Practical Design and Applications of Ultrafast Semiconductor Disk Lasers

by
Caleb W. Baker

A Dissertation Submitted to the Faculty of the
College of Optical Sciences
In Partial Fulfillment of the Requirements
For the Degree of
Doctor of Philosophy
In the Graduate College
The University of Arizona

2 0 1 7

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Caleb W. Baker entitled:

Practical Design and Applications of Ultrafast Semiconductor Disk Lasers

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

_____ Date: August 21st, 2017
Jerome V. Moloney

_____ Date: August 21st, 2017
R. Jason Jones

_____ Date: August 21st, 2017
Robert A. Norwood

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

_____ Date: August 21st, 2017
Jerome V. Moloney

Statement by Author

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

Signed: Caleb W. Baker

Acknowledgements

I would like to express my deepest gratitude to our research professors, Dr. Maik Scheller and Dr. Alexandre Laurain, to our collaborator, Prof. R. Jason Jones, and to my advisor, Prof. Jerome V. Moloney. Their guidance, assistance, and mentorship have been invaluable to me throughout the years I have spent at the university.

I would also like to thank the other current and former members of our VECSEL group during my time here, Prof. Stephan W. Koch, Dr. Jorg Hader, Dr. Cody Mart, Isak Kilen, Hwang-Jye Yang, Hsiu-ting Chan, and Robert Rockmore, for their many and varied contributions to the work presented here.

I would like to acknowledge our collaborators at Philipps-Universitat Marburg, Dr. Wolfgang Stolz and Antje Ruiz-Perez, who grow our VECSEL devices and our collaborators at University of New Mexico, Dr. Ganesh Balakrishnan and Dr. Sadhvikas Addamane, who grow our saturable absorbers.

A special thank you to Prof. Robert Norwood for his involvement in my dissertation committee and his feedback and advice during the preparation of this dissertation. Thank you, as well, to Dr. Paris Panagiotopoulos for his help editing certain portions of this dissertation.

I would like to thank the administrative staff at the College of Optical Sciences, Mark Rodriguez, Ruth Corcoran, Hector Garcia, Melissa Ayala, Jennifer Garcia, and Lindsay Loebig, for their tireless support of all of us students.

This work has been supported by multiple grants from the Air Force Office of Scientific Research: FA9550-10-1-0064 ("Ultrafast Nonlinear Optics on Macroscopic and Sub-Wavelength Scales"), FA9550-14-1-0062 ("Ultrafast Light COupling to Nonequilibrium Carriers in Extended Semiconductor Media"), and FA9550-17-1-0246 ("Nonequilibrium Semiconductor Disk Laser Theory and Simulation: Robust, Novel Multi-Functional Source for Femtosecond Pulsed and Multi-Wavelength CW Generation").

Finally, I would like to thank all of my family and friends for the support they have given me during my time in graduate school.

Dedication

This thesis is dedicated to the following people, all of whom had the opportunity to shape a developing mind and did not squander it.

—

Mrs. Barbara Frasier
Mrs. Patricia Beatty
Mr. David Carter
Mr. Lee Derr
Mr. Steve Rischling
Mr. Scott Davis
Mr. Samuel Milazzo
Dr. Bill Kiele
Dr. John Jackson
Dr. Jennifer Clarke
Dr. Sandy Wurtele
Dr. Zbigniew Celinski

—

Table of Contents

Table of Contents—*Continued*

TABLE OF CONTENTS—*Continued*

List of Figures

List of Figures—*Continued*

List of Figures—*Continued*

List of Figures—*Continued*

Abstract

Vertical External Cavity Surface Emitting Lasers (VECSELs) have become well established in recent years for their design flexibility and promising power scalability. Recent efforts in VECSEL development have focused heavily on expanding the medium into the ultrafast regime of modelocked operation. Presented in this thesis is a detailed discussion regarding the development of ultrafast VECSEL devices. Achievements in continuous wave (CW) operation will be highlighted, followed by several chapters detailing the engineering challenges and design solutions which enable modelocked operation of VECSELs in the ultrafast regime, including the design of the saturable absorbers used to enforce modelocking, management of the net group delay dispersion (GDD) inside the cavity, and the design of the active region to support pulse durations on the order of 100 fs. Work involving specific applications - VECSELs emitting on multiple wavelengths simultaneously and the use of VECSEL seed oscillators for amplification and spectral broadening - will also be presented.

Key experimental results will include a novel multi-fold cavity design that produced record-setting peak powers of 6.3 kW from a modelocked VECSEL, an octave-spanning supercontinuum with an average power of 2 W generated using a VECSEL seed and a 2-stage Yb fiber amplifier, and two separate experiments where a VECSEL was made to emit on multiple wavelengths simultaneously in modelocked and highly stable CW operation, respectively. Further, many diagnostic and characterization measurements will be presented, most notably the in-situ probing of a VECSEL gain medium during stable modelocked operation with temporal resolution on the order of 100 fs, but also including characterization of the relaxation rates in different saturable absorber designs and the effectiveness of different methods for managing the net GDD of a device.

Chapter 1

Introduction

The first functional laser was built in 1960 by Theodore Maiman. It was based on a ruby crystal and by today's standards was a fairly disappointing laser. Charles Townes, one of the fathers of the theoretical foundations for lasers, recounted being disappointed that, while Maiman had definitely proven the existence of lasing, he did not report a "bright beam of light," one of the most expected characteristics of a theoretical laser at the time [1]. Townes famously remarked that the general reception to the first functional laser was that it was "a solution looking for a problem."

Since that original ruby laser, lasers have found the problems to which they are solutions and have grown to become one of the most important inventions of the 20th century. Townes shared the 1964 nobel prize in physics with Nikolay Basov and Aleksandr Prokhorov for work on laser theory; lasers have enabled several other nobel prizes (e.g. the 1997 prize shared by Steven Chu, Claude Cohen-Tannoudji, and William Phillips for laser trapping and cooling); and laser products cover an extremely wide range of applications from laser surgery to industrial machining, communications and entertainment, military hardware, and an exceptionally wide suite of scientific devices.

Lasers have also grown from their origins in ruby crystals and Helium-Neon tubes to encompass an impressively wide variety of materials. Gas lasers now include Carbon Dioxide, Argon-ion, and gases pumped by chemical reactions. Solid state gain media have grown so diverse that classification as a solid state laser must usually also be accompanied by a sub-classification between crystal (e.g. YAG and Titanium-Sapphire), fiber (e.g. Yb-, Er-, and Th-doped glass fibers), and semiconductor (e.g. junction diode and quantum confinement lasers). A very commonly told apocryphal joke among scientists and engineers working in laser-related fields is that anything will lase... if you pump it hard enough.

It is the very last type mentioned above in the semiconductor sub-classification of solid state lasers that this dissertation will focus on. The Vertical External Cavity Surface Emitting Laser (VECSEL) is a quantum well laser which was first conceived in the 1990s as an attempt to design a semiconductor laser with high beam quality [2]. Up to that point, most semiconductor lasers were designed to be edge-emitting, following tradition from the first light-emitting p-n junction diodes. Even today, though, edge-emitters usually feature very divergent, asymmetric, poor-quality beams and are not the best for good power extraction due to thermal limitations. By shifting to a surface-emitting design and using external optical cavities and sophisticated thermal management techniques, beam qualities could be greatly improved and heat

dissipation could be more effectively assisted, allowing for higher power operation.

In 2012, the culmination of work in power scaling for VECSELs yielded 100 W of CW operation at 1 μm from a single VECSEL chip [3, 4] and as much as 27 W in fundamental TEM_{00} operation [5]. And as with most laser media, when the question "does it work?" is answered affirmative, the question to immediately follow will likely be "can you modelock it?". The work presented here, with brief exceptions, is devoted to thoroughly answering that second question by detailing the principles, challenges, and practical engineering concerns in modelocked VECSELs.

1.1 Dissertation Outline

This dissertation will possess the following general structure:

Chapter 1, which you are presently reading, is a very brief introduction and contains mostly a small portion of context for where VECSELs fit into the larger field of laser study.

Chapter 2 will present background information and literature review for the development of VECSELs up to the work included in this dissertation. Topics discussed will cover the design and processing of a VECSEL device, general knowledge regarding operation of VECSELs, and the current status of the field with regard to power records, wavelength coverage of VECSEL lasers, and pulsed operation benchmarks.

Chapter 3 will contain a discussion of the theoretical background for modelocking lasers. The general E&M theory will be presented first, followed by several strategies for achieving a modelocked laser (including saturable absorption, Kerr lensing, and re-injection) and finally, a couple of concerns will be addressed which are unique to the VECSEL medium's carrier lifetimes and the inherent spectral filtering and field enhancement of a VECSEL structure's microcavity resonance.

Chapter 4 will begin the discussion of what I consider to be the three main challenges to constructing a modelocked VECSEL. It shall be devoted to the first major challenge of designing a proper semiconductor saturable absorber (SESAM). A more detailed discussion of the principle of saturable absorption will be followed by a presentation of the two main SESAM design strategies (low-temperature GaAs growth and surface quantum well growth) and the chapter will finish with the presentation of ultrafast measurements of carrier relaxation in different SESAM structures.

Chapter 5 will cover the second main challenge of dispersion management. A brief overview of general dispersion theory and its effects on modelocking behavior will precede the bulk of the chapter, which will detail dispersion management techniques both inside and outside of the laser cavity.

Chapter 6 will cover the final main challenge, which is unique to the semiconductor media's band structures. Non-equilibrium dynamics within the excited states of carriers in a VECSEL's band structure will be introduced and experimental results of ultrafast pump/probe and in-situ probing measurements demonstrating carrier

behavior on timescales indicative of non-equilibrium dynamics will be presented in detail.

Chapter 7 will cover work in potential applications for modelocked VECSELs as well as introducing the concept of 2-color operation and its applications. Work employing the master-oscillator, power-amplifier (MOPA) design strategy (common in fiber lasers) for will be presented in the context of super continuum generation and potential VECSEL-based frequency combs and 2-color operation will be discussed in terms of its potential for Terahertz generation.

Closing out the main body of the dissertation, Chapter 8 will conclude the dissertation with a brief summary of the experimental results contained in this dissertation as well as the future outlook for the field of modelocked VECSELs.

