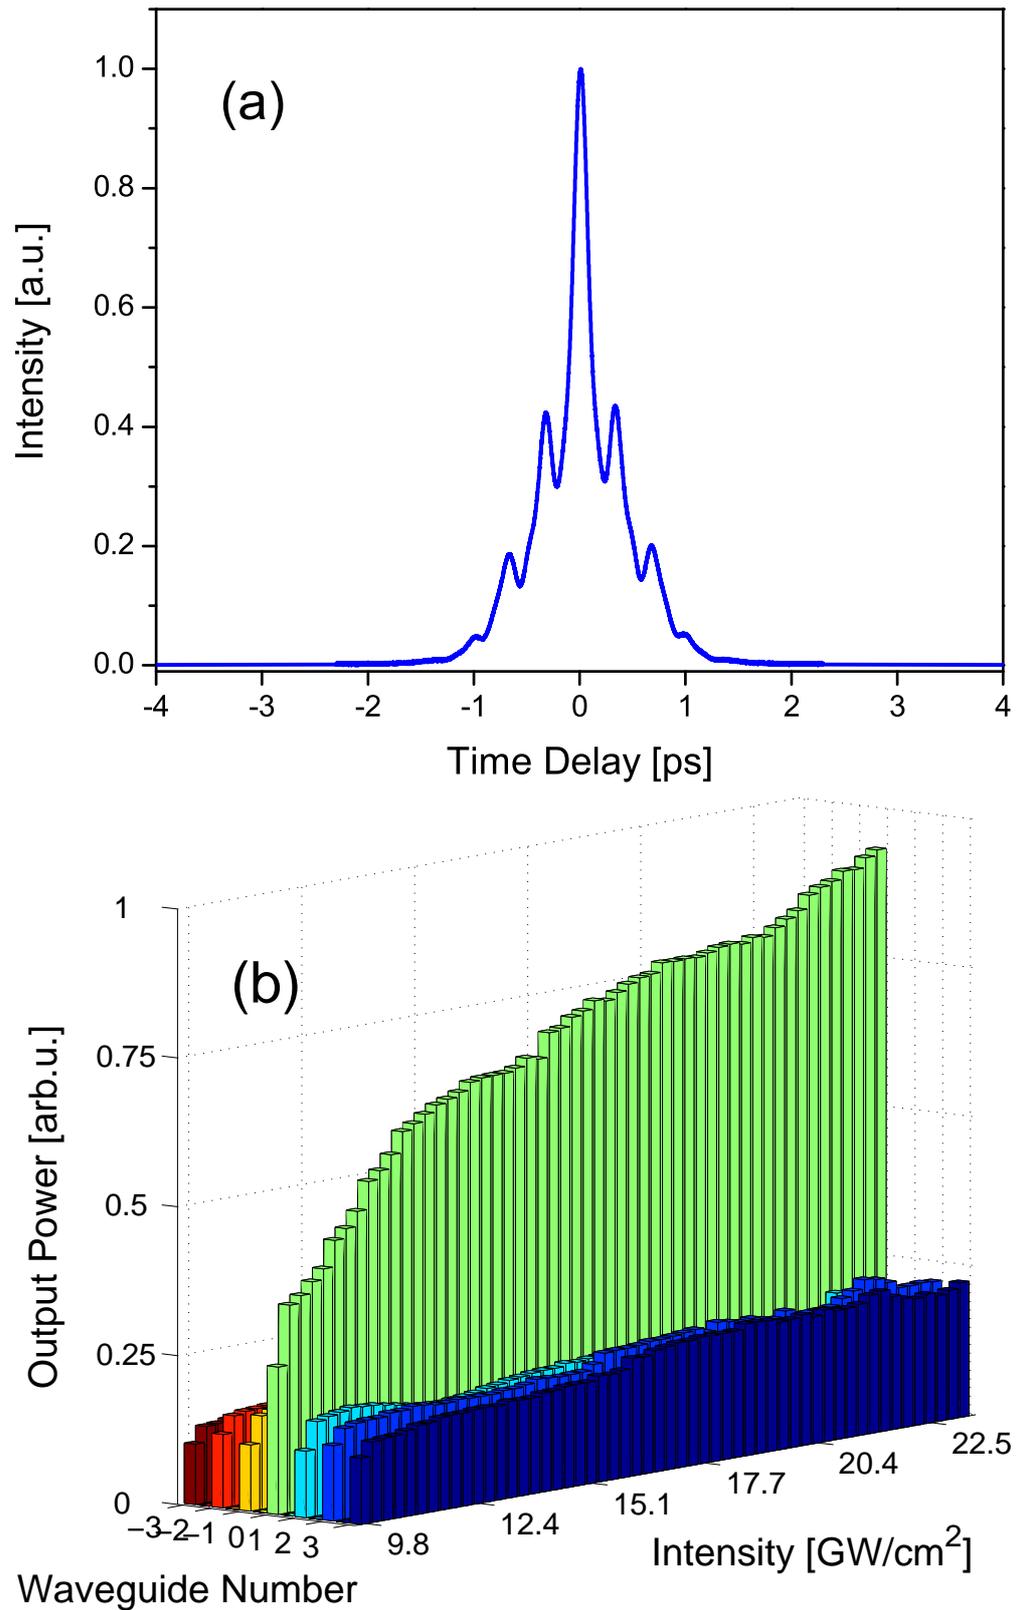


Figure 6.4: (a) Autocorrelation of the shortest input pulse, and (b) the resulting 3D power distribution as a function of peak intensity. As the input peak power is increased, the relative power distribution clamps to a set distribution. The overall power of this distribution still increases slightly, but the relative distribution is fixed (see Fig. 6.6).

undergo a higher level of 3PA than a normally chirped pulse of the same length.

To make a clear and direct comparison between all of the cases studied above, Fig. 6.6 plots a power ratio of the power in the neighboring waveguides divided by the power in the central waveguide. Thus, if clamping does indeed occur, the ratio should level off to some set value. As can be seen in Fig. 6.6, this effect occurs relatively early in the power scan for the shortest input pulse. This ratio for the anomalously chirped pulse moves toward the asymptote at a slightly higher average power than the shortest input pulse, while the ratio for the normally chirped pulse just barely reaches the asymptote.

Perhaps the most surprising part of this plot is how well the shapes of the ratios for symmetric waveguides match up. I should also note that the red curve here has a high enough peak power that the little bump indicative of the transition from discrete diffraction to self-focusing (see blue and green curves) is shifted off to the left. In other words, even the lowest intensity points of this data are in the self-focused regime.

In conclusion, the distribution of power in the discrete spatial soliton reaches a set point when the three photon absorption is strong. For chirped pulses you simply have to go to high average power (and thus high peak power) to see the effect. By looking at three different input chirps and calculating the power ratio as we have done, an asymptote line becomes clear, and all three chirps approach it from different slopes. The peak intensity that is needed to reach this asymptote is around  $13 \text{ GW/cm}^2$ . The effect of 3PA on the discrete spatial soliton will be useful for cavity design of a mode-locked laser utilizing waveguide arrays. Also, the power clamping effect may find application as a power locking device to be used just before an amplifier. The large downshift in power fluctuations due to the 3PA in the waveguide array would serve to reduce amplitude noise on the optical field before it is amplified.

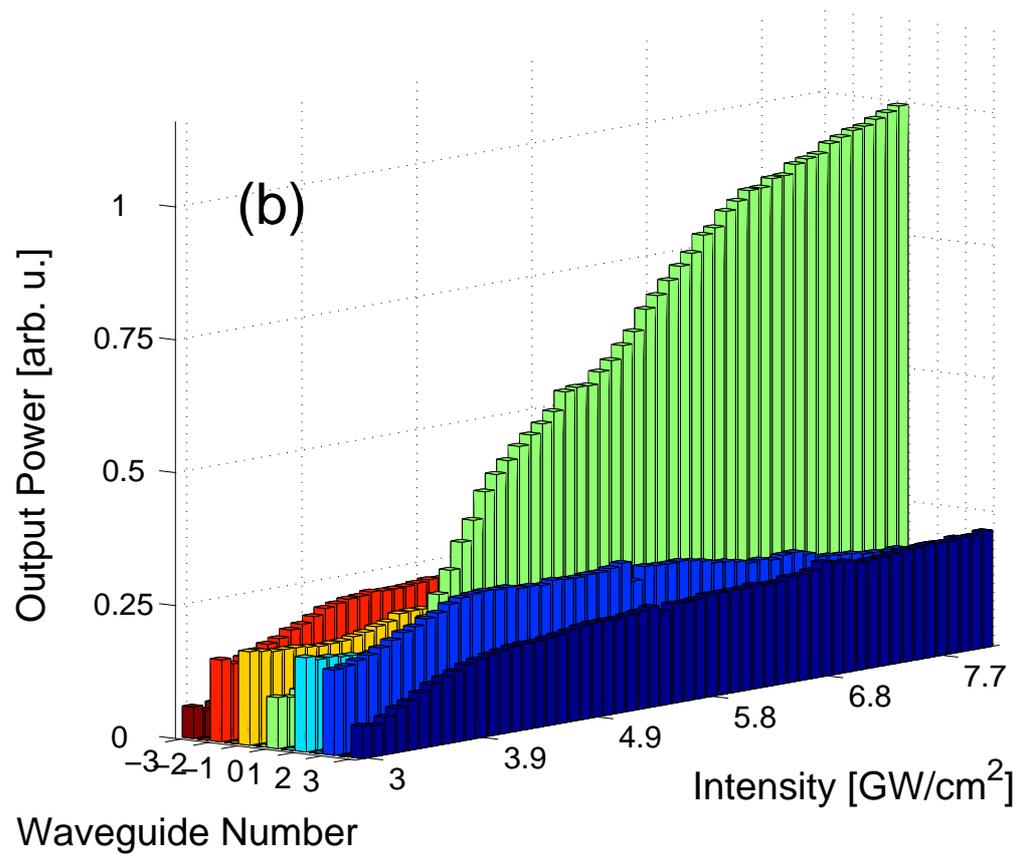
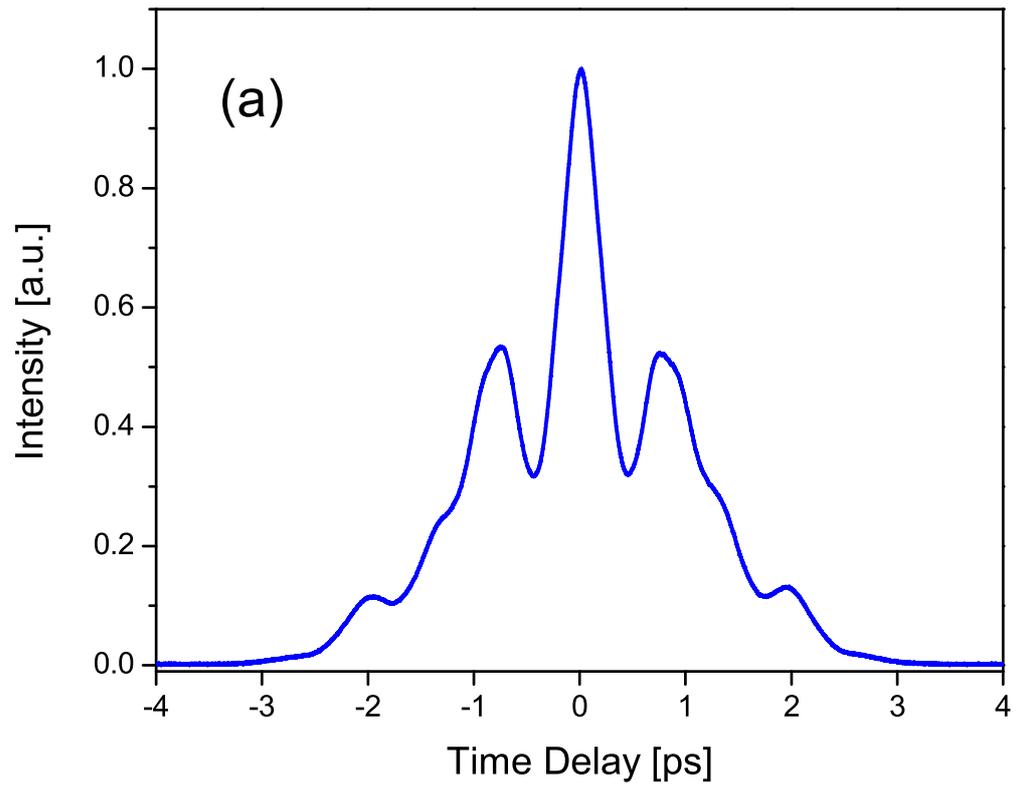


Figure 6.5: (a) Autocorrelation of the anomalously chirped input pulse, and (b) the resulting 3D power distribution as a function of peak intensity.