And finally, if you're a new graduate student reading this dissertation, please refer straight to Appendix A, where I have compiled a collection of practical laboratory techniques and devices which I have found useful during my grad school tenure. They are presented here in a more informal manner that I hope allows for a little easier approachability than formal academic text.

Chapter 2

The VECSEL

The Vertical External Cavity Surface Emitting Laser (VECSEL) is a class of optically pumped semiconductor lasers based on quantum confinement within a stack of semiconductor materials. The following chapter presents a broad overview of VECSELs as a whole, starting from the principle of operation, following through device processing and laser operation, and concluding with a brief look at the current status of the field.

2.1 Structure Design and Principle of Operation

Shown in Fig 2.1, the most common design of VECSEL produces gain by allowing carriers excited in the quantum barriers to diffuse into the neighboring quantum wells and recombine across the quantum well band gap [2]. The difference between the photon energy necessary to excite a carrier in the barriers and the photon energy released by recombination in the wells is an intrinsic inefficiency known as the "quantum de-

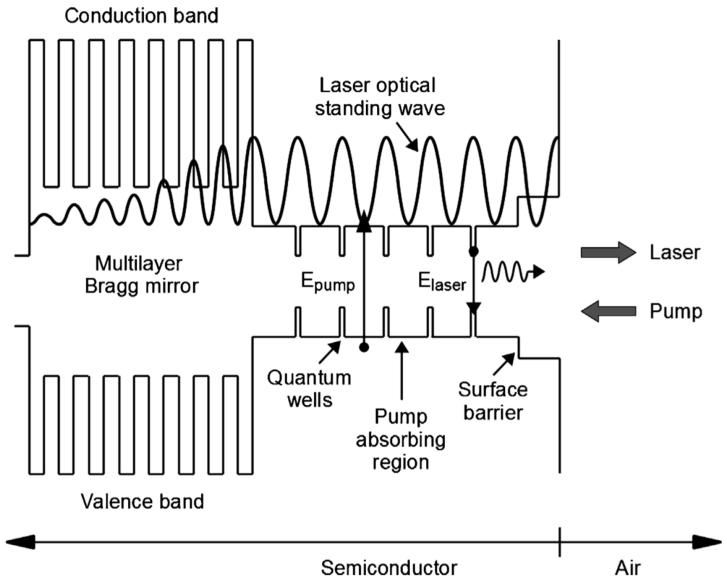

Figure 2.1: Schematic showing the principle of operation for a barrier-pumped VECSEL. Carriers are excited in the quantum barriers and diffuse into the quantum wells where they recombine across the quantum wells' bandgap.

fect" of the device. Minimizing the quantum defect is important as it affects not only the overall efficiency of the device, but because the lost energy must be dissipated as heat, can make thermal management more difficult at high powers.

To truly minimize the quantum defect, a slightly different pumping scheme is required, pumping to a higher energy sub-band within the well itself rather than in the barriers [6]. By doing so, one can utilize pump wavelengths that are not so far separated from the lasing band and thus minimize the difference in photon energies. Despite this minimizing the intrinsic inefficiency of the gain media, however, well pumping is not a common practice as the quantum wells comprise an effective thickness up to an order of magnitude lower than the barrier regions. Thus, the absorbed pump light per round trip in a well-pumping scheme will be a fraction of a barrier-pumping scheme and multiple passes of the pump through the structure will be necessary (often referred to as "recycling" the pump).

On one side of the active region, a Distributed Bragg Reflector (DBR) serves as a highly reflective mirror (>99.9% reflectivity) at the emission wavelength - so that the VECSEL chip as a whole may be treated as a "mirror" inside the cavity with >100% reflectivity - and at least a partial reflector for the pump wavelength - so that pump light not absorbed in the active region may be partially recycled. Because the carriers excited in the quantum barriers are mobile, a surface "cap" layer with higher bandgap energy is necessary at the surface opposite the DBR to prevent surface scattering from depleting the population of excited carriers.

The design of the active region features several details that one must carefully consider. A typical VECSEL design possesses a highly reflective DBR on one side

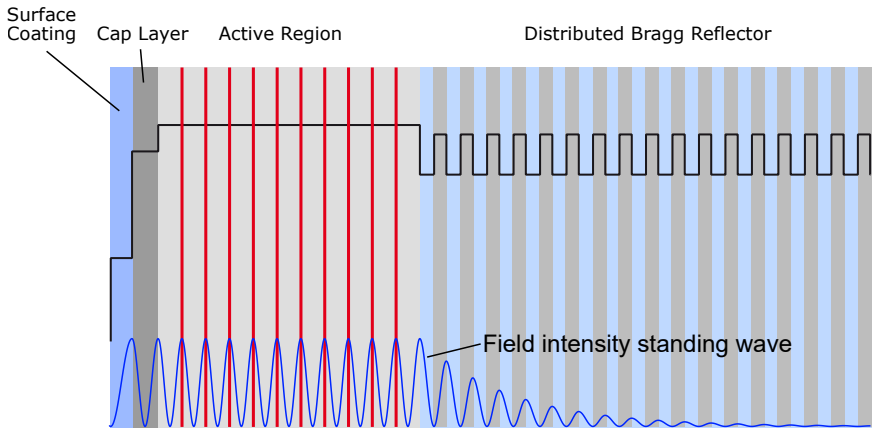

Figure 2.2: Design of a standard Resonant Periodic Gain (RPG) device. In such structures, quantum wells are placed periodically at intracavity field anti-nodes to maximize field enhancement, and therefore gain, but the structure will suffer from significant spectral filtering and as such will not be optimized for wide tunability or broad-spectrum modelocking.

and a surface interface on the other, meaning the active region comprises essentially a low-Q Gires-Tournois etalon [7]. This is a very important aspect of designing the active region as this effect, often called the device's "microcavity," will produce a field enhancement effect at resonant wavelengths, resulting in both higher gain and spectral filtering. In high-gain devices, this can be leveraged to one's advantage [5], but it must be mitigated as much as possible in devices meant for broad-spectrum, ultrafast modelocked operation (e.g. [8]; see also chapter 5).

Another equally important concern regarding the active region is the placement of the quantum wells relative to the cavity's standing wave. Fig 2.2 shows a standard design for high power CW operation. To maximize the interaction between the quantum wells and the intracavity field, wells are placed periodically at the anti-nodes where the field intensity of the design wavelength will be the strongest. This is generally referred to as a Resonant Periodic Gain structure, or RPG.

Similar to the Gires-Tournois Interference (GTI) field enhancement, resonant placement of the quantum wells in the active region is favorable for high power operation, but is not optimal for ultrafast modelocked operation. Because the quantum well placement strongly favors a specific target wavelength, the resulting gain profile will strongly discriminate against the spectral widths required for ultrafast operation

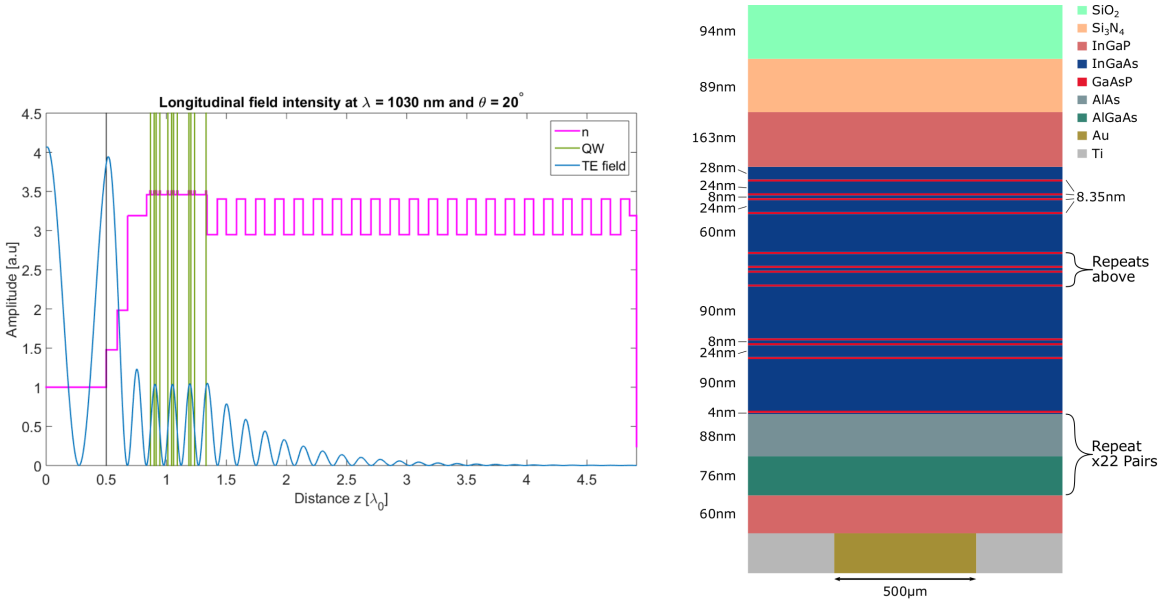

Figure 2.3: Design for a ultrafast modelocked VECSEL. (a) Shows the positions of the various layers relative to the intracavity field for a design wavelength and incident angle of 1030 nm and 20° , respectively. (b) Shows the material composition and layer thicknesses in nanometers. Also displays metal patterning for the hybrid metallization technique (see section 2.2). [9]

in the 100 fs regime and below. To address this, active region designs for ultrafast VECSELs have been adapted to either include fewer quantum wells, shift the placement of the wells off of the anti-nodes of the design wavelength's standing wave, or both [9–13].

An example of such a structure is shown in Fig. 2.3. From Fig. 2.3a, it is clear to see that while the placement of the wells is still somewhat periodic within the active region, they are certainly not placed one-to-one at anti-node peaks, as in Fig. 2.2. Also shown in both Fig. 2.3a and Fig. 2.3b is a double-layer anti-reflection coating, which will reduce the effect of the device's microcavity, helping attain a flatter, broader gain profile for modelocking.

2.2 Device Processing and Thermal Management

At the bottom of the material stack in Fig 2.3b, there is pictured a metallic layer of patterned Titanium and Gold. This is a vital part of the post-growth processing of a VECSEL device. Because the structure is relatively thin and not used in transmittance, heat can be extracted with decent efficiency out of the backside of the DBR. This is done by soldering a diamond heat spreader to the chip with Titanium, which has good adhesion to semiconductor material, Indium, which has is nicely malleable for the bonding process, and Gold, which has good thermal properties and can form alloys with both previous metals [14].

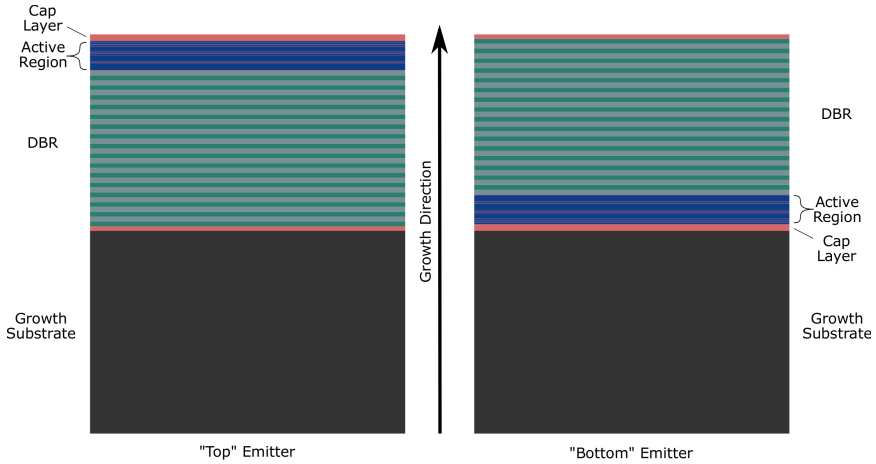

Figure 2.4: Growth styles of a VECSEL wafer.

The general procedure for accomplishing this depends on the growth of the device. Figure 2.4 shows the two types of growth, referred to as "top-" and "bottom-emitter," respectively. For bottom emitters, the DBR is grown last and therefore metal deposition can be done immediately on the top layer of the device once it is cleaved to size from the full wafer (typical sizes for our lab: 2-inch or 4-inch wafers are cleaved

down to 3-4 mm x 3-4 mm chips). By contrast, the growth substrate must first be removed from a top-emitter before any metal can be deposited.

Because good surface contact between the semiconductor and the diamond heat spreader, without any air pockets or solder voids, is the highest priority when processing a VECSEL for any high-power operation, bottom-emitter growth is generally preferred as the metal can be deposited directly onto an epitaxial layer after the device is grown. Having to etch away the growth substrate introduces uncertainty in the quality of the surface, which can be much more easily tolerated in the cap layer (which may alter slightly the emission wavelength and microcavity properties) than in the bonding surface (which may cause burns or prevent the device from operating at all).

Assuming the more common bottom-emitter growth, once the metal layers are deposited on both the back of the DBR and the diamond (in order: a titanium layer directly on the surfaces, a gold layer on the titanium layer, and indium on the gold layer), they are placed in a bonding chamber under vacuum where the temperature is raised to the indium melting point and then slowly lowered to room temperature over a couple hours (to prevent thermal stress from cracking the device) [5].

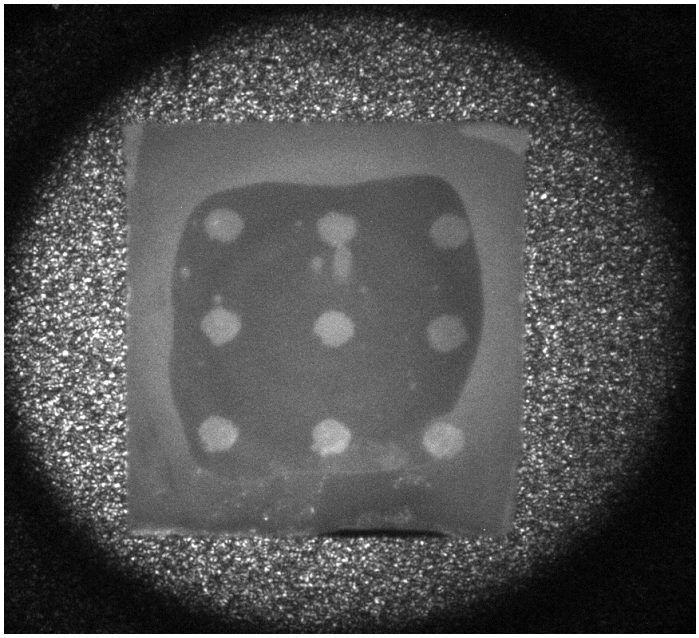

Figure 2.5: Photoluminescence from a VECSEL chip with Au/Ti patterned bonding. Spots with high quality gold surfaces (light grey circles) are potential sites of laser operation; Titanium matte (dark grey) maintains adhesion with the semiconductor material.

The metal deposited does not have to be strictly in the form of the Titanium/Gold/Indium layering described above. For example, recent work in ultrafast VECSELs has demon-

strated that patterned bonding with regions of gold can enable growth with fewer DBR pairs [15]. DBRs as thick as 20+ layer pairs are often required for high reflectivity, but thicker DBRs introduce more dispersion, have poorer thermal properties, and are more difficult to grow. A layer of Gold with high surface quality can be used to offset the reduction in reflectivity if the DBR is grown with fewer pairs. Gold, however, has poor adhesion to semiconductor material, and thus patterning with Titanium is required to prevent separation between the VECSEL and the diamond. Fig 2.5 displays such a patterned-bonding scheme.

After bonding, the growth substrate is removed by a combination of coarse lapping and more precise chemical wet-etching in acid so that the cap layer is revealed. Some designs will feature a thin "etch-stop" layer of poor solubility in the acid used to remove the substrate to make this process more robust at the expense of slightly complicating the growth. Wax can be used at the edges of the chip to prevent lateral etching when the device is submerged in the etchant.

Finally, the fully processed device is clamped to a water-cooled heat sink (commonly copper) with an intermediary layer of indium foil or high-quality thermal paste. A thermoelectric cooler is used to control the heat sink temperature.

Several simple tests can be used after processing to assess the quality of the resulting device. Emission wavelengths significantly shorter than expected strongly indicate over-etching of the cap layer, which will blue-shift the microcavity resonance. Very poor bonding will be immediately obvious from surface liftoff or the chip burning easily when pumped.

A more refined description of bonding quality can be obtained with good precision by measuring the device's thermal impedance [16]. Procedurally, this requires the measurement of the emission wavelength with respect to range of pump powers (heat sink temperature constant) and with respect to a range heat sink temperatures (pump power constant). These measurements should be taken at a linear point in the laser's power curve (i.e. not near threshold or rollover). The slope of the latter, $\frac{\Delta\lambda_{HS}}{\Delta T_{HS}}$, is then divided by the slope of the former, $\frac{\Delta\lambda_{Pump}}{\Delta P_{Pump}}$, to attain the thermal impedance, R_{th}:

$$R_{th} = \left(\frac{\Delta\lambda_{HS}}{\Delta T_{HS}}\right)\left(\frac{\Delta P_{Pump}}{\Delta\lambda_{Pump}}\right). \tag{2.1}$$

This quantity represents the internal temperature of the structure with increasing pump power, which should be inversely proportional to the ability of the device to dissipate heat (and thus the quality of the bonding). What is considered a "good" thermal impedance is heavily dependent on the pump spot size, pump absorption levels, etc., but as a ballpark figure, a well-bonded VECSEL should have a thermal impedance on the order of $\simeq 1\frac{K}{W}$.

A fully processed device is shown in Fig 2.6. Regions of both good and poor bonding quality can be seen in the contrast between the nice, flat surface in the

foreground and the more wavy surface behind. An example of partial lift off, likely due to minor lateral etching during substrate removal, can be seen at the edges of the semiconductor material.

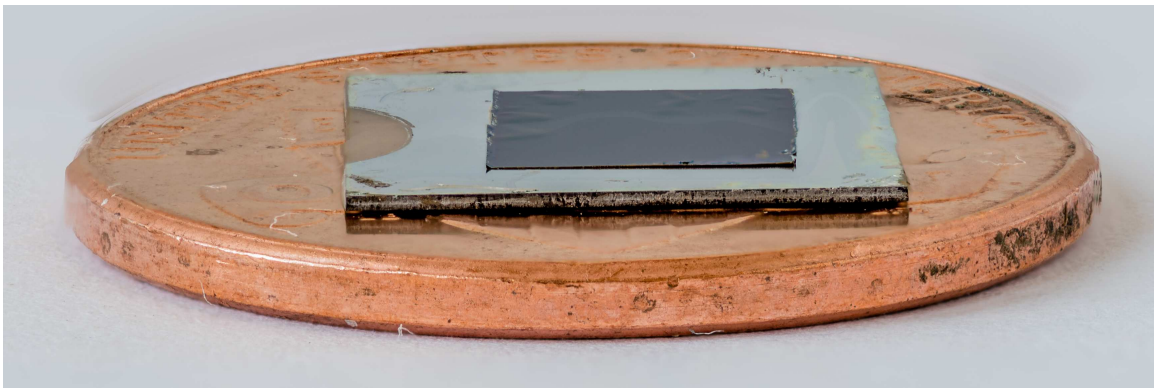

Figure 2.6: Processed VECSEL Device. Semiconductor portion (black) is indium-soldered to a diamond heat spreader (silver/clear). Penny included for scale.

2.3 Operation of VECSELs

Because the VECSEL is an external cavity laser, there is a large amount of flexibility in its cavity design, from the cavity geometries that can be employed to the intra-cavity elements that can be included for a particular application or desired mode of operation.

Geometric layouts include everything from simple linear designs that are 1.5 mm long [17] to complicated multi-pass cavities exceeding 1 m in length [18]. Cavity construction must of course obey basic laser rules: the net loss must be no more than the round-trip gain, the cavity mode and the pump mode must be properly mode-matched, and the cavity must satisfy the stability condition of

$$-1 \;\leq\; \frac{A+D}{2} \;\leq\; 1, \tag{2.2}$$

where A and D are the [0,0] and [1,1] elements of the ray transfer matrix describing a full round-trip through the cavity (see Appendix A.3.1 for details on ray transfer matrices and laser resonators).

The gain of a VECSEL device can vary significantly based on the design of the structure and the quality of the growth and bonding, but is in general less than that of solid state or fiber lasers. A VECSEL maximized for narrow-band CW emission (probably an RPG with maximized field enhancement from its microcavity) may sustain as high as 5% output coupling per pass through the gain, but devices designed

for ultrafast modelocked operation will likely require lower a output coupling of $\simeq$ 1-2% or less.