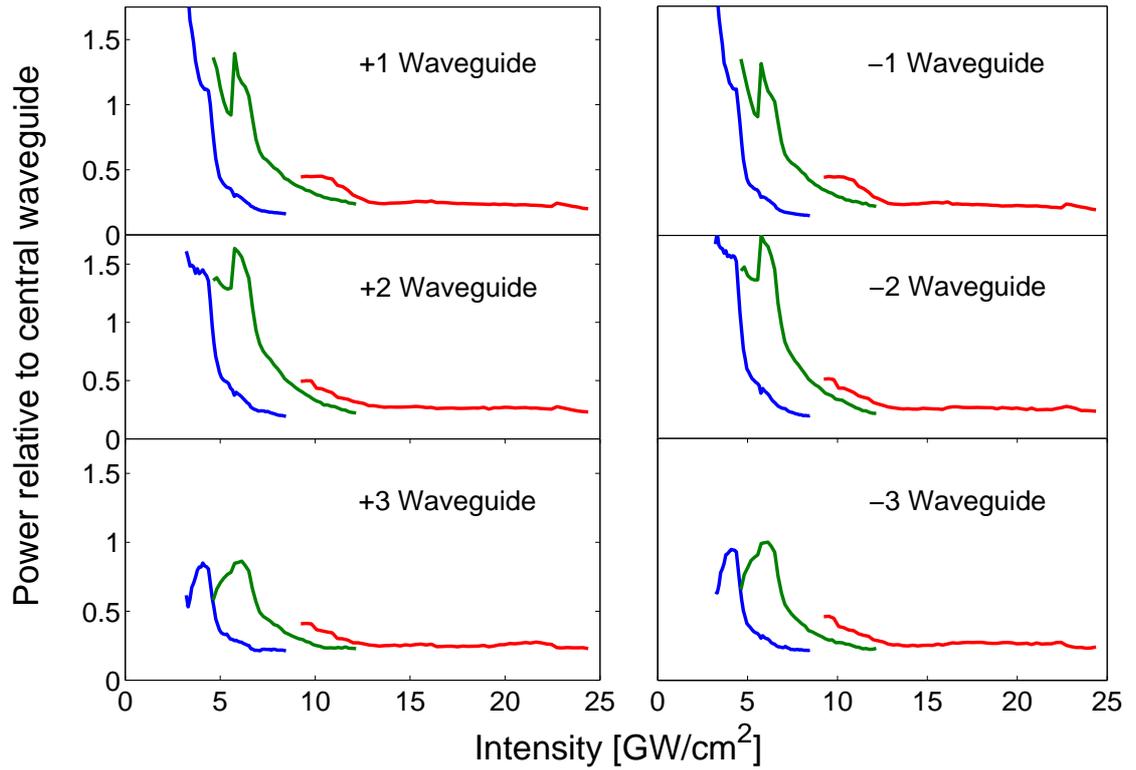


Figure 6.6: Relative power in the neighboring waveguides. The red curve corresponds to the shortest pulse, the green curve corresponds to the normal chirp pulse, and the blue curve corresponds to the anomalous chirp pulse.

## 6.4 Chirp effects on self-focusing

To characterize the self-focusing as a function of peak power at various input chirps, we scanned the pulse compressor through its entire range while recording the power distributions. This measurement of self-focusing versus chirp is summarized in Fig. 6.7. As expected, at the shortest input pulse, the multi-photon absorption is dominant and creates a relative power minimum. The peaks to the left and right of the minimum represent the best self focusing for anomalous (left peak) and normal (right peak) chirp on the input pulse. As the chirp is further increased, the peak power is reduced and the self-focusing begins to spread out. For application as a mode-locking mechanism, this measurement will help design the dispersion map of the cavity.

## 6.5 Spectral Phase Clamping

This section covers an experiment aimed at measuring the full electric field shaping due to the waveguide array. In the previous chapter we examined this shaping in terms of an autocorrelation measurement. This type of measurement, while useful for getting a qualitative picture of the physics occurring, has several drawbacks. First of all, the autocorrelation is a convolution of the intensity profile of a pulse with itself, thus it is always symmetric even if the real pulse is not symmetric. Also, the autocorrelation is a phase insensitive technique. Thus, we do not get a clear picture of how the waveguide array is effecting the spectral phase of the pulses. A much better method of characterizing the pulse shaping effects of the waveguide array is Second Harmonic Generation-Frequency Resolved Optical Gating (SHG-FROG) [107]. This technique involves recording a spectrally resolved autocorrelation (see Fig. 6.8), or spectrogram (Spectrum vs. Time Delay). The information contained in the spectrogram is sufficient to retrieve the full pulse

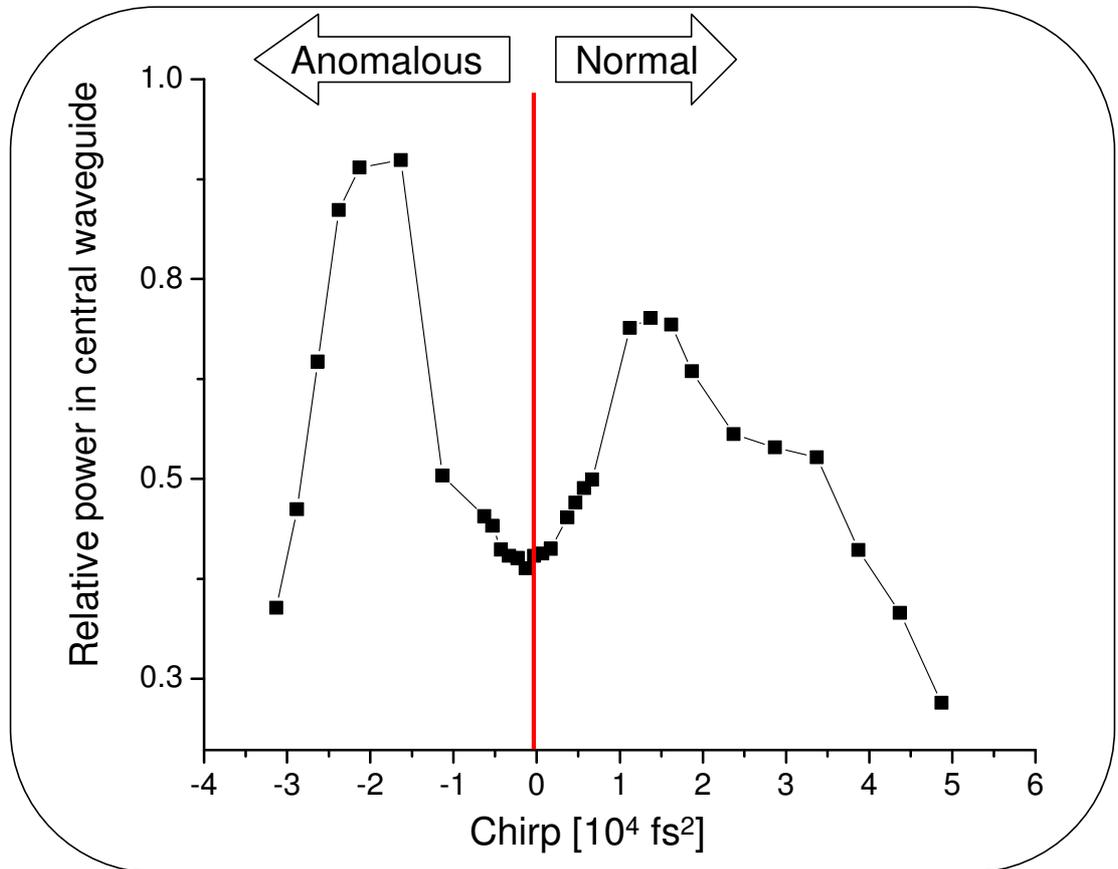


Figure 6.7: Relative output power versus input pulse chirp for the highest average power available. The vertical line indicates the minimum due to multi-photon absorption. The two peaks on either side of the minimum show that tight self-focusing occurs for both anomalous and normal chirp. Moving beyond these peaks, the pulse begins to spread out to a level at which the self-focusing degrades due to the lower peak power of the pulse. In principle, a higher average power would shift the peaks further away from the center minimum.

shape and phase. We employed a commercial computer algorithm to analyze the spectrograms and retrieve the electric field amplitude and phase.

Using the FROG system, we explored the extended parameter space of chirp and average power by carefully examining the waveguide array's effect on the full electric field of an ultrashort pulse as a function of these variables. Surprisingly, we find that the waveguide array has the ability, at sufficiently high peak power, to set the output pulse's chirp to a fixed value, irregardless of the chirp on the input pulse. The input pulses were again generated using the system found in Fig. 6.8. This system produces pulse energies of 3.5 nJ and allows us to change the pulse chirp from normal (+200,000 fs<sup>2</sup>) to anomalous (-100,000 fs<sup>2</sup>).

The first test of this system was to reproduce the shortening results of the previous chapter. Accordingly, we tuned the compressor to the same setting and measured the spectrogram as a function of input peak power (varied using the variable power controls of Fig. 6.8). The spectrograms were then analyzed using the FROG inversion algorithm to produce the electric field envelopes seen in Fig. 6.8. This plot gives us a much clearer view than the autocorrelation of what is happening to the pulse in the central waveguide. As the input peak power is increased, the structured input pulse (which is not shown but is equivalent to the envelope shown in the leftmost pane) is transformed into a shortened and truly symmetric electric field. With this confirmation, we moved on to measuring the pulse shaping effects as a function of pulse chirp.

To get an accurate picture of how the waveguide array was affecting the spectral phase, we first mapped out the range of electric fields from the compressor. By tuning the distance between gratings in the compressor system we varied the pulsewidth over a wide range as shown in Fig. 6.10. As is shown in the figure, the retrieved spectral phase yielded maximum and minimum GDD values of +200,000 fs<sup>2</sup> to -100,000 fs<sup>2</sup>, respectively. We then systematically coupled each

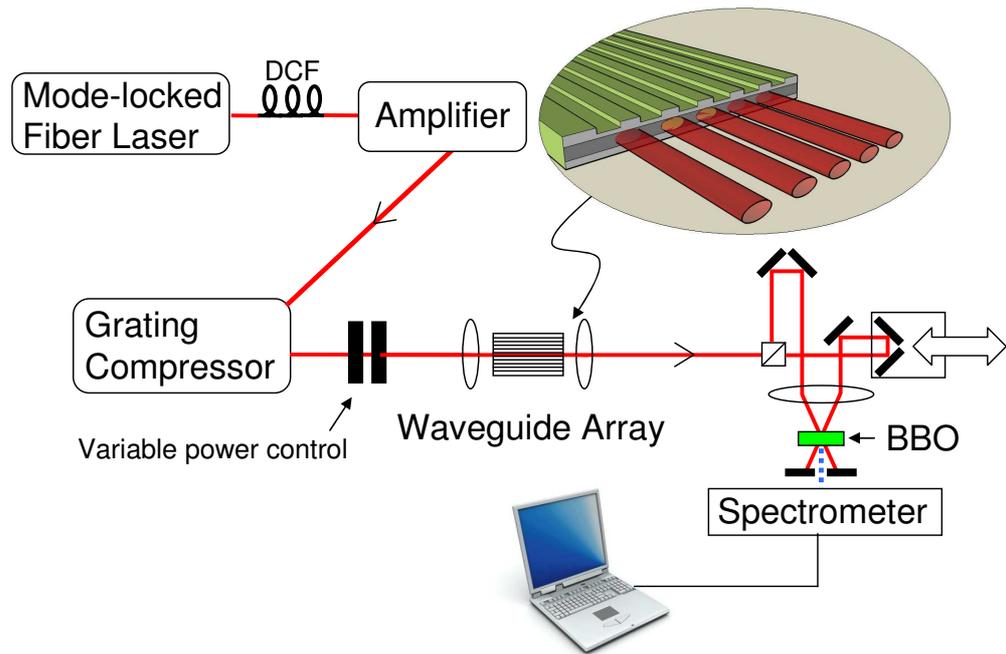


Figure 6.8: Experimental setup. The output of the fiber laser is broadened by dispersion compensating fiber (DCF) to avoid nonlinearities in the amplifier. The grating compressor allows tuning of the chirp for the input pulses. The variable power control consists of a half-wave plate and a polarizer. Light in the central waveguide is apertured and sent to a background-free SHG-FROG. The data from the spectrometer (FROG trace) is stored on a computer and analyzed with a commercial algorithm.