Mode-matching conditions in VECSELs are significantly less demanding than in other free-space cavity lasers such as Ti-Sapphire. Because VECSEL active regions have thicknesses on the other of hundreds of nanometers to microns, comparing the Gaussian beam width of the pump and the cavity mode at the gain is the only essential requirement, eliminating the need to consider focusing parameters.

The most common pump utilized (particularly in 1 μm VECSELs) is an electrically-pumped, fiber-coupled diode bar, which will have a spot size dependent on the size of the fiber core it is coupled into and the properties of the focusing optics used to image the output of the fiber onto the VECSEL surface. Pump spot sizes (and therefore cavity mode sizes) can range from $\simeq$100 μm in diameter to larger than 1 mm [5].

The method of selecting pump and cavity mode sizes is largely dependent on the device being used and the desired mode of operation. The lasing threshold of any laser is dependent on the power density supplied by the pump within the spatial extent of the cavity mode. A ballpark estimate of the pump power density necessary to make a low-threshold VECSEL lase in an all-HR-mirror cavity is around 1 $\frac{kW}{cm^2}$, but can increase significantly depending on the design and quality of the structure and the output coupling losses. The relative sizes between the pump and cavity modes at the gain is often leveraged to create a soft aperture enforcing TEM_{00} operation. This is accomplished by having a slightly larger cavity mode than pump (by $\simeq$10%). That said, power extraction is generally improved by increasing the relative size of the pump spot and allowing the cavity beam to contain multiple transverse modes, along with increasing the mode size in general.

Intracavity elements play a large roll in operating VECSELs with specific applications in mind. Fabry-perot etalons can be used to enforce polarizations or spectrally filter the lasing band and enforce single-frequency operation [19]. Certain uniaxial crystals (e.g. $Ba(BO_2)_2$ and $LiNbO_3$) can be used intracavity to very efficiently generate second or higher harmonics of the fundamental emission of the VECSEL (e.g. [20]). THz wavelengths can even be generated using intracavity crystals with specialized VECSEL designs [21].

Finally, one of the more practical operation concerns of the quantum foundations of VECSELs comes in the "detuning" of the quantum well absorption line. For a given temperature and pump level, the quantum wells comprising the active region will possess a given band gap energy and fermi distribution of excited carriers, which will give the quantum well a specific absorption edge or gain profile (depending on if the well is inverted). As the temperature of the well increases, these properties will change in a manner that red-shifts that absorption edge or gain peak of the well. This effect is shown for the high-gain VECSEL which yielded 100W of output power in Fig 2.7, from [4].

Because there are other factors affecting the performance of the laser which do

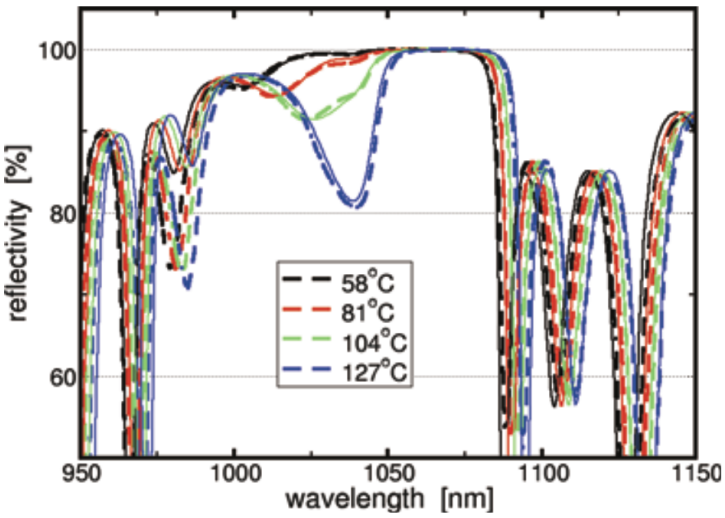

Figure 2.7: Temperature-dependent-reflectivity (TDR) measurement of a high-gain RPG VECSEL (from [4]). The shifting of the quantum well's absorption edge is displayed by the red shift in the absorption feature in the TDR and the interaction between QW red shift and the microcavity resonance is shown in the increasing absorption depth of said feature.

not shift as dramatically with temperature (e.g. the DBR stopband, the microcavity resonance, etc.) and the temperatures inside the structure will reach above of 100° C, the resulting gain bandwidth of the quantum well can easily shift out of agreement with those other design parameters, significantly increasing loss and causing the VECSEL to stop lasing. This is referred to as the device's thermal "rollover" point and is generally the limiting factor on the emission power of a device if it is bonded and cooled well enough to prevent burning.

If the quantum well were designed such that its gain peak coincided with the desired emission wavelength at room temperature, it would reach rollover at very low pump powers, severely limiting the performance of the device. For this reason, quantum wells are often "detuned," or grown such that they are optimized for shorter wavelengths at room temperature under the assumption that they will red-shift into agreement with other laser parameters during operation, extending the operational regime prior to thermal rollover. A minor consequence when lower power operation is desired is that sometimes the VECSEL must be heated via heatsink if its quantum well structure was widely detuned from the design wavelength on the assumption that it would be used for higher power operation.

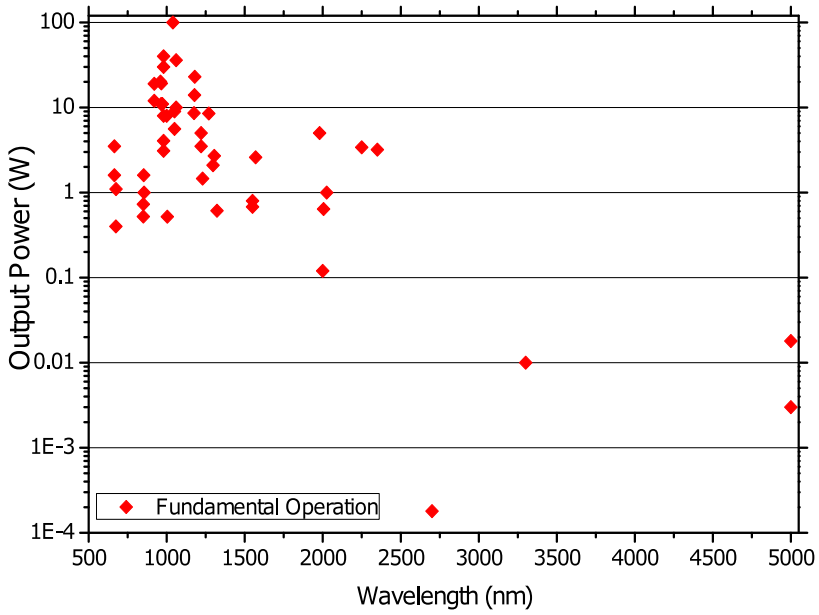

Figure 2.8: Wavelength coverage of published, fundamental-operation VECSEL results. (inset) zoom of the regime near 1 μm.

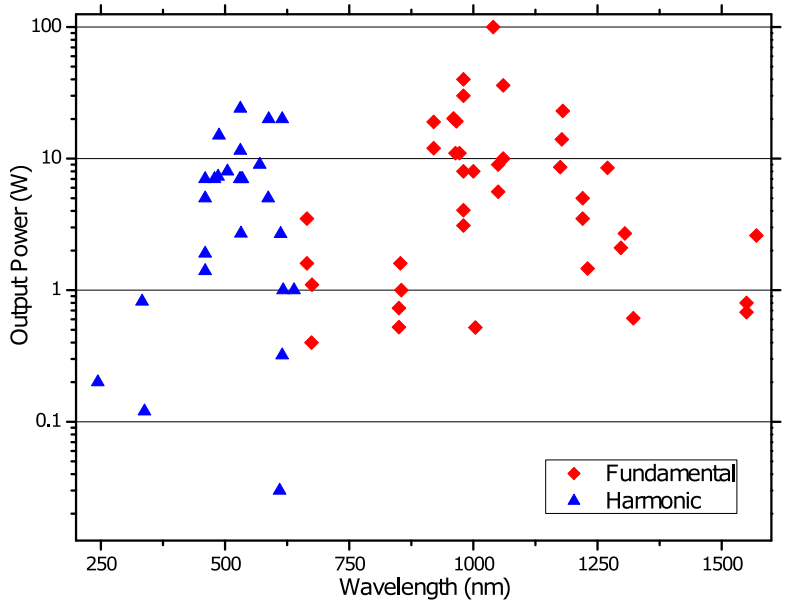

Figure 2.9: Wavelength coverage of published VECSEL results, including fundamental operation near $1\mu m$ and harmonic generation results as short as 244 nm.

2.4 Current Status of the Field

There are many different perspectives from which to gauge the current status of the VECSEL field. I am choosing here to pick three of them: Wavelength coverage and

CW power records for VECSELs in both fundamental operation and harmonic generation; Modelocking results, including the current record pulse durations and peak powers; and the applications for which VECSELs are currently being investigated.

Fig 2.8 displays, to the best of my knowledge, the current wavelength coverage and power records at the time of this writing [2,22–32]. The overall CW power record remains 100 W [3, 4]. In single frequency, 15W has been attained with a sub-MHz linewidth [19]. While the bulk of the work has centered around 1 μm, fundamental operation in VECSELs presently spans from 665 nm [28] into the mid-IR at 3 μm [31] 5 μm [29].

Fig 2.9 shows the work done in intracavity frequency conversion [2,20,33,34]. Although most harmonic generation work is done in the visible, the shortest wavelength generated in this manner is 244 nm, from the 4th harmonic of a 976 nm VECSEL [35]. The highest generated harmonic power is 24 W from a frequency-doubled 1062 VECSEL [20].

In terms of modelocking performance, Fig 2.10a shows published results in terms of pulse durations and average output powers and Fig 2.10b shows the same results plotted against their recorded peak powers [11, 12, 36, 37]. The current record pulse duration is 96 fs [12] with minor external compression. The current record peak power is 6.3 kW [37], a result discussed later in this thesis. In repetition rates, VECSELs have been constructed to operate as low as 85 MHz [18] and as high as 100GHz [17].

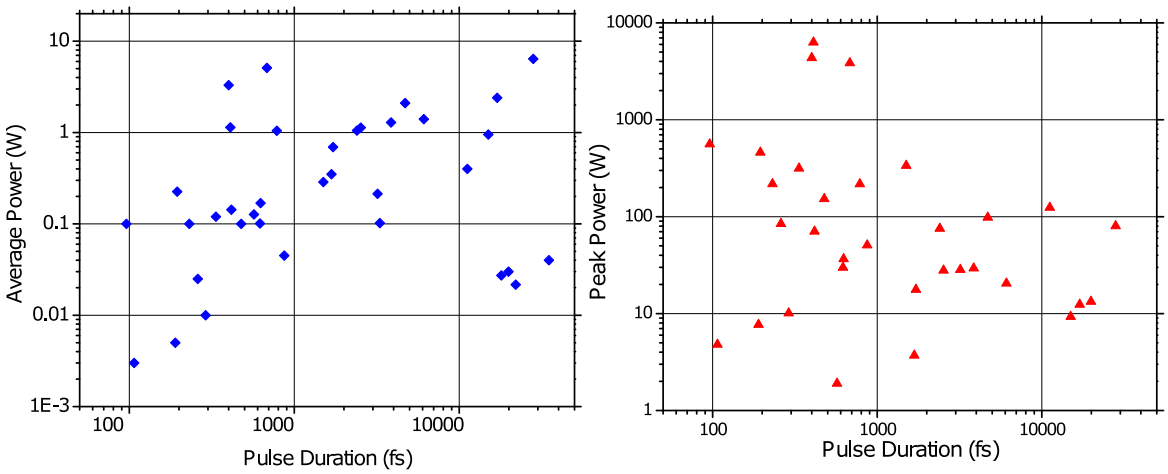

(a) (b)

Figure 2.10: Current status of modelocking results in VECSELs with respect to (a) their pulse durations (record: 96 fs [12]) and (b) their peak powers (record: 6.3 kW [37]).

Finally, applications for which VECSELs are being investigated are numerous. Harmonic generation at 588 nm has been shown, allowing VECSELs to be a potential

candidate for compact Guidestar Adaptive Optics systems on large telescopes [34,38]. Work is being done on amplifying and spectrally broadening VECSEL seed oscillators for use in frequency comb applications [39]. Two-color VECSELs are a candidate for a potential compact THz source [40–42]. There is even preliminary work being done investigating the possibility of solar pumping a VECSEL device [43].

2.5 Chapter Summary

This chapter has covered a brief overview of VECSELs from the principle of operation through to the current status of the field, including the cutting edge efforts in pulse shortening and peak power scaling which will be discussed in much greater detail in successive chapters. The following chapter will begin an in-depth look at the modelocking of these devices by presenting the theoretical background for modelocking in general and discussing some specific challenges inherent to modelocking VECSELs specifically.

Chapter 3

Modelocking Background

Without any special perturbation, a laser will tend to operate in continuous wave - a state of operation where the lasing longitudinal modes all possess random phases and thus for all classical purposes, power is emitted from the laser in a steady, continuous manner. While CW operation is often sufficient, it is nevertheless possible to obtain field intensities orders of magnitude higher by forcing the laser to emit the same amount of average power in short, bright pulses rather than a continuous stream.

This can be done by two broad categories of mechanisms, the first of which relies on controlling the loss of the cavity, allowing brief periods where stimulated emission is possible and otherwise allowing as much energy to accumulate in the gain medium as its carriers' upper-state lifetime allows. This method is referred to as "Q-switching," as the technique revolves around controlling the cavity's loss, or Q-factor. Q-switching can produce very high pulse energies, but is generally limited to pulse durations on the order of nanoseconds or longer [44].

For ultrafast pulse generation in the picosecond regime and shorter, it is necessary to generate a pulse by phase-locking a laser's supported longitudinal modes. By locking numerous modes in phase, one enforces localized constructive interference at a specific point in their relative distributions. That point of constructive interference can then oscillate within the cavity and be emitted through an output coupler without breaking the phase lock.

This section will cover the theoretical foundation for enforcing such a phase lock by various means. It will begin with possibly the most intuitive picture, active modelocking, and by a series of modifications to the active modelocking framework, arrive at soliton solutions for saturable absorption in a system with realistic dispersion.

3.1 Active Modelocking

When a sine wave is subjected to an external modulation, its frequency spectrum gains side bands corresponding to the sum and difference frequencies between the sine wave's frequency and the modulation frequency. If this sine wave is a supported mode of a laser's cavity and gain bandwidth, and if the modulation is tuned to be the free spectral range of the laser's cavity, the generated side bands will overlap in frequency with the neighboring supported modes of the cavity. This allows power transfer between the central mode and its neighbors, thus enabling the central mode to injection lock its neighboring mode in the gain [45]. Further, the same process can be extended to those neighboring modes, allowing them to injection lock their

neighbors and so on until all frequencies supported by both the cavity and the gain bandwidth are locked. The master equation defining this process is:

$$\Delta A_n = \left\{ \frac{g}{1 + \left(\frac{n\Delta\Omega}{\Omega_g}\right)^2} - l \right\} A_n + \frac{1}{2} M(A_{n-1} - 2A_n + A_{n+1}), \tag{3.1}$$

where ΔA_n defines the change in amplitude of the mode 'n' steps away from the center frequency. This equation essentially defines a discrete set of modes, A_n, interacting with a linear gain, 'g', a linear loss, 'l', and a modulation-induced power exchange with neighboring modes, $\frac{1}{2}M(A_{n-1} - 2A_n + A_{n+1})$, where M is the modulation strength. The gain term is limited as the mode gets further from the center frequency in steps of the modulation frequency, $n\Delta\Omega$, due to the finite width of the gain, Ω_g (this is actually the half-width, to describe the distance from the center frequency).

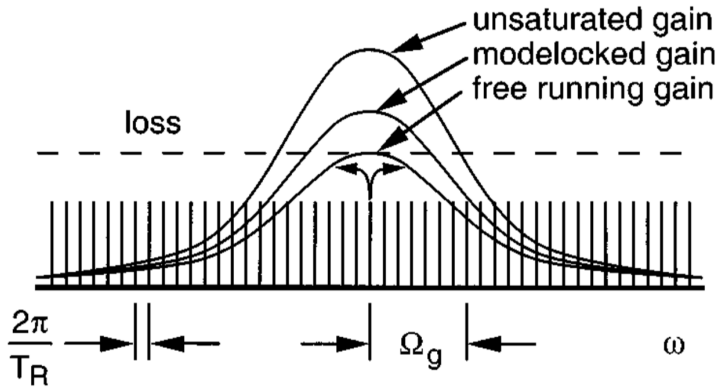

Figure 3.1: Supported longitudinal modes of a cavity overlayed with gain curve. For active modelocking, the modulation frequency must be exactly the spacing between two "teeth" of modes. From [45] .

If we assume that the modes are tightly packed, we can then make a couple approximations:

1. The amplitude of discrete modes, $A_n \Longrightarrow A(\Omega)$, a continuous amplitude distribution.

2. Discrete frequency steps from center frequency, $n\Delta\Omega \Longrightarrow \Omega$, a frequency continuum.

3. Discrete injection lock relationship, $(A_{n-1} - 2A_n + A_{n+1}) \Longrightarrow \Omega_m^2 \frac{\partial^2 A}{\partial \Omega^2}$, a second derivative in frequency, taking the modulation frequency, Ω_m, into account.

With one further step, which is to expand the gain filtering as a power series and take the first two terms, $\frac{g}{1+\left(\frac{\Omega}{\Omega_g}\right)^2} \Rightarrow g\left(1-\left(\frac{\Omega}{\Omega_g}\right)^2+...\right)$, the master equation in Eq. 3.1 becomes

$$\Delta A(\Omega) = (g-l)A(\Omega) - g\left(\frac{\Omega}{\Omega_g}\right)^2 A(\Omega) + \frac{1}{2}M\Omega_m^2\frac{\partial^2 A}{\partial \Omega^2}. \tag{3.2}$$

This equation still contains roughly the same structure as the discrete master equation, Eq. 3.1: a linear term dependent on the round-trip gain and loss, a $\frac{1}{\Omega_g}$ term describing the finite spectral window supported by the gain bandwidth, and a term describing some scheme for triggering modelocking. If $\Delta A(\Omega)$ is set to 0, one can find the steady-state solution for an established pulse in this cavity:

$$\begin{aligned} A(\Omega) &= A_0 e^{-\Omega^2\tau^2} \\ \tau &= \sqrt[4]{\frac{2g}{M\Omega_g^2\Omega_m^2}} \end{aligned} \tag{3.3}$$

There is one more major step to complete the picture of active modelocking. All of the above description is done in a frequency domain; this is a good way to begin the derivation, as it is a nicely intuitive picture of how the constituent modes are locked together in a pulse, but the time domain is more often desirable when considering the formation of the pulse and calculating its temporal width. The attainment of a master equation and solution similar to Eqns. 3.2 and 3.3 in the time domain is fortunately relatively simple, owing to the fourier relation between time and frequency [46]:

$$\begin{aligned} a(t) &= \int e^{i\Omega t}A(\Omega)d\Omega \\ A(\Omega) &= \frac{1}{2\pi}\int e^{-i\Omega t}a(t)dt \end{aligned} \tag{3.4}$$

The main changes to Eq. 3.2 because of this fourier transform are the $\left(\frac{\Omega}{\Omega_g}\right)^2 A(\Omega)$ term and the second derivative in the modulation portion. It is a widely known property of Fourier transform pairs that a derivative in one domain transforms to multiplication by the domain variable of the other domain (see [46] for a proof of this). Therefore, the second derivative in the frequency domain becomes a factor of $-t^2$ when transformed and conversely, the factor of Ω^2 becomes a negative second derivative in time. All constants (e.g. g, l, Ω_g, etc.) are unaffected by the transform and as the transform is a linear operator, the overall structure of the master equation does not change. The master equation in the time domain is therefore:

$$\frac{1}{T_R}\frac{\partial}{\partial T}a(T,t) = (g-l)a(T,t) + g\left(\frac{1}{\Omega_g}\right)^2\frac{\partial^2}{\partial t^2}a(T,t) - \frac{1}{2}M\Omega_m^2 t^2 a(T,t). \tag{3.5}$$

The new term on the left is introduced as a long-term evolution term. Its time domain, T, is a distinct time domain from the time domain utilized on the right-hand side, t. The long-time variable is intended to account for time scales on the order of the round-trip time, T_R, and longer, allowing for tracking of the pulse formation and long-time dynamics. The short-time variable is intended to describe the temporal characterization on the order of the pulse duration itself. Realistically, these time scales are disparate by several orders of magnitude; barring exceptional cavity designs, pulse durations are concerned with picosecond and femtosecond timescales, contrasting with long-time variation that occurs on the order of hundreds of picoseconds to nanoseconds or longer.