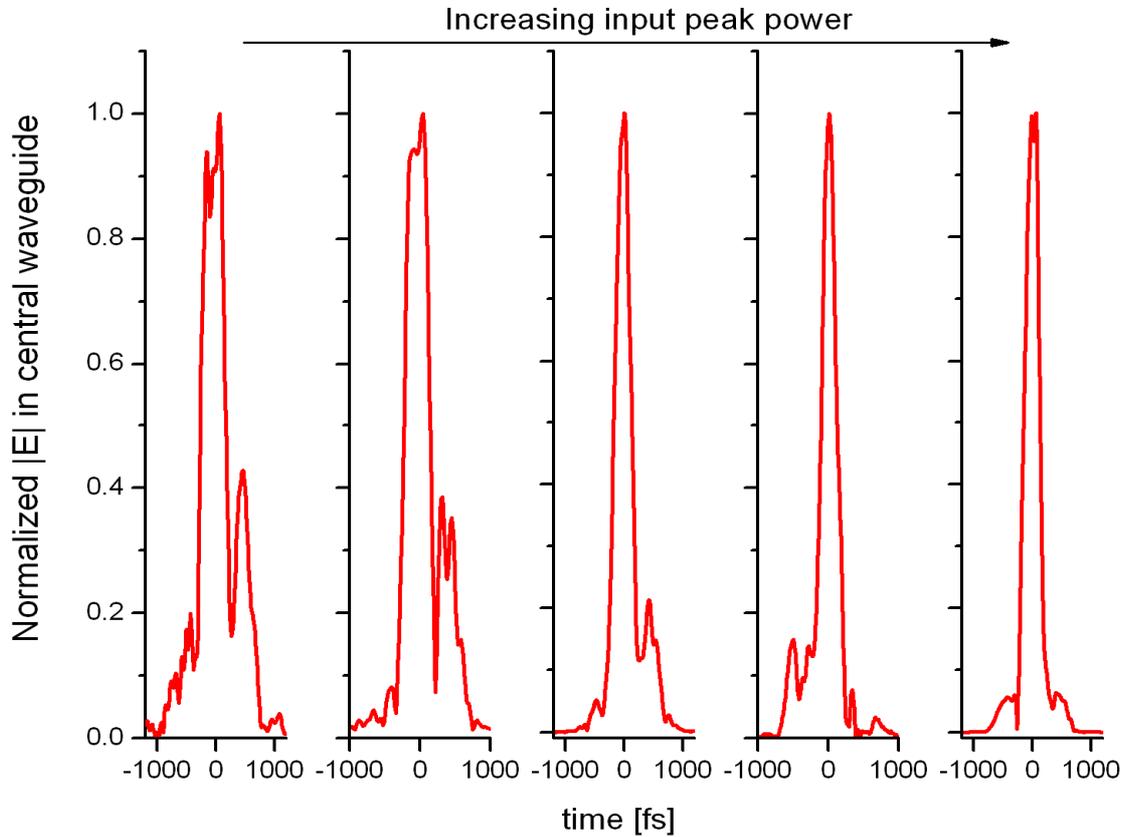


Figure 6.9: Electric field envelope shortening in center waveguide as a function of input peak power. This measurement confirms the pulse shortening effects of the previous chapter's autocorrelation data. At low power, the highly structured input electric field is basically unchanged after propagation through the center waveguide (leftmost plot). However, as the peak power is increased the pulse shaping in the center waveguide begins to take hold and clean up the wings of the pulse. At the highest power, the output electric field is clean and symmetric (rightmost plot).

of these fields (at the highest average power available) into the waveguide array and measured the field that emerged from the central waveguide. The results of this measurement are shown in Fig. 6.11. The effect of the waveguide array is clear, it chops the pulses in the time domain roughly the same amount regardless of input chirp. As Fig. 6.11 shows, this electric field reshaping occurs over a wide range of input chirps.

Since this reshaping phenomenon is explained by a power dependent process, we investigated the chopping effect of the central waveguide at several average power levels. The data from this experiment, which is summarized in Fig. 6.12, show the chopping effect in the frequency domain rather than the time domain. Viewing the phenomenon in this domain is instructive as it tells us the amount of linear frequency chirp on the pulse's spectral phase. With this knowledge we can say how much spectral phase distortion the pulse has relative to its Fourier limit.

The lowest average power level shown (26 mW) corresponds to the threshold of the self-focusing regime for this waveguide array. The reshaping phenomenon vanishes at the lowest power, and the output chirp of the pulse in the central waveguide closely follows that of the input chirp. At the medium power level (45 mW), the reshaping is observed for all input chirps except for input chirps beyond  $+130,000 \text{ fs}^2$ . At these chirp levels, the pulse is too stretched out in time to achieve the necessary peak power to undergo full reshaping. Thus, the outlier points rapidly move toward the input chirp from the compressor. At the highest average power (60 mW), the waveguide array sets the output GDD to around  $+17,000 \text{ fs}^2$  for all inputs. Particularly interesting is that the waveguide array does this for either sign of input chirp (normal or anomalous).

We can classify the effect of the waveguide array as a fixed point attractor for the spectral phase. As the average power of the input pulses is increased, the attraction towards a fixed point GDD ( $\sim +17,000 \text{ fs}^2$ ) becomes stronger.

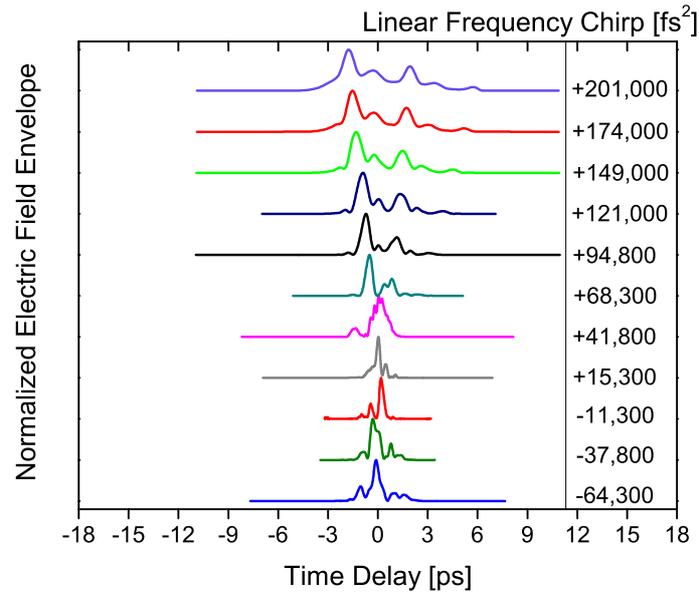


Figure 6.10: Electric field envelopes before the waveguide as a function of time and compressor setting. The plot shows the electric fields as the compressor is tuned from the most normal dispersion setting (top trace) through the zero group delay dispersion (GDD) point (fourth trace from bottom) to the most anomalous setting (bottom trace).

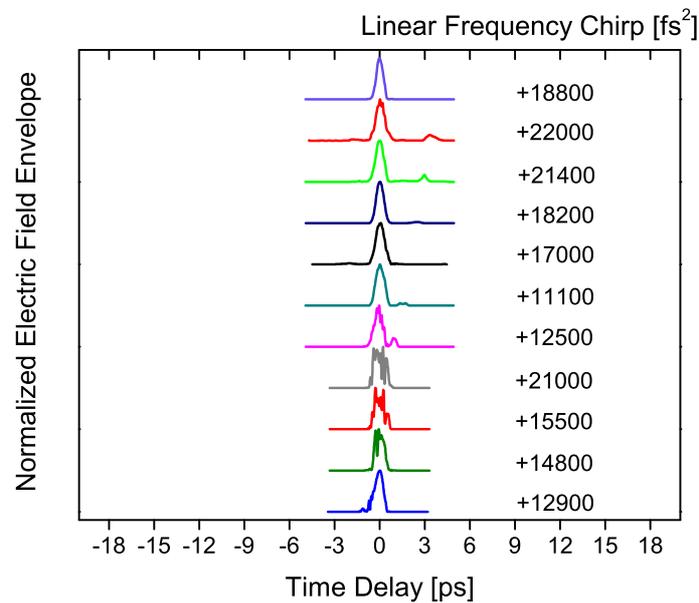


Figure 6.11: Electric field envelopes after the waveguide as a function of time and compressor setting. Each trace corresponds to an input shown in Fig. 6.10. The rapid oscillations near the zero dispersion point are due to self phase modulation in the material.

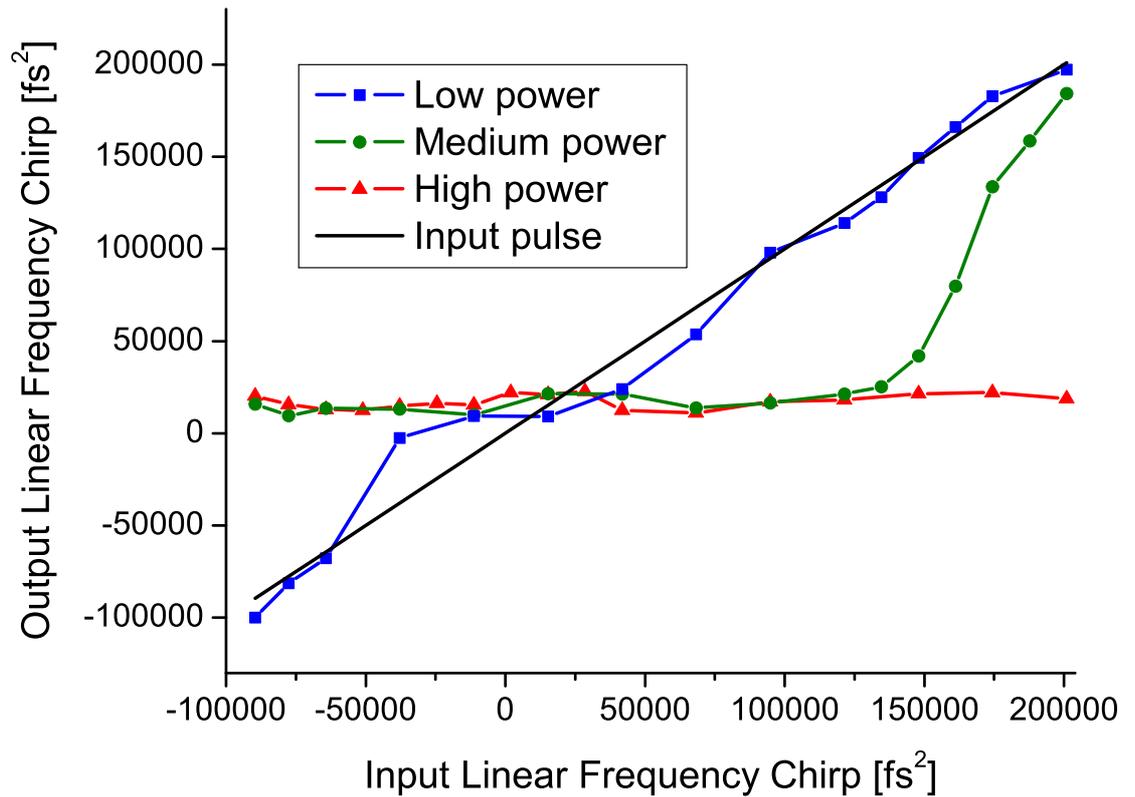


Figure 6.12: Second order spectral phase (GDD) of the input pulses and output pulses at various average power levels. The black line is the input pulse spectral phase.

Fig. 6.13 helps to visualize this statement. In this figure, the input average power and spectral phase is plotted as a circle, and the vector arrow points toward the output average power and spectral phase of the pulse in the central waveguide. For the low average power case (26 mW), the spectral phase is relatively unchanged for each pulse (Fig. 6.13-(a)). However, near the zero GDD point, the pulses still have enough peak power to self-focus and hence the output GDD of the central waveguide is attracted to the fixed point. As the average power is increased to 45 mW, the attractor pulls in more of the data points (Fig. 6.13-(b)). The pulses around 10 mm, however, are still too long to cause self-focusing and thus do not feel a strong attraction to the fixed point. At the highest average power (60 mW), all of the input chirps are attracted to the fixed point (Fig. 6.13-(c)). The width of the attraction region ( $\sim$ several thousand  $\text{fs}^2$ ) is currently limited by error in the FROG retrieval algorithm.

The physical mechanism behind this spectral phase attraction process is currently unclear, but is being investigated via simulation. We believe this process is not related to soliton-like pulse shaping due to the material dispersion of the waveguide array being normal. Furthermore, the spectral phase attraction occurs for either sign of input chirp, a characteristic not observed in the soliton case. A full simulation of the electric field evolution will give us insight into the dynamic evolution of the pulse in the waveguide array, which will be extremely useful for understanding the mechanism driving this process.

## 6.6 Conclusions

The characterization of the waveguide for use as a mode-locking mechanism is now complete. We have shown the waveguide array shortens the pulse significantly in a single-pass through the waveguide array. We have thoroughly investigated the effects of multi-photon absorption on the self-focusing in the

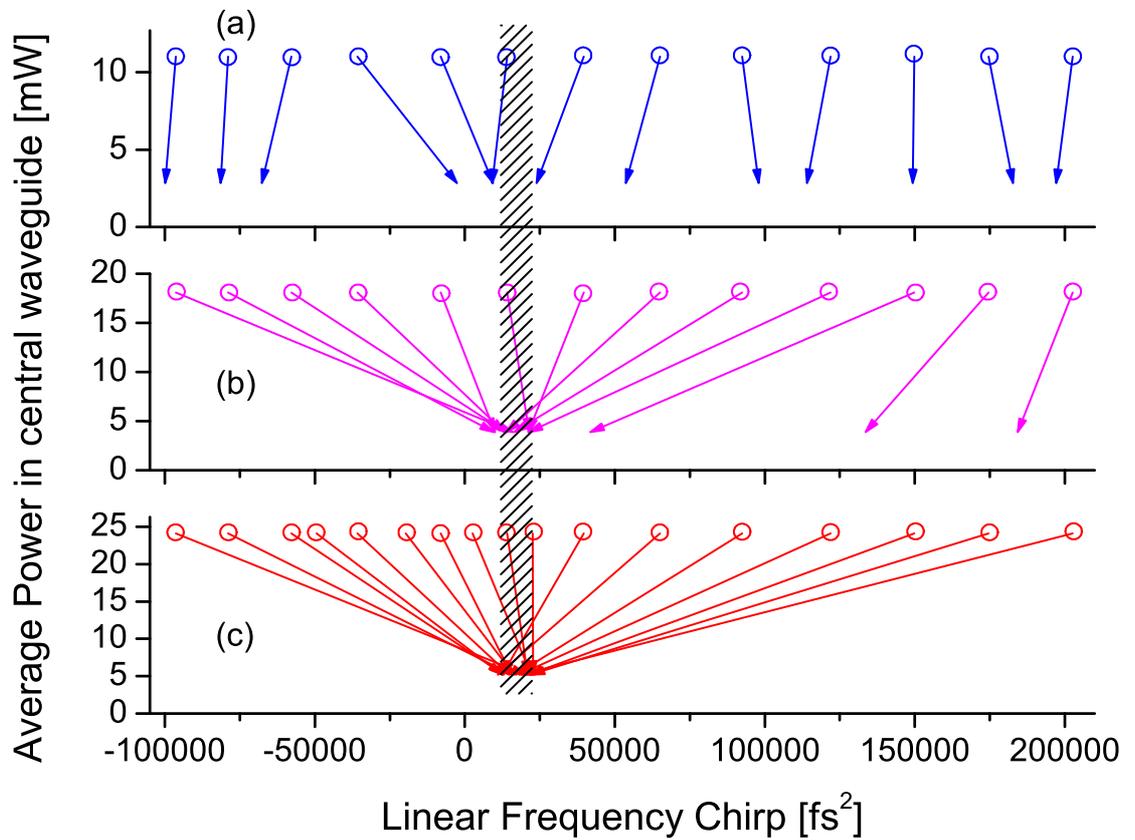


Figure 6.13: A vector plot showing the spectral phase attractor. The attraction region is denoted by the dashed lines. Part (a) shows data for the 26 mW average power case, (b) is data from the 45 mW case, and (c) is data from the 60 mW case. At high average power, all inputs go to the fixed point of  $+17,000 \text{ fs}^2$ . Recent measurements have shown that a closer spaced array (i.e. higher coupling) tends to lower the attraction region to a lower value of normal dispersion. Further theoretical study is required to get a clearer picture of exactly how the attraction region depends on coupling between waveguides.

waveguide array. Finally, we have shown that the waveguide array has a spectral phase attractor point that depends on the coupling between waveguides. In future experiments we will use all of this knowledge to build a fiber laser cavity with an intra-cavity waveguide array as the saturable absorber. This goal will require some technical skill as we will need to construct our own waveguides photolithographic techniques (see Appendix B). We are well on our way towards this goal, and have already demonstrated self-focusing in home-built devices for several arrays with different waveguide spacings. The next technical challenge will be to anti-reflection coat the input and output facets of the array. Once this is successful, we will be able to put the waveguide array in our fiber cavity and hopefully achieve mode-locking.

## Chapter 7

### Summary and Outlook

The underlying theme of this thesis is the development and application of mode-locked fiber lasers. In terms of application of these lasers, I presented a wide range of experiments we performed in which a mode-locked fiber laser was crucial: synchronization of remotely located fiber lasers, broadband molecular detection, high-field physics and high-order harmonic generation. While these projects highlight the wide range of application of this type of laser, this list of applications is still very small compared to the list of potential applications. Thus, there is quite a lot of interest in developing extremely stable, compact, and robust versions of these lasers. Toward this end, we built and successfully mode-locked a fiber laser with a fast intra-cavity actuator (EOM). This advance led to two orders of magnitude reduction in the timing jitter noise of the laser. While this laser was quite stable, it was hardly compact or robust in a mode-locking sense. In fact, one week of my life was devoted to finding the mode-locking state after a slight bump caused the laser to drop out of mode-locking.

To address the robustness issue we have turned to waveguide arrays, which theoretically should be more stable than NPR lasers. During the process of studying the waveguide arrays we discovered several interesting phenomena that we think could have many applications. The first experiment we performed showed pulse chopping in the center waveguide. This could have application in telecomm

systems where pulse broadening leads to loss of data fidelity. The second experiment we performed demonstrated a power distribution clamping effect on the discrete spatial soliton. While these results are interesting for fundamental soliton physics, we feel it could also lead to application as an amplitude-noise reducing device. In this scheme, a waveguide array placed before an amplifier would be driven into highly non-linear absorption resulting in a clamped output. This stable output could then be amplified. In the last experiment we demonstrated that the waveguide array has a spectral phase attraction point. Thus, for spectral phase the waveguide is a many-input/one-output device. The most obvious application of this would be to build a pulse compression system based on waveguide arrays. In this scenario the device would receive any initial spectral phase, convert that to the attractor point spectral phase, then compress the pulse to its Fourier limit (ignoring higher-order dispersion for the moment) by sending the pulse through a pre-calculated length of material that completely cancels the attractor point spectral phase. In this way, any input pulse with sufficiently high peak power to induce self-focusing would be transformed into the shortest pulse possible for the given bandwidth (ignoring higher-order dispersion).

Based on the results from the previous experiments there appears to be a very good chance that we can achieve a mode-locked laser based on waveguide arrays in the near future. In fact, a setup aimed at putting the waveguide array in a fiber cavity is already well underway. If such a laser based on waveguide arrays is successful, it will need a name. I would like to propose that we call it a Waveguide Array Femtosecond Fiber Laser, or WAFFL. Perhaps in the near future, a WAFFL will be in a lab near you.

## Bibliography

- [1] A. Ostendorf, G. Kamlage, and B. Chichkov, Precise deep drilling of metals by femtosecond laser pulses, *RIKEN Review: Focused on Laser Precision Microfabrication* **50**, 87 (2003).
- [2] W. Oskay *et al.*, Single-Atom Optical Clock with High Accuracy, *Phys. Rev. Lett.* **97**, 020801 (2006).
- [3] J. Ye, Absolute measurement of a long, arbitrary distance to less than an optical fringe, *Opt. Lett.* **29**, 1153 (2004).
- [4] T. Schibli, Combs for dark energy, *Nature Photonics* **2**, 712 (2008).
- [5] F. Hong, K. Minoshima, A. Onae, H. Inaba, H. Takada, A. Hirai, H. Matsumoto, T. Sugiura, and M. Yoshida, Broad-spectrum frequency comb generation and carrier-envelope offset frequency measurement by second-harmonic generation of a mode-locked fiber laser, *Opt. Lett.* **28**, 1516 (2003).
- [6] D. D. Hudson, K. W. Holman, R. J. Jones, S. T. Cundiff, and J. Ye, Mode-locked fiber laser frequency-controlled with an intracavity electro-optic modulator, *Opt. Lett.* **30**, 2948 (2005).
- [7] A. L. Cavalieri *et al.*, Clocking femtosecond X rays, *Phys. Rev. Lett.* **94**, 114801 (2005).
- [8] D. D. Hudson, S. M. Foreman, S. T. Cundiff, and J. Ye, Synchronization of mode-locked femtosecond lasers through a fiber link, *Opt. Lett.* **31**, 1951 (2006).
- [9] M. Thorpe, D. Hudson, K. Moll, J. Lasri, and J. Ye, Cavity-ringdown molecular spectroscopy based on an optical frequency comb at 1.45-1.65  $\mu\text{m}$ , *Opt. Exp.* **32**, 307 (2007).
- [10] I. Hartl, T. Schibli, A. Marcinkevicius, D. Yost, D. Hudson, M. Fermann, and J. Ye, Cavity-enhanced similariton Yb-fiber laser frequency comb:  $3 \times 10^{14}$  W/cm<sup>2</sup> peak intensity at 136 MHz, *Opt. Lett.* **32**, 2870 (2007).

- [11] D. Hudson, K. Shish, T. Schibli, J. N. Kutz, D. Christodoulides, R. Morandotti, and S. T. Cundiff, Nonlinear femtosecond pulse reshaping in waveguide arrays, *Opt. Lett.* **33**, 1440 (2008).
- [12] T. H. Maiman, Stimulated optical radiation in ruby, *Nature* **187**, 493 (1960).
- [13] L. E. Hargrove, R. L. Fork, and M. A. Pollack, Locking of HeNe laser modes induced by synchronous intracavity modulation, *Appl. Phys. Lett.* **5**, 4 (1964).
- [14] H. W. Mocker and R. J. Collins, Mode competition and self-locking effects in a q-switched ruby laser, *Appl. Phys. Lett.* **7**, 270 (1965).
- [15] E. P. Ippen, C. V. Shank, and A. Dienes, Passive mode locking of the cw dye laser, *Appl. Phys. Lett.* **21**, 348 (1972).
- [16] R. L. Fork, C. H. B. Cruz, P. C. Becker, and C. V. Shank, Compression of optical pulses to six femtoseconds by using cubic phase compensation, *Opt. Lett.* **12**, 483 (1987).
- [17] L. F. Mollenauer and R. H. Stolen, The soliton laser, *Opt. Lett.* **9**, 13 (1984).
- [18] D. E. Spence, P. N. Kean, and W. Sibbett, 60-fsec pulse generation from a self-mode-locked Ti:sapphire laser, *Opt. Lett.* **16**, 42 (1991).
- [19] R. Ell *et al.*, Generation of 5-fs pulses and octave-spanning spectra directly from a Ti:sapphire laser, *Opt. Lett.* **26**, 373 (2001).
- [20] R. H. Stolen, Nonlinearity in fiber transmission, *Proceedings of the IEEE* **68**, 1232 (1980).
- [21] H. A. Haus, J. G. Fujimoto, and E. P. Ippen, Structures for additive pulse mode locking, *J. Opt. Soc. Am. B* **8**, 2068 (1991).
- [22] E. P. Ippen, in *Ultrafast Dynamics of Quantum Systems: Physical Processes and Spectroscopic Techniques*, edited by D. Bartolo and Gambarota (Plenum Press, New York, 1998).
- [23] I. N. Duling, Subpicosecond all-fibre erbium laser, *Elec. Lett.* **27**, 544 (1991).
- [24] H. A. Haus and E. P. Ippen, Self-starting of passively mode-locked lasers, *Opt. Lett.* **17**, 1331 (1991).
- [25] D. J. Richardson, R. I. Laming, D. N. Payne, M. W. Phillips, and V. J. Matsas, 320fs soliton generation with passively mode-locked erbium fibre laser, *Elec. Lett.* **27**, 730 (1991).
- [26] K. Tamura, H. A. Haus, and E. P. Ippen, Self-starting additive pulse mode-locked erbium fibre ring laser, *Elec. Lett.* **28**, 2226 (1992).