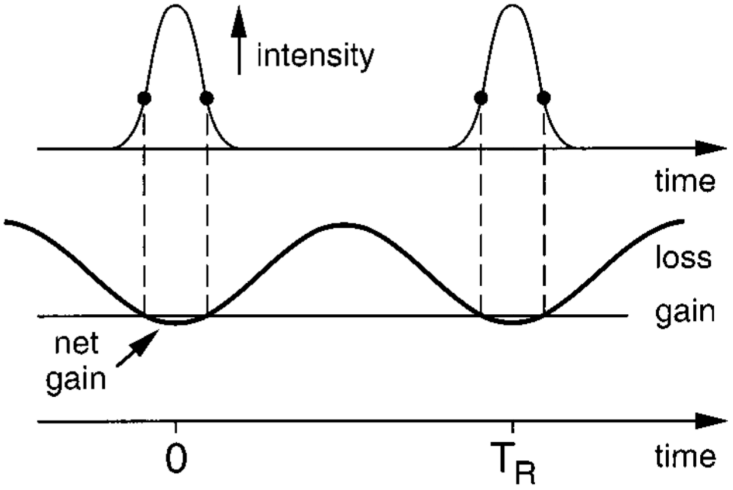

Figure 3.2: Time domain picture of the net gain and loss in the case of active mode-locking. From [45].

Despite including the long-time evolution for completeness' sake, the steady state solution to the time-domain master equation is found by setting the left side of Eq. 3.5 to 0:

$$\begin{aligned} a(t) &= a_0 e^{\frac{-t^2}{2\tau^2}} \\ \tau &= \sqrt[4]{\frac{2g}{M\Omega_g^2\Omega_m^2}} \end{aligned} \tag{3.6}$$

3.1.1 Practical Implementation

Active modelocking can be achieved both with intracavity modulators and an external modulation that is re-injected to the cavity. Unfortunately, the time constant of the produced pulses has a second order dependence on the modulation frequency, Ω_m, which imposes a severe lower limit to the achievable pulse duration as modulators do not exist which are fast enough to support fs-scale pulses. As such, active modulation is generally reserved for Q-switch lasers or systems where long pulses are tolerable.

3.2 Fast Saturable Absorption

Although active modelocking is a nicely intuitive picture of how the modes of a laser cavity can be locked together to form a pulse, virtually all ultrafast (ps duration and shorter) operation is done by means of saturable absorption. The key method to saturable absorption is introducing some loss mechanism to the laser which can be defeated more easily by an intense pulse than by CW operation. Modelocking operation is then preferred as the laser will see less round trip losses in a pulsed state than a CW one [45].

This can be done using an absorber with a saturable loss, as the name implies, but the dynamics are actually identical to other methods such as using an intensity-dependent kerr lens to focus through an aperture or relying on a nonlinear polarization rotation in a length of fiber to produce the correct polarization state for an isolator. It is therefore extremely important to understand the dynamics of modelocking by saturable absorption if one is intent on working with ultrafast lasers.

There are two broad categories: "fast" saturable absorbers - discussed here - where the saturation is fast enough for the gain to be considered constant on comparable and "slow" saturable absorbers - discussed in the next section - where saturation of the gain must also be taken into account.

A simple picture of a saturable absorber will dictate that its losses are defined as

$$s(t) = \frac{s_0}{1 + \frac{I(t)}{I_{sat}}}. \tag{3.7}$$

Expanding in a power series and taking the first two terms,

$$s(t) = s_0 \left(1 - \frac{I(t)}{I_{sat}}\right) = s_0 - \frac{s_0 |a(t)|^2}{I_{sat} A_{eff}}, \tag{3.8}$$

one can see the physics of the process: the absorber's overall loss, $s(t)$, contains some unsaturable portion, s_0, and some other term which depends on the power density of the lasing modes at the absorber, $I(t) = \frac{|a(t)|^2}{A_{eff}}$. The saturable term will serve to reduce the overall loss if the laser's power density is high, limited by the possible saturation in the absorber, I_{sat}. If we condense the constant terms into a single constant for simplicity's sake, the loss of the absorber can be written as

$$s(t) = s_0 - \gamma|a(t)|^2$$
$$\gamma = \frac{s_0}{I_{sat}A_{eff}} \tag{3.9}$$

where γ is often called the "Self Amplitude Modulation" (or SAM) coefficient, and the master modelocking equation for fast saturable absorption can be written:

$$\frac{1i}{T_R}\frac{\partial}{\partial T}a(T,t) = (g-l)a(T,t) + g\left(\frac{1}{\Omega_g}\right)^2\frac{\partial^2}{\partial t^2}a(T,t) + \gamma|a(t)|^2a(T,t). \tag{3.10}$$

Here, the unsaturable portion of the loss, s_0, is simply included in the linear loss of the cavity, l. Again, one can see the same general structure as in the active modelocking case. There is a long-time pulse formation term on the left, a $(g-l)$ term for the linear gain and loss of the cavity, a $\left(\frac{1}{\Omega_g}\right)^2$ term do describe the bandwidth limits of the gain, and a final term to describe the mechanism inducing modelocking, now depending on $\gamma|a(t)|^2$ to represent saturable absorption.

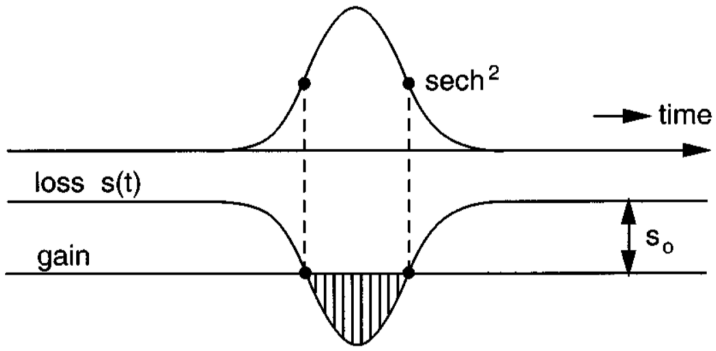

Figure 3.3: Net gain and loss for a modelocking in the case of fast saturable absorption. The gain remaining constant is the key feature differentiating fast and slow SAM modelocking. From [45].

If we again set the long-time dynamics to 0 to indicate a steady-state system, the solution to the remaining diff. eq. is in the form of hyperbolic secants:

$$a(t) = a_0 sech\left(\frac{t}{\tau}\right)$$
$$\tau = \sqrt{\frac{2g}{\gamma A_0^2\Omega_g^2}} \tag{3.11}$$

3.2.1 Practical Implementation

Although it is named "fast saturable absorption," the physical systems which work fast enough for the constant-gain approximation to remain valid are ironically not absorbers [45]. Rather, they are artificial absorbers which are still accurately described by saturable absorption theory.

The two chief examples are polarization rotation and kerr lens modelocking. The former is based on a nonlinear polarization rotation inside of an isotropic Kerr material (often inside an optical fiber). Included in the cavity should be a waveplate for manual adjustment of the polarization and an isolator, or some other device with high, polarization-dependent loss.

While linear and circular polarization will acquire a simple phase shift, if the polarization is elliptical inside the nonlinear material, it will rotate through a certain angle depending on the field intensity. Therefore, input parameters can be found such that the nonlinear polarization rotation inside the kerr medium produces low loss at the isolator only for high intensities. Thus, the round-trip losses of the cavity greatly favor a pulsed state rather than a CW one. As one could guess from the reference to fiber, this strategy is very common in modelocked fiber lasers.

Modelocking based on the kerr lens effect is very similar but instead of a polarization-dependent loss mechanism, a hard aperture is added to the cavity. This scheme relies on the difference in spatial intensity across laser's transverse mode to produce an spatially-varying index profile in a kerr medium. Because the center of a TEM_{00} profile is much higher than in the wings, the kerr effect is excited more strongly in the center than the wings and, depending on the material, the resulting phase shift to the field can mimic a lens, prompting a self-induced focusing of the beam. The aperture is maintained such that in a pulsed state, the beam will self-focus enough to cleanly pass the aperture, but in a CW state, it will be clipped and therefore incur heavy losses. This is the most commonly used method to modelock Titanium:Sapphire lasers.

3.3 Slow Saturable Absorption

When it comes to saturable absorbers which are based on actual absorption of light by a material, time scales are unfortunately too long to be adequately described by the mechanics in the previous section. When the time scales involved are not short enough, the gain can no longer be approximated as constant and saturation of the gain must be taken into account. To do this, a more detailed description of saturation is employed for the gain than was used in Eq. 3.7. The gain must obey the following dynamics

$$\frac{dg}{dt} = -g\frac{|a(t)|^2}{W_g}, \tag{3.12}$$

where g and $a(t)$ are still the linear gain and pulse amplitude, respectively, and W_g is the saturation energy of the gain. This is an approximation of the *relaxation equation* of the gain, where it is assumed that the depletion of the gain is much stronger and faster than its relaxation during the period where the pulse is not present [47]. While a linear recovery term must be added to describe the gain per round trip, when considering only the time period where the pulse is interacting with the gain, that term can be discarded. The solution in this case is given by:

$$g(t) = g_i e^{-\int_0^t \frac{|a(t)|^2}{W_g} dt}. \tag{3.13}$$

A completely analogous process gives the saturation for the absorber, as well, where W_s is the analogous saturation energy for the absorber:

$$s(t) = s_i e^{-\int_0^t \frac{|a(t)|^2}{W_s} dt}. \tag{3.14}$$

The next step is ideally to use both Eq. 3.13 and Eq. 3.14 in a master modelocking equation, as in Eq. 3.10. Unfortunately, no analytical solutions can be obtained without an approximation for both gain and loss [45]. In this case, a Taylor series expansion around the exponent argument out to the second order:

$$\begin{aligned} g(t) = g_i e^{-\int_0^t \frac{|a(t)|^2}{W_g} dt} &\simeq g_i \left[1 - \left(\int_0^t \frac{|a(t)|^2}{W_g} dt \right) + \frac{1}{2} \left(\int_0^t \frac{|a(t)|^2}{W_g} dt \right)^2 \right] \\ s(t) = s_i e^{-\int_0^t \frac{|a(t)|^2}{W_s} dt} &\simeq s_i \left[1 - \left(\int_0^t \frac{|a(t)|^2}{W_s} dt \right) + \frac{1}{2} \left(\int_0^t \frac{|a(t)|^2}{W_s} dt \right)^2 \right] \end{aligned} \tag{3.15}$$

With one more approximation of a fixed spectral filter rather than a gain-induced one, $\Omega_g \Rightarrow \Omega_f$ (necessary because Ω_g would actually depend on Eq. 3.13), we can write the master modelocking equation for slow saturable absorption,

$$\frac{1}{T_R} \frac{\partial}{\partial T} a(T,t) = (g(t) - l) a(T,t) + \frac{1}{\Omega_f^2} \frac{\partial^2}{\partial t^2} a(T,t) - s(t) a(T,t), \tag{3.16}$$

where I have left the gain and loss terms in Eq. 3.15 in their $g(t)$ and $s(t)$ forms to preserve the ability to read the equation. At the risk of belaboring the point, once again the general structure of the master equation is the same as in sections 3.1 and 3.2; there is a long-time evolution term, a linear gain and loss term, a spectral bandwidth term, and a pulse shaping term, which in this case includes the saturation of both gain and loss. The steady state solution to this equation is again a hyperbolic secant in form,

$$a(t) = a_0 sech\left(\frac{t}{\tau} \right), \tag{3.17}$$

but this time contains a rather complicated pair of constraints to its coefficients [45]:

$$\tau = \sqrt[4]{\frac{4}{\Omega_f^2 a_0^4}} \left(\frac{s_i}{W_s^2} - \frac{g_i}{W_g^2} \right)^{-\frac{1}{4}} \tag{3.18}$$

$$g_i - s_i - l = -\frac{1}{\Omega_f^2 \tau^2} + g_i \left[\frac{a_0^2 \tau}{W_g} - \left(\frac{a_0^2 \tau}{W_g} \right)^2 \right] - s_i \left[\frac{a_0^2 \tau}{W_s} - \left(\frac{a_0^2 \tau}{W_s} \right)^2 \right] \tag{3.19}$$

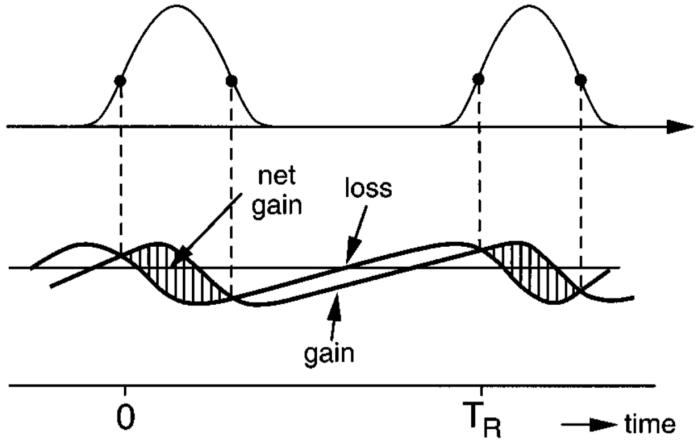

Figure 3.4: Net gain and loss in the case of slow saturable absorption. Note that the loss *must* saturate faster than the gain for net gain to be possible within the pulse. From [45].

From the first constraint we can note something important concerning the relative saturation of the gain and loss, namely that $\frac{s_i}{W_s^2} > \frac{g_i}{W_g^2}$ is required for τ to be real. Hence, the absorber must saturate more strongly than the gain for a pulse to form (this can also be seen in Fig 3.4). The second constraint gives the net gain of the system prior to interaction with the pulse. Because Eq. 3.18 dictates that the saturation terms are larger than the gain terms, this net gain is therefore negative. We can therefore draw the conclusion that laser operation under slow saturable absorption is stable to noise buildup between the pulses.

3.4 Saturable Absorption in the Presence of Dispersion

Up to this point, this discussion of modelocking has lacked one key element to the realistic behavior of modelocked systems in the absence of any dispersion terms.

Because the crux of a modelocked system is a collection of wavelengths locked together in time, accounting for all dispersive effects' tendency to pull the locked wavelengths apart in time (referred to as "chirping" the pulse) is very important to understanding modelocked behavior. It should be noted that this section uses a convention for dispersion that works in units of fs^2, common for discussions of material dispersion parameters. This differs from convention commonly used in waveguide discussions (such as optical fibers), that works in units of $\frac{ps}{nm*km}$.

3.4.1 Dispersion Principles

There are, in general modelocked systems, two types of dispersion that are of particular importance: Group Velocity Dispersion (GVD) and Self-Phase Modulation (SPM) [48].

Self-Phase Modulation arises from the Kerr effect - a second order non-linearity which causes an intensity-dependent variation in a material's index. Self-focusing is the ultimate result if SPM is considered with respect to a laser mode's transverse profile. However, if considered in the time domain, the net effect is a self-induced and spectrally-dependent phase shift between the constituent wavelengths of a high-peak-power pulse, causing a negative chirp to form. I.e. shorter wavelengths will shift towards the lead edge of the pulse, with longer wavelengths shifting to the back edge. The effect is described for a medium of length, L, and nonlinear index, n_2, by the Kerr coefficient,

$$\delta = \frac{2\pi}{\lambda}\frac{n_2 L}{A_{eff}}, \tag{3.20}$$

which can then be added to the master modelocking equation on the $|a(t)|^2$ term to be properly dependent on intensity [45].

Group Velocity Dispersion arises from the wavelength dependence of a material's linear refractive index. GVD - also often discussed as Group Delay Dispersion (GDD), which is the net GVD for a length of material - generally induces a positive chirp (i.e. long wavelengths shift to the lead edge of the pulse; short wavelengths shift to the back edge). The calculation of GDD is a bit more complicated than SPM as it can stem from multiple sources. It can be the result of simple material dispersion in a bulk medium,

$$D = L * GVD = L\left[\frac{2}{c}\left(\frac{\partial n}{\partial \omega}\right)_{\omega=\omega_0} + \frac{\omega_0}{c}\left(\frac{\partial^2 n}{\partial \omega^2}\right)_{\omega=\omega_0}\right], \tag{3.21}$$

waveguide dispersion in a waveguide with propagation constant β,

$$D = \frac{L}{2}\frac{\partial^2 \beta}{\partial \omega^2}, \tag{3.22}$$

or as the result of the Kramers-Kronig relation in materials with a significant complex portion to their index, α, indicating gain or absorption,

$$n(\omega) = 1 + \frac{c}{\pi} \wp \int_0^{+\infty} \frac{\alpha(\Omega)}{\Omega^2 - \omega^2} d\Omega, \tag{3.23}$$

where $\wp$ denotes the principle integral and the variable index in Eq. 3.23 would be used in Eq. 3.21 to attain the GDD.

It is critical to understand the different sources of GDD as different gain media have different dominant GDD mechanisms. For example, Erbium doped silica has a broad, flat gain bandwidth, so an Erbium fiber laser (with fiber lengths on the meter scale) will have significant contribution from material and waveguide GDD, but the smooth gain bandwidth yields a complex index with minimal variance and therefore less GDD results from the Kramers-Kronig contribution. By contrast, VECSELs contain comparatively minimal material thicknesses (on the micron scale) and corresponding simple GDD, but have more limited bandwidths, which make the Kramers-Kronig contribution to the GDD much more important to consider. GDD features are also enhanced by the Q-factor if they occur inside a resonator, so structures like the microcavity of a VECSEL must be carefully considered.

3.4.2 Soliton Modelocking

In general, both types of dispersion cause deterioration in the modelocking behavior of a laser and managing the dispersion with a variety of techniques discussed in Chapter 5 is of the utmost importance when attempting to generate ultrashort pulses on femtosecond time scales. That said, one can note that for shorter wavelengths (e.g. below 1.3 μm in Silica), SPM produces a negative chirp (blue leads red) while most GDD mechanisms produce a positive chirp (red leads blue). Thus, it is possible to achieve operation in what is called the "soliton" regime, where SPM and GDD effects balance the net dispersion for the laser as a whole. The master equation for this regime is [48]:

$$\frac{1}{T_R}\frac{\partial}{\partial T}a(T,t) = (g-l)a(T,t) + \left(\frac{1}{\Omega_f^2} - iD\right)\frac{\partial^2}{\partial t^2}a(T,t) - \left(s(t) + i\delta|a(T,t)|^2\right)a(T,t). \tag{3.24}$$

This equation is strikingly similar to the case of slow saturable absorption with two distinct changes: the saturable gain in the slow absorber case has been replaced with a constant linear gain term, assuming that gain saturation will take place over several pulses rather than in the duration of a single pulse. Additionally, dispersion terms describing the net GDD, D, and SPM, δ, have been introduced. As was the case for slow saturable absorption, there is no known analytical solution when real gain filtering is included, and therefore a fixed spectral filter has been included instead.

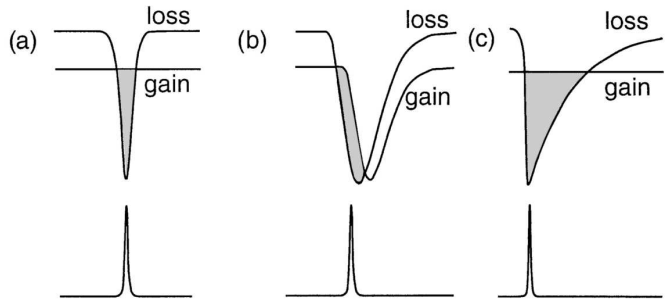

Figure 3.5: Comparison of net gain and loss for (a) fast saturable absorption, (b) slow saturable absorption, and (c)dispersion-managed soliton modelocking. From [49]

Like the slow saturable absorber case, the solution will be in the shape of a hyperbolic secant. However, due to the introduction of dispersion, a phase term is now also necessary:

$$a(T,t) = a_0 sech\left[x(T,t)\right] e^{i\theta(T,t)}, \tag{3.25}$$

with time and phase functions:

$$\begin{aligned} x(T,t) &= \frac{1}{\tau}\left(t + 2Df_0T - t_0\right) \\ \theta(T,t) &= -f_0 t - D\left(\frac{1}{\tau^2} - f_0^2\right)\frac{T}{T_R} + \theta_0. \end{aligned} \tag{3.26}$$

A frequency offset, f_0, a time shift, t_0, and an initial phase θ_0 are required by the introduction of dispersion. One can note that this solution is close to Eq. 3.17 for a slow absorber if these terms, along with the dispersion, are set to 0.