- [27] J. Rauschenberger, T. M. Fortier, D. J. Jones, J. Ye, and C. S. T., Control of the frequency comb from a modelocked Erbium-doped fiber laser, *Opt. Exp.* **10**, 1404 (2002).
- [28] B. Washburn, R. Fox, N. Newbury, J. Nicholson, K. Feder, P. Westbrook, and C. Jrgensen, Fiber-laser-based frequency comb with a tunable repetition rate, *Opt. Exp.* **12**, 4999 (2004).
- [29] <http://www.cisco.com/en/US/prod/collateral/modules/ps5455/images/whitepaperc1146366107.jpg>.
- [30] P. C. Becker, N. A. Olsson, and J. R. Simpson, *Erbium-doped Fiber Amplifiers*, 1<sup>st</sup> ed. (Academic Press, Boston, 1999), p. 460.
- [31] D. J. Jones, S. A. Diddams, J. K. Ranka, A. Stentz, R. S. Windeler, J. L. Hall, and S. T. Cundiff, Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis, *Science* **288**, 635 (2000).
- [32] S. T. Cundiff, J. Ye, and J. Hall, Optical frequency synthesis based on mode-locked lasers, *Rev. Sci. Instrum.* **72**, 3749 (2001).
- [33] D. R. Walker, T. Udem, C. Gohle, B. Stein, and T. W. Hansch, Frequency dependence of the fixed point in a fluctuating frequency comb, *Appl. Phys. B.* **89**, 535 (2007).
- [34] T. Fortier, D. J. Jones, and S. T. Cundiff, Phase stabilization of an octave-spanning Ti:sapphire laser, *Opt. Lett.* **28**, 2198 (2003).
- [35] S. T. Cundiff and J. Ye, Colloquium: Femtosecond optical frequency combs, *Review of Modern Physics* **75**, 325 (2003).
- [36] K. Evenson, J. Wells, F. Petersen, B. Danielson, and G. Day, Accurate frequencies of molecular transitions used in laser stabilization: the 3.39- $\mu\text{m}$  transition in  $\text{CH}_4$  and the 9.33- and 10.18- $\mu\text{m}$  transitions in  $\text{CO}_2$ , *Appl. Phys. Lett.* **22**, 192 (1973).
- [37] M. J. Thorpe, Cavity-enhanced direct frequency comb spectroscopy, Ph.D. thesis, University of Colorado, 2008.
- [38] L. Enloe and J. Rodda, Laser phase-locked loop, *IEEE* **53**, 165 (1964).
- [39] T. Hansch and B. Couillaud, Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity, *Opt. Comm.* **35**, 441 (1980).
- [40] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, Laser phase and frequency stabilization using an optical resonator, *Appl. Phys. B* **31**, 97 (1983).
- [41] H. Black, Wave Translation System, 1937.

- [42] A. D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye, Compact, thermal-noise-limited optical cavity for diode laser stabilization at  $1 \times 10^{-15}$ , *Opt. Lett.* **32**, 641 (2007).
- [43] S. M. Foreman, Femtosecond Frequency Combs for Optical Clocks and Timing Transfer, Ph.D. thesis, JILA and the Department of Physics, University of Colorado, Boulder, CO, 2007.
- [44] A. Hati, D. A. Howe, F. L. Walls, and D. Walker, *Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum* (IEEE, Piscataway, New Jersey, 2003), pp. 516–520.
- [45] P. J. Winzer, Shot-noise formula for time-varying photon rates: a general derivation, *J. Opt. Soc. Am. B* **14**, 2424 (1997).
- [46] E. D. Black, An introduction to Pound-Drever-Hall laser frequency stabilization, *Am. J. Phys.* **69**, 79 (2001).
- [47] E. M. Lally, Master’s thesis, Virginia Polytechnic Institute and State University, 2006.
- [48] D. W. Allan, Statistics of atomic frequency standards, *Proc. of the IEEE* **54**, 221 (1966).
- [49] J. L. Hall and M. Zhu, in *Laser Manipulation of Atoms and Ions, Proc. of the International School of Phys. Enrico Fermi: Course 118*, edited by E. Arimondo, W. D. Phillips, and F. Strumia (North Holland, Amsterdam, 1992), p. 671.
- [50] R. K. Shelton, S. M. Foreman, L. S. Ma, J. L. Hall, H. C. Kapteyn, M. M. Murnane, M. Notcutt, and J. Ye, Subfemtosecond timing jitter between two independent, actively synchronized, mode-locked lasers, *Opt. Lett.* **27**, 312 (2002).
- [51] A. Bartels, S. A. Diddams, T. M. Ramond, and L. Hollberg, Mode-locked laser pulse trains with subfemtosecond timing jitter synchronized to an optical reference oscillator, *Opt. Lett.* **28**, 663 (2003).
- [52] T. R. Schibli, J. Kim, O. Kuzucu, J. Gopinath, S. N. Tandon, G. S. Petrich, L. A. Kolodziejski, J. G. Fujimoto, E. P. Ippen, and F. X. Kaertner, Attosecond active synchronization of passively mode-locked lasers by balanced cross correlation, *Opt. Lett.* **28**, 947 (2003).
- [53] B. Shillue, S. AlBanna, and L. D’Addario, *2004 IEEE International Topical Meeting on Microwave Photonics Technical Digest* (IEEE, Piscataway, NJ, 2004).

- [54] ALMA/ESO/NRAO/NAOJ.
- [55] R. W. Schoenlein, W. P. Leemans, A. H. Chin, P. Volfbeyn, T. E. Glover, P. Balling, M. Zolotarev, K. J. Kim, S. Chattopadhyay, and C. V. Shank, Femtosecond X-ray Pulses at 0.4 Å generated by 90° Thomson Scattering: A Tool for Probing the Structural Dynamics of Materials, *Science* **274**, 236 (1996).
- [56] M. F. DeCamp *et al.*, Coherent control of pulsed X-ray beams, *Nature* **413**, 825 (2001).
- [57] LCLS/SLAC.
- [58] J. P. Gordon and H. A. Haus, Random walk of coherently amplified solitons in optical fiber transmission, *Opt. Lett.* **11**, 665 (1986).
- [59] P. Kubina, P. Adel, F. Adler, G. Grosche, T. W. Hansch, R. Holzwarth, A. Leitenstorfer, B. Lipphardt, and H. Schnatz, Long term comparison of two fiber based frequency comb systems, *Opt. Exp.* **13**, 904 (2005).
- [60] E. Hecht, in *Optics*, 4 ed., edited by A. Black (Addison Wesley, Boston, 2002), p. 368.
- [61] L. A. Jiang, M. E. Grein, H. A. Haus, E. P. Ippen, and H. Yokoyama, Timing jitter eater for optical pulse trains, *Opt. Lett.* **28**, 78 (2003).
- [62] L. E. Nelson, D. J. Jones, K. R. Tamura, H. A. Haus, and E. P. Ippen, Ultrashort-pulse fiber ring lasers, *Appl. Phys. B* **65**, 277 (1997).
- [63] R. K. Shelton, L. S. Ma, H. C. Kapteyn, M. M. Murnane, J. L. Hall, and J. Ye, Phase-coherent optical pulse synthesis from separate femtosecond lasers, *Science* **293**, 1286 (2001).
- [64] A. C. Yu, X. Ye, D. Ionascu, W. X. Cao, and P. M. Champion, Two-color pump-probe laser spectroscopy instrument with picosecond time-resolved electronic delay and extended scan range, *Rev. Sci. Instrum.* **76**, 114301 (2005).
- [65] H. d. Riedmatten, I. Marcikic, W. Tittel, H. Zbinden, D. Collins, and N. Gisin, Long distance quantum teleportation in a quantum relay configuration, *Phys. Rev. Lett.* **92**, 047904 (2004).
- [66] BRAN - Boulder Research and Administration Network <http://www.branfiber.net>.
- [67] K. W. Holman, D. J. Jones, D. D. Hudson, and J. Ye, Precise frequency transfer through a fiber network by use of 1.5- $\mu\text{m}$  mode-locked sources, *Opt. Lett.* **29**, 1554 (2004).

- [68] K. W. Holman, D. D. Hudson, J. Ye, and D. J. Jones, Remote transfer of a high-stability and ultralow-jitter timing signal, *Opt. Lett.* **30**, 1225 (2005).
- [69] J. Ye *et al.*, Delivery of high-stability optical and microwave frequency standards over an optical fiber network, *J. Opt. Soc. Am. B* **20**, 1459 (2003).
- [70] E. Crosson *et al.*, Stable Isotope Ratios Using Cavity Ring-Down Spectroscopy: Determination of  $^{13}\text{C}/^{12}\text{C}$  for Carbon Dioxide in Human Breath, *Anal. Chem.* **74**, 2003 (2002).
- [71] T. Risby and S. F. Solga, Current status of clinical breath analysis, *Appl. Phys. B* **85**, 421 (2006).
- [72] A. O’Keefe and D. A. G. Deacon, Cavity ring-down optical spectrometer for absorption measurements using pulsed laser sources, *Rev. Sci. Instrum.* **59**, 2544 (1988).
- [73] J. Ye and J. Hall, Cavity ringdown heterodyne spectroscopy: High sensitivity with microwatt light power, *Phys. Rev. A* **61**, 061802 (2000).
- [74] S. M. J. Kelly, Characteristic sideband instability of periodically amplified average soliton, *Elec. Lett.* **28**, 806 (1992).
- [75] S. Xiao and A. Wiener, 2-D wavelength demultiplexer with potential for  $\geq 1000$  channels in the C-band, *Opt. Exp.* **12**, 2895 (2004).
- [76] R. J. Jones, K. Moll, M. Thorpe, and J. Ye, Phase-coherent Frequency Combs in the Vacuum Ultraviolet via High-Harmonic Generation inside a Femtosecond Enhancement Cavity, *Phys. Rev. Lett.* **94**, 193901 (2005).
- [77] C. Gohle, T. Udem, M. Herrmann, J. Rauschenberger, R. Holzwarth, H. Schuessler, F. Krausz, and T. Hansch, A frequency comb in the extreme ultraviolet, *Nature* **436**, 234 (2005).
- [78] P. Corkum, Plasma perspective on strong field multiphoton ionization, *Phys. Rev. Lett.* **71**, 1994 (1993).
- [79] M. Hentschel, R. Kienberger, C. Spielman, G. Rieder, N. Milosevic, T. Brabec, P. Corkum, U. Heinzmann, M. Drescher, and F. Krausz, Attosecond metrology, *Nature* **414**, 509 (2001).
- [80] J. Itatani, J. Levesque, D. Zeidler, H. Niikura, H. Pepin, J. Kieffer, P. Corkum, and D. Villeneuve, Tomographic imaging of molecular orbitals, *Nature* **432**, 867 (2004).
- [81] M. Ferray, A. L’Huillier, X. Li, L. Lomprk, G. Mainfray, and C. Manus, Multiple-harmonic conversion of 1064 nm radiation in rare gases, *J. Phys. B: At. Mol. Opt. Phys.* **21**, L31 (1988).

- [82] P. Corkum, N. Burnett, and F. Brunel, Above-threshold ionization in the long-wavelength limit, *Phys. Rev. Lett.* **62**, 1259 (1989).
- [83] M. Lewenstein, P. Balcou, M. Ivanov, A. L'Huillier, and P. Corkum, Theory of high-harmonic generation by low-frequency laser fields, *Phys. Rev. A* **49**, 2117 (1994).
- [84] M. Fermann, A. Galvanauskas, and G. Sucha, *Ultrafast Lasers: Technology and Applications* (Marcel Dekker, New York, 2002).
- [85] F. O. Ilday, J. R. Buckley, W. G. Clark, and F. W. Wise, Self-Similar Evolution of Parabolic Pulses in a Laser, *Phys. Rev. Lett.* **92**, 213902 (2004).
- [86] A. L'Huillier, L. A. Lompre, G. Mainfray, and C. Manus, Multiply charged ions induced by multiphoton absorption processes in rare-gas atoms at 1.064  $\mu\text{m}$ , *J. Phys. B* **16**, 1363 (1983).
- [87] T. Schibli, I. Hartl, D. Yost, M. Martin, A. Marcinkevicius, M. Fermann, and J. Ye, Optical frequency comb with submillihertz linewidth and more than 10 W average power, *Nature Photonics* **2**, 355 (2008).
- [88] D. Yost, T. Schibli, and J. Ye, Efficient output coupling of intracavity high-harmonic generation, *Opt. Lett.* **33**, 1099 (2008).
- [89] J. Proctor and J. Kutz, Passive mode-locking by use of waveguide arrays, *Opt. Lett.* **30**, 2013 (2005).
- [90] H. S. Eisenberg, Y. Silberberg, R. Morandotti, A. R. Boyd, and J. S. Aitchison, Discrete spatial optical solitons in waveguide arrays, *Phys. Rev. Lett.* **81**, 3383 (1998).
- [91] H. Furuta, H. Noda, and A. Ihaya, Novel Optical Waveguide for Integrated Optics, *Appl. Opt.* **13**, 322 (1974).
- [92] J. U. Kang, A. Villeneuve, M. Sheik-Bahae, G. I. Stegeman, K. Al-hemyari, J. S. Aitchison, and C. N. Ironside, Limitation due to three-photon absorption on the useful spectral range for nonlinear optics in AlGaAs below half band gap, *Appl. Phys. Lett.* **65**, 147 (1994).
- [93] D. N. Christodoulides and R. I. Joseph, Discrete self-focusing in nonlinear arrays of coupled waveguides, *Opt. Lett.* **13**, 794 (1988).
- [94] J. Meier, G. I. Stegeman, D. N. Christodoulides, Y. Silberberg, R. Morandotti, H. Yang, G. Salamo, M. Sorel, and J. S. Aitchison, Experimental observation of discrete modulational instability, *Phys. Rev. Lett.* **92**, 163902 (2004).

- [95] S. Suntsov, K. G. Makris, D. N. Christodoulides, G. I. Stegeman, A. Hach, R. Morandotti, H. Yang, G. Salamo, and M. Sorel, Observation of discrete surface solitons, *Phys. Rev. Lett.* **96**, 063901 (2006).
- [96] S. Droulias, K. Hizanidis, D. N. Christodoulides, and R. Morandotti, Waveguide array-grating compressors, *App. Phys. Lett.* **87**, 131104 (2005).
- [97] D. Marcuse, *Theory of dielectric optical waveguides*, 2nd ed. (Academic Press, San Diego, 1991).
- [98] A. B. Aceves, C. De Angelis, T. Peschel, R. Muschall, F. Lederer, S. Trillo, and S. Wabnitz, Discrete self-trapping, soliton interactions, and beam steering in nonlinear waveguide arrays, *Phys. Rev. E* **53**, 1172 (1996).
- [99] H. S. Eisenberg, R. Morandotti, Y. Silberberg, J. M. Arnold, G. Pennelli, and J. S. Aitchison, Optical discrete solitons in waveguide arrays. I. Soliton formation, *J. Opt. Soc. Am. B* **19**, 2938 (2002).
- [100] U. Peschel, R. Morandotti, J. M. Arnold, J. S. Aitchison, H. S. Eisenberg, Y. Silberberg, T. Pertsch, and F. Lederer, Optical discrete solitons in waveguide arrays. 2. Dynamics properties, *J. Opt. Soc. Am. B* **19**, 2637 (2002).
- [101] F. Tauser, A. Leitenstorfer, and W. Zinth, Amplified femtosecond pulses from an Er: fiber system: Nonlinear pulse shortening and self-referencing detection of the carrier-envelope phase evolution, *Opt. Exp.* **11**, 594 (2003).
- [102] E. B. Treacy, Optical pulse compression with diffraction gratings, *IEEE J. of Quant. Elec.* **5**, 454 (1969).
- [103] K. Naganuma, K. Mogi, and H. Yamada, Group-delay measurement using the Fourier transform of an interferometric cross correlation generated by white light, *Opt. Lett.* **15**, 393 (1990).
- [104] S. Diddams and J.-C. Diels, Dispersion measurements with white-light interferometry, *J. Opt. Soc. Am. B* **13**, 1120 (1996).
- [105] R. Morandotti, H. S. Eisenberg, Y. Silberberg, M. Sorel, and J. S. Aitchison, Self-Focusing and Defocusing in Waveguide Arrays, *Phys. Rev. Lett.* **86**, 3296 (2001).
- [106] S. T. Cundiff, Coherent spectroscopy of semiconductors, *Opt. Exp.* **16**, 4639 (2008).
- [107] D. Kane and R. Trebino, Characterization of arbitrary femtosecond pulses using frequency-resolved optical gating, *IEEE J. of Quant. Elec.* **29**, 571 (1993).

## Appendix A

### Matlab codes

#### A.1 A simple model of a mode-locked laser

The following code is a simple model for a mode-locked laser. This model adds up the electric field modes of a cavity (note cavity parameters can be set in the declaration section of the code), and computes the total intensity pattern. The mode-locking process is modeled by adding a phase term to each cavity mode. The user can choose to have an arbitrary amount of phase noise (set by the 'factor' parameter) between comb components. A factor equal to 0 yields a perfect comb structure, while a factor  $\geq 2\pi$  yields random intensity fluctuations.

```
Em = 1;           % Electric field amplitude
factor = 0;       % noise factor (0 = phase locked, >> 0 large phase noise)
c = 3e8;          % speed of light
L = 1.5;          % cavity length
FSR = c/(2*L);    % free spectral range of the cavity
t = -27e-9:1e-12:27e-9; % go from -27 ns to +27 ns in steps of 1 ps.
total = 0;
Upper = 1.9355e14+80*FSR; % Upper frequency limit:
                        % c/lambda + % (80 comb modes * FSR)
for f = 1.9355e14:FSR:Upper % go from fundamental laser frequency to
```

```

                                % higher freqs by steps of 1 FSR
    phase = rand*factor;          % calculate the phase for the mth mode
    E = Em*cos(2*pi*f*t+phase); % calculate electric field for mth mode
    total = total+E;

end

time = t./(1e-9);
plot(time,abs(total).^2/max(abs(total).^2))
xlim([-27 27])

```

## A.2 Pound Drever Hall simulation

The following declarations define the modulation depth, carrier power, side-band power cavity free-spectral range, reflection amplitude, and reflection coefficient, respectively.

---


$$\beta := 1.08$$

$$P_c := N [\text{BesselJ}[0, \beta]^2]$$

$$P_s := N [\text{BesselJ}[1, \beta]^2]$$

$$\text{FSR} := \frac{(3 \cdot 10^8)}{2 \cdot 0.5}$$

$$r := .995$$

$$F[\omega_-] := \frac{r \left( \text{Exp} \left[ \frac{i \cdot \omega}{\text{FSR}} \right] - 1 \right)}{1 - r^2 \cdot \text{Exp} \left[ \frac{i \cdot \omega}{\text{FSR}} \right]}$$


---

The following line allows the user to manipulate the modulation frequency real-time. In the small modulation case, the error signal looks like the derivative

of a resonance. In the large modulation case we get the classic PDH error signal:

---

```

Manipulate [Plot [(-2 * sqrt[Pc * Ps] * Im [(F [omega * 10^9] * Conjugate [F [omega * 10^9 + Omega * 10^6]]
-Conjugate [F [omega * 10^9]] * F [omega * 10^9 - Omega * 10^6]])]/0.685, {omega, -.200, .200},
PlotRange -> {-1, 1}, AxesLabel -> {(omega - omega_0) [GHz], ErrorSignal},
TextStyle -> {FontSize -> 23}, PlotStyle -> {Thickness[0.005]}, {Omega, 1, 120}]

```

---

Fig. A.1 shows the PDH error signal when the modulation frequency is very low (i.e.  $\Omega \sim 3\text{MHz}$ ). For comparison, the figure in Ch. 2 is the error signal for a modulation frequency of 110 MHz.

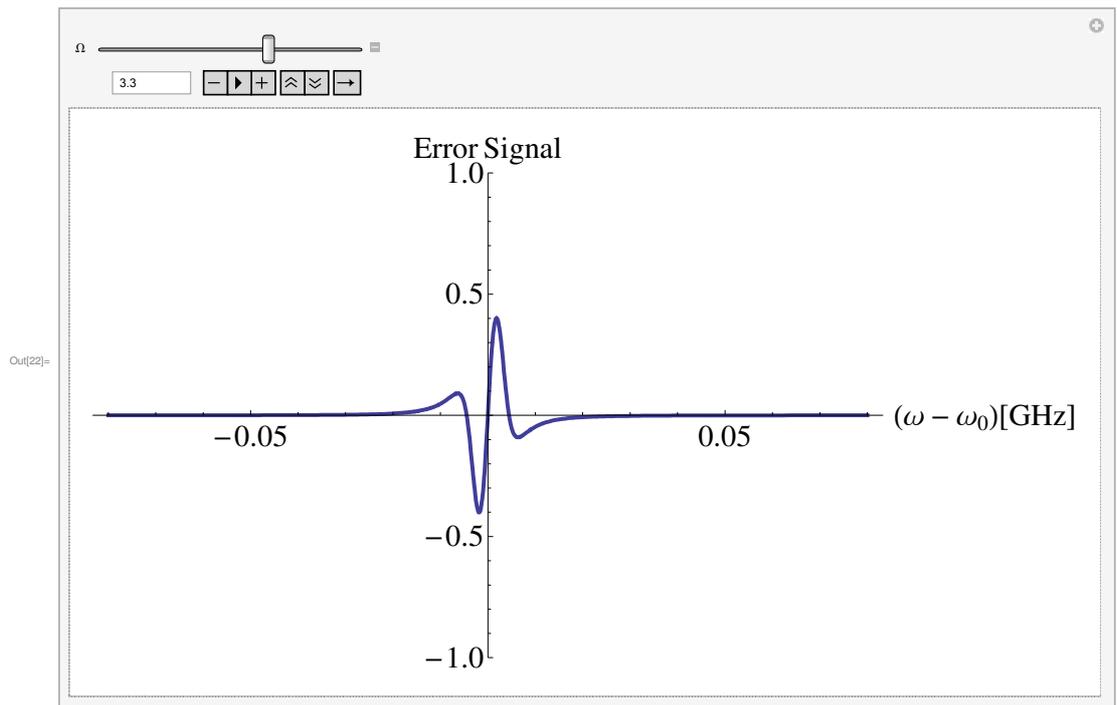


Figure A.1: Simulation output for low modulation frequency (PDH).

## Appendix B

### Fabricating Waveguide Arrays

This appendix discusses the fabrication of the waveguide arrays. While much of the data in this thesis was recorded using a waveguide array built outside of JILA, the next generation of waveguide arrays used in the Cundiff lab will be homemade. One of the nice properties of AlGaAs is that it lends itself to processing via photolithography. To achieve waveguiding in the vertical dimension, the AlGaAs wafer must be grown with varying layers of alloy concentration. The recipe that appears to be best at minimizing multi-photon absorption is a core layer of  $\text{Al}_{0.18}\text{Ga}_{0.82}\text{As}$  with cladding layers of  $\text{Al}_{0.24}\text{Ga}_{0.76}\text{As}$ . For lateral confinement, photolithography is used to write ridges into the top cladding layer. It is the purpose of this Appendix to describe the writing of these ridges into the wafer.

#### B.1 Growing the wafer: Molecular Beam Epitaxy

The wafer is typically grown via Molecular Beam Epitaxy (MBE), a process invented in the 1960s at Bell Laboratories. In this process a heated crystalline substrate (i.e. GaAs) is fixed in a high vacuum chamber, and effusion cells with mechanical shutters are arranged around the substrate. The shutters can be opened and closed to introduce the Aluminum, Gallium, and Arsenide atoms to the chamber. The atoms, which have a large mean free path in the high vacuum chamber and a large GaAs sticking coefficient, then stick to the surface of the

substrate. A well controlled machine with a good vacuum can achieve mono-layer precision using this process. Thus, the vertical guiding layers in our wafer can be grown to exact thickness specification.

## B.2 Processing the wafer

Once the wafer is grown, the waveguide array has to be written into the wafer using photolithography. To achieve nice ridge profiles, each step in the following process should be followed carefully. The first step is to cleave off a smaller piece of the AlGaAs wafer from the main wafer. The wafer should be grown in the (001) direction, meaning that the (001) vector is pointing up out of the top of the wafer. The orthogonal directions (110) and (1 $\bar{1}$ 0) will cleave nicely, leaving a mirror like finish on the cleaved surface. To help understand how these vectors relate to the actual crystal orientation, we can use 3D models as seen in the following figures. The following models show the lattice for GaAs (i.e. two different atoms). AlGaAs is basically the same lattice, just with aluminum atoms randomly placed at some of the nodes.

AlGaAs/GaAs is a form of the so-called *zinc blende* structure. These crystals belong to the face-centered cubic lattice group. The orthogonal cleave directions, (110) and (1 $\bar{1}$ 0), are shown on the right hand side of Fig. B.1. Typically, if the wafer is grown in the (001) direction, the major flat of the wafer will be either (110) or (1 $\bar{1}$ 0), while the minor flat will indicate the other direction. As will be seen in the next section, these two directions have quite different etch properties, so it is important to know the orientation of your wafer. A larger view of the crystal is shown in Fig. B.2, the orientation of the bonds becomes clearer in this picture as you can see many unit cells together.

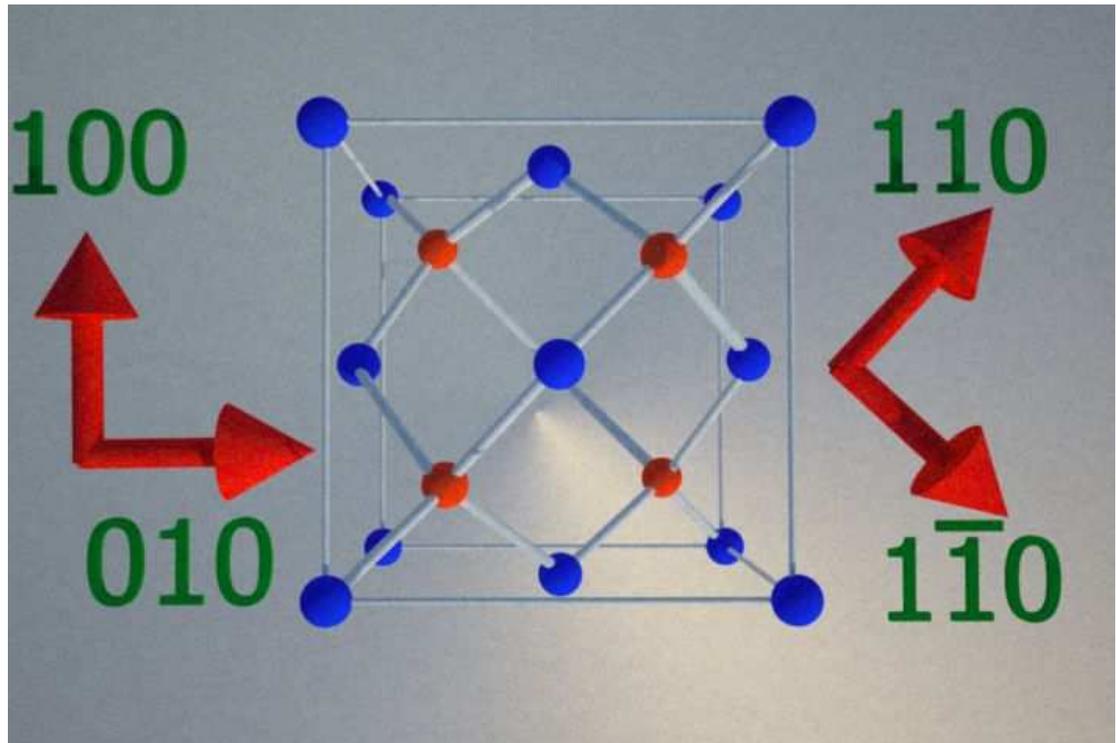


Figure B.1: Model of the lattice orientation of GaAs (top down view). The orthogonal vectors for cubic symmetry are shown on the left (with the (001) direction out of the page). The natural axis of the crystal is rotated  $45^\circ$  from the cubic lattice vectors. These natural basis vectors are shown on the right.

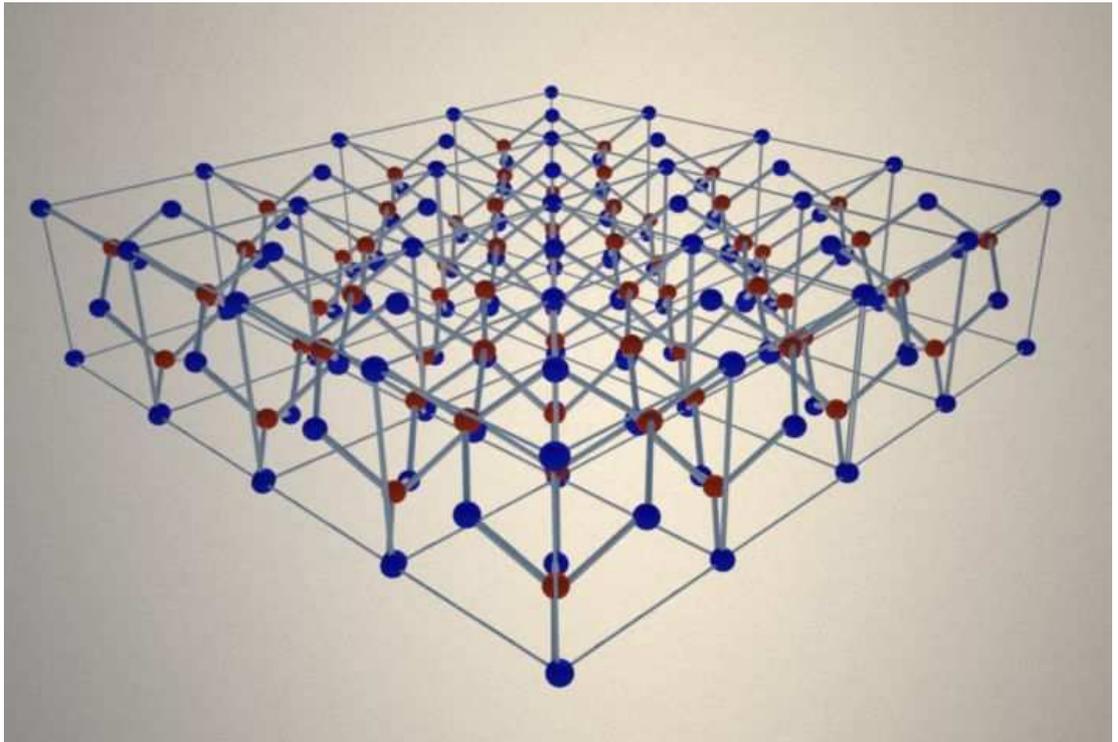


Figure B.2: Large view of the GaAs crystal. Looking down the 110 direction.

### B.3 Photolithography

Once the crystal is cleaved nicely, a standard degreasing procedure must be performed. A typical degreasing process is: put the cleaved chip in an Acetone soak (5 minutes), then Methonal soak (5 minutes), then de-ionized (DI) water soak, then rinse chip under free flowing DI water, then spin-dry the wafer using the photoresist spinner device. Once this process is complete, the chip can be inspected for specs of dirt. If the surface appears to be free of any particles, then the chip can then be baked at high temperature (i.e. 600 C). This removes any moisture that is adsorbed by the surface of the AlGaAs.

The photoresist can then be spun onto the surface using the spin-coating device. Using the Shipley S1813 positive photoresist, the spinner should rotate at 3000 rpm for 45 seconds, producing a  $1.3 \mu\text{m}$  layer of photoresist on the chip surface (see Fig. B.3-(a)). Once the chip is coated properly, a soft bake should be performed to drive away the solvent from the photoresist. This soft bake also helps improve adhesion of the resist to the chip surface, and anneals the shear stresses caused by the spin-coating. Typically, soft bakes times are on the order of 1 minute.

The next step is to write a pattern in the photoresist using UV light (see Fig. B.4-(b)). The photoresist polymers are 1000's of molecules long. When exposed to light, however, the polymers are broken up into chains of 10's of molecules and are thus much weaker. The developer essentially lifts off the photoresist at different rates, with the weak polymer sections lifting off at a faster rate. The chip should be put in the developer for around 40s to allow the exposed parts of the photoresist to be lifted off completely.

After the device has reached the stage shown in Fig. B.4-(b), it is ready for the *hard bake*. This bake hardens the remaining photoresist.

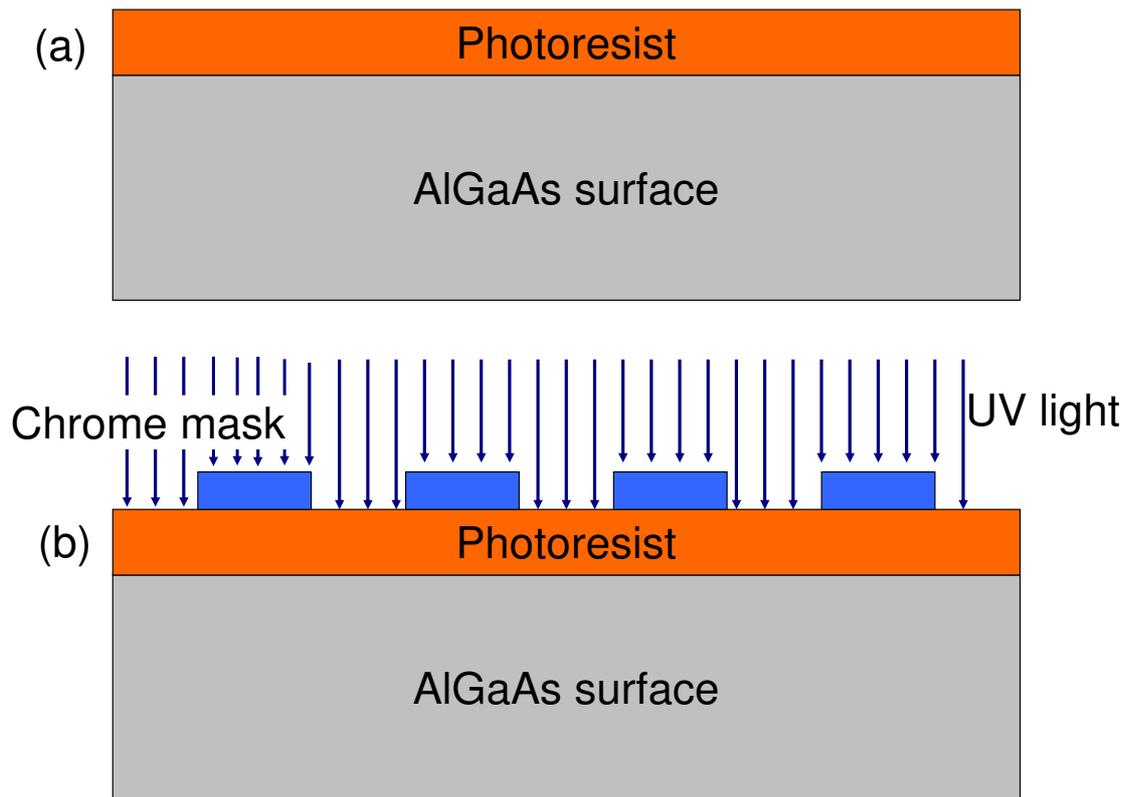


Figure B.3: Schematic of the first step of patterning. (a) The photoresist is spun onto the wafer, achieving a thickness of  $1.3\mu\text{m}$ . (b) Patterning the photoresist. The chrome mask is brought into soft contact with the photoresist, then the UV light is fired and the exposed polymers of the photoresist are broken into smaller chains of molecules.

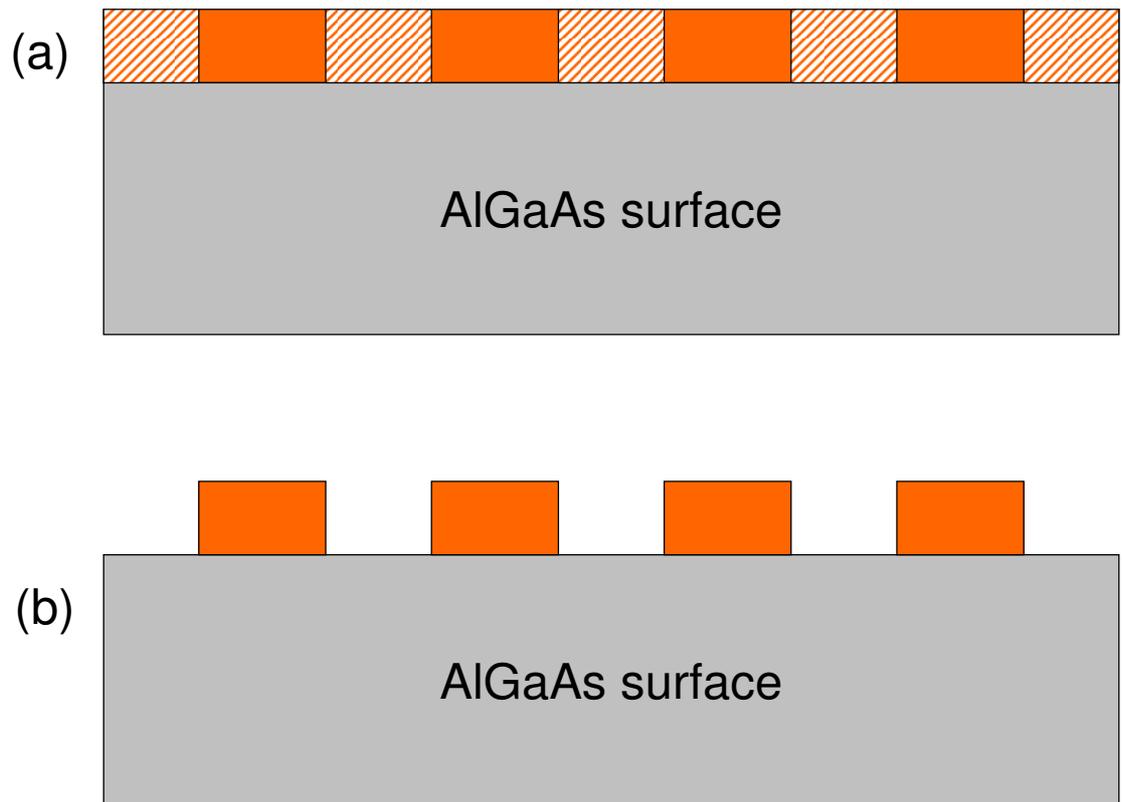


Figure B.4: Schematic of the second step of patterning. (a) The exposed parts of the photoresist are then washed away using the developer chemical. (b) After developing, the remaining photoresist undergoes a hard bake to harden the material for the etching process.

## B.4 Etching

The last step in this whole process is wet etching of the AlGaAs (see Fig. B.5). While it is possible (and perhaps even desirable) to etch the waveguides using a reactive-ion etch (RIE) machine, a wet acid etch seems to work well enough and is much less complicated. The wet etching formula that seemed to work well for the devices made here at JILA was: 10 parts  $\text{H}_2\text{O}$ , 1 part  $\text{H}_2\text{O}_2$ , and 1 part  $\text{H}_2\text{SO}_4$  (Sulfuric Acid). An etch time of 60 seconds resulted in good coupling between adjacent waveguides.

The etch direction has a significant impact on the waveguide profile. While an RIE etch would in principle give a near vertical sidewall profile for etching along either direction, the wet etch profile is very dependent on the etch direction. Etching along  $(110)$  direction gives a high aspect ratio sidewall profile (see Fig. B.7), while etching along the  $(1\bar{1}0)$  direction gives an aspect ratio much closer to one (see Fig. B.6).

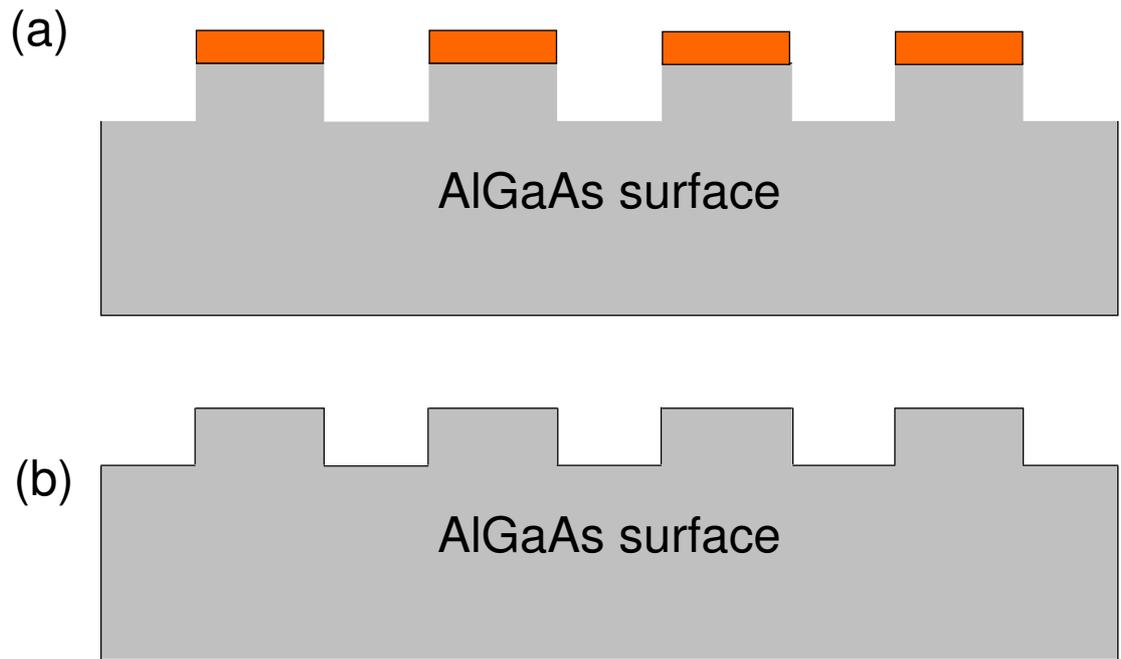


Figure B.5: Schematic of the last step of patterning. (a) Post-etch profile of the waveguides. The wet etch eats away the AlGaAs and the photoresist at different rates. The ideal situation (as shown) is to use an etch time that gives a nice etch depth while not fully eating through the photoresist. (b) Final device profile. After the wet etching, an organic solvent such as Acetone can be used to remove the remaining photoresist.

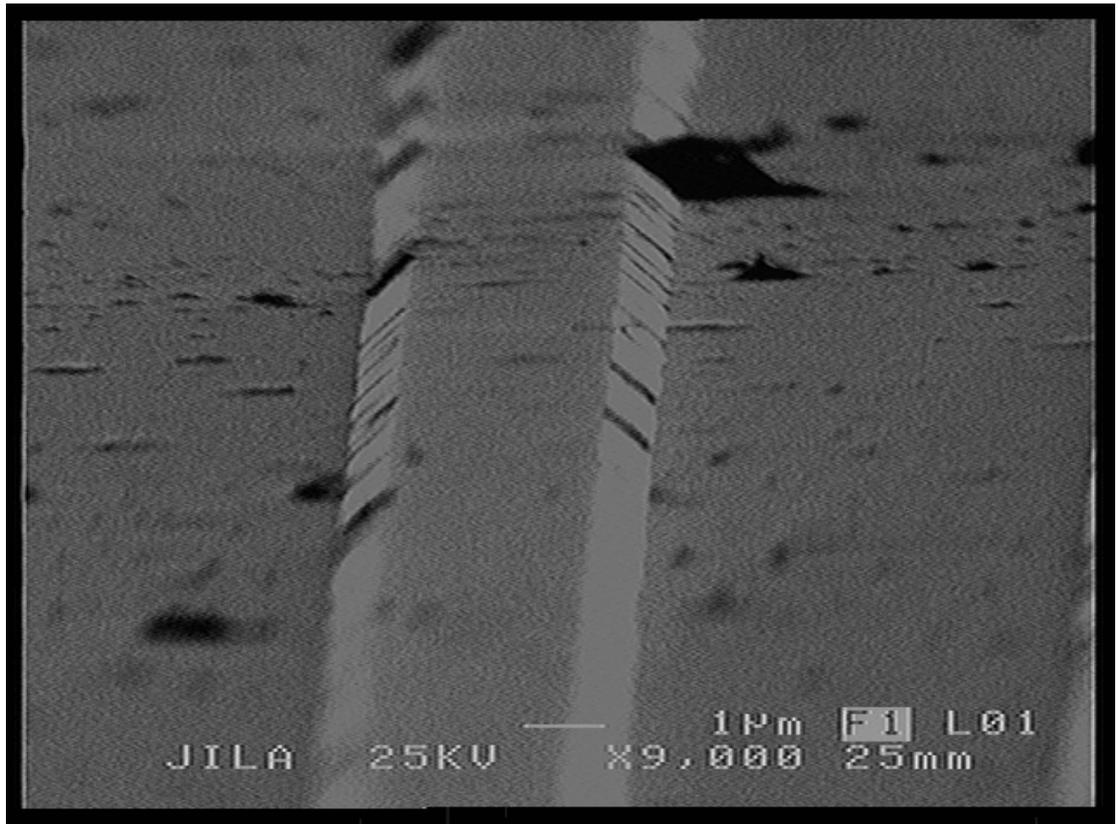


Figure B.6: Etching along the  $(1\bar{1}0)$  direction in AlGaAs. This direction gives a v-groove type sidewall profile.

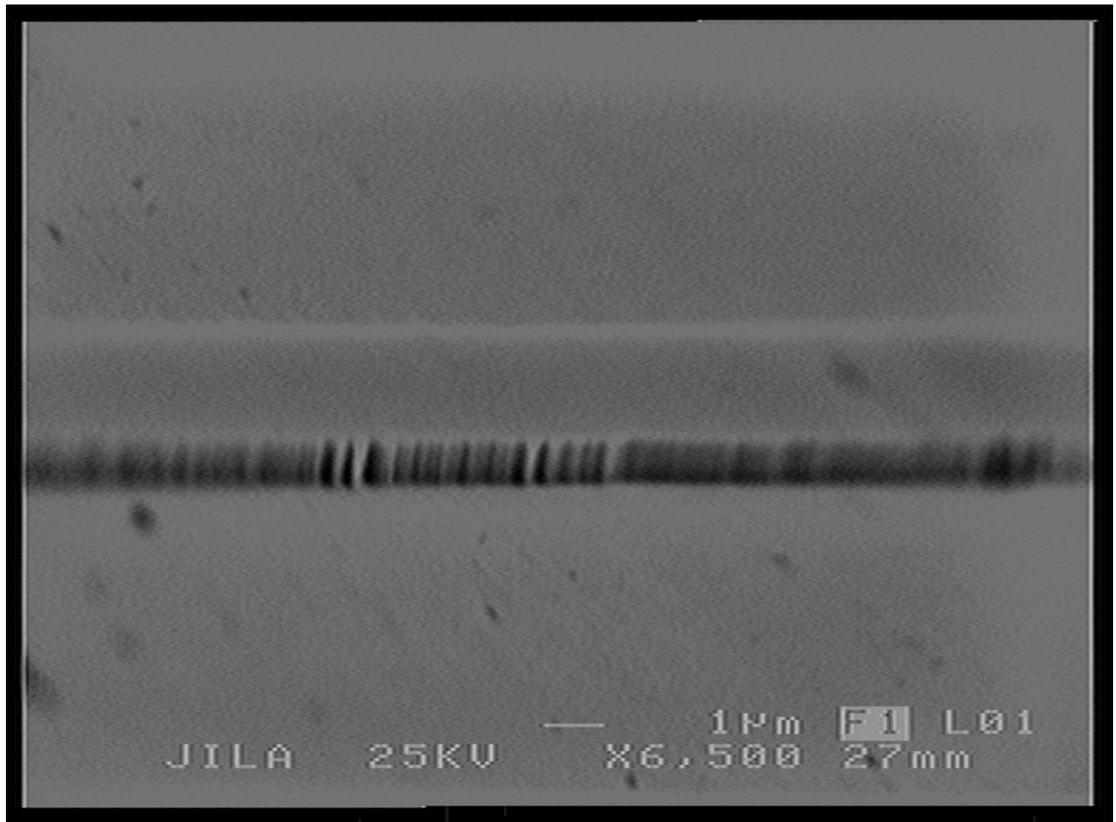


Figure B.7: Etching along the (110) direction in AlGaAs. This direction gives a nice sidewall profile. This is the etch direction that is used in our fabrication.