A stable pulse width is achieved by the balance of positive GDD, and intensity-dependent SPM [45, 48]:

$$\frac{1}{2}\delta\tau^2|a_0|^2 = |D|. \tag{3.27}$$

This is admittedly a surface-level description of soliton modelocking. There are many further concerns regarding soliton modelocking which unfortunately exceed the scope of this thesis. There is a full treatment using soliton perturbation theory in [48] which delves much deeper than this section.

3.5 Numerical Methods for Modelocking in VECSELs

As the numerous approximations in the previous sections indicate, in order to realistically describe modelocking in an actual VECSEL structure, analytical solutions must

be discarded in favor of numerical approaches. At the time of this thesis, there have been published two major approaches to numerically simulating modelocked VECSELs: a phenomenological rate equation model [50] and a full, many-body quantum model [51, 52].

3.5.1 Rate Equation Model

At the foundation of the rate equation model is a pair of saturation equations for gain and loss similar to those found in Section 3.3 covering slow saturation. In this case, however, the long time domain previously ignored is included and pulse shaping is going to be allowed to occur over many round trips of the cavity. The general form of the model, shown in Fig 3.6, is to start with a pulse expression which is then input successively to different modules describing the various cavity elements [50]. Covering each module in turn...

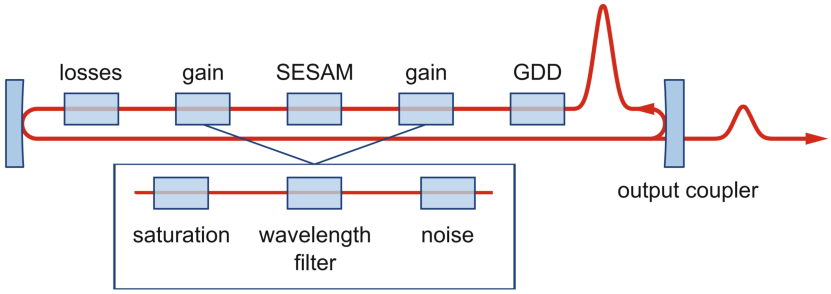

Figure 3.6: Schematic of rate equation model; each block represents one of the operators described below. From [50].

Pulse Expression
The expression for the field of the simulated pulse is a simple use of the slowly varying envelope approximation,

$$E(t) = Re\left\{a(t)e^{-i\omega_0 t}\right\}. \tag{3.28}$$

Fast Fourier Transform algorithms make it fairly trivial to transform this expression between time and frequency domains as is necessary for application of the various operators described below. As there is a Fourier relationship between time window and frequency resolution, care should be taken that the time window and resolution are chosen such that there is adequate sampling in both domains (at least 20 sampling points in each domain is recommended in [50]).

The simulation can either be seeded with a particular pulse shape by applying a choice of $a(t)$, or allowed to form a pulse from noise. To reduce calculation load, the

pulse can be defined over a particular time window, in which the full calculations are run, and simpler expressions (which assume $a(t) = 0$ in the equations below) can be used to update the status of any element between interactions with the pulse. For example, the full gain saturation only needs to be calculated during the time interval across which the pulse is expected to exist in the gain; gain recovery when the pulse is not present can be calculated linearly, using less resources.

Gain

The gain is governed by the differential equation:

$$\frac{dg}{dt} = -\frac{g(t) - g_{ss}}{\tau_g} - g(t)\frac{|a(t)|^2}{E_{sat,g}}, \tag{3.29}$$

where g_{ss} is the small signal gain, $E_{sat,g}$ is the saturation energy, and τ_g is the recovery time. The second term governs the saturation; as was approximated in section 3.3, this term will dominate the dynamics when the pulse is present (although there is no necessity to make that approximation here), but its $a(t)$ dependence should rapidly fall off as the pulse's peak exits the gain. This leaves the first term - the longer-term recovery of the gain due to the pump - to govern the dynamics during the time where the pulse is completing its round trip in the rest of the cavity.

As this is a numerical model, it is conducted in time steps and each step's $\frac{dg}{dt}$ is obtained from the previous time step's values for the gain and field. The change in gain and the gain's effect on the pulse's field are then applied before calculating the next time step.

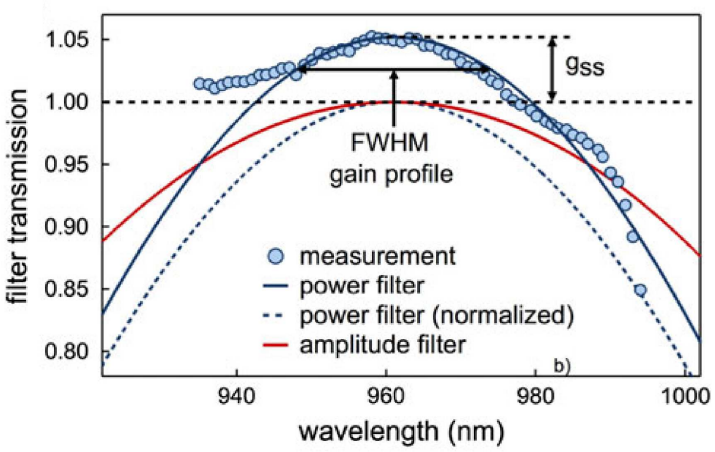

Figure 3.7: Example of how to attain a parabolic filter from experimental data. Fit parabola to experimental gain (solid blue), normalize it (dashed blue), and convert from intensity to field amplitude (red). From [50].

There are three other important aspects of the gain: a parabolic filter function (shown in Fig. 3.7), which is obtained from fitting to experimental data; a small, randomly-distributed noise amplitude; and a phase shift,

$$\Delta\phi(t) = -\frac{\alpha_g}{2}(t), \tag{3.30}$$

that arises from the Kramers-Kronig relation between the complex and real portions of the refractive index (α_g is a "linewidth enhancement factor" discussed in [53]). Each of these aspects is included as independent operators when the changes to the gain and field amplitude are applied.

SESAM

Similar to the gain, the loss is subject to a saturation diff. eq.:

$$\frac{ds}{dt} = -\frac{s(t) - \Delta R}{\tau_s} - s(t)\frac{|a(t)|^2}{E_{sat,s}}, \tag{3.31}$$

where ΔR is the modulation depth, and E_{sat} and τ_s are the analogous saturation energy and recovery time constant, respectively, for the loss. There will also be a phase shift analogous to Eq. 3.30, but depending on the linewidth enhancement factor of the absorber.

The complication in the absorber case is discussed more thoroughly in Chapter 4, with the result being that the relaxation time constant, τ_s, is actually better described by two separate time constants accounting for fast and slow carrier relaxation mechanisms. As such, the term relying on τ_s should actually consist of two terms relying on fast and slow time constants, τ_{fast} and τ_{slow}:

$$\frac{ds}{dt} = -\frac{s(t) - (A-1)\Delta R}{\tau_{fast}} - \frac{s(t) - (A)\Delta R}{\tau_{slow}} - s(t)\frac{|a(t)|^2}{E_{sat,s}}, \tag{3.32}$$

where A is an experimentally determined constant that balances the contribution to the overall recovery of the absorber between fast and slow processes (this will vary by design and growth quality of the SESAM). Again, because this equation is solved via numerical methods, it is relatively simple to calculate the contribution of each term per timestep, with each term being calculated separately from the previous timestep.

Miscellaneous terms

In addition to the saturable elements of the cavity, there are a handful of effects that need to be accounted for per round trip. The first is the group delay dispersion of the cavity, which is applied with a phase shift based on the GDD of the cavity, D:

$$\Delta\phi(\omega) = \frac{1}{2}D(\omega - \omega_0)^2. \tag{3.33}$$

The next effect that must be accounted for is the unsaturable losses in the cavity (output coupling, scattered light, etc.), which can be applied with a simple, constant percentage loss to the pulse envelope.

And finally, as mentioned above, to lighten the computation resources required for this model, the pulse can be defined within a specific time window and outside that time window the system can be updated with simpler expressions. If that is done, however, the pulse can drift within the time window. To keep the pulse from drifting out of the defined time window, then, it must be re-centered per round trip.

3.5.2 Many Body Model

In contrast to the rate equation model, which is heavily based on experimentally-measured input parameters, the many body model described in [51,52] relies on first-principle physics and requires significantly more computing power (super computer resources are required for any reasonable computing times).

With this model, the field is calculated all the way from the fundamental level of Maxwell's wave equation, which in one dimension (taken to be the center of a Gaussian spacial distribution) is:

$$\left[\frac{\partial^2}{\partial z^2} - \frac{n^2}{c_0^2}\frac{\partial^2}{\partial t^2}\right] E(z,t) = \mu_0 \frac{\partial^2}{\partial t^2} P(z,t). \tag{3.34}$$

The left side is standard propagation and the index is assumed to be invariant in frequency and possessing of no nonlinear component. This is not a bad assumption; as stated in section 3.4.1, the dominant portion of the GDD in a VECSEL comes from the Kramers-Kronig relation between the gain/absorption and the real portion of the index, which is accounted for in the macroscopic polarization, $P(z,t)$.

The VECSEL structure is set when defining the domain of the simulation, with the quantum wells being essentially delta functions in z where the polarization is not set to 0. At these points, the field is connected to the dynamics of the gain through the macroscopic polarization, which is the sum of the microscopic polarization states, $P(z,t) = \Sigma_{\lambda,\nu,k} p_{\lambda,\nu,k}$, across the band structure of the material(s) (where λ and ν are indices denoting specific conduction and valence bands, respectively, and k is the momentum in the plane of the quantum well). The dynamics of the microscopic polarization are then solved for using the semiconductor Bloch equations (SBE):

$$\begin{aligned}\frac{\partial}{\partial t} p_{\lambda,\nu,k} = -\frac{i}{\hbar} \sum_{\lambda_1,\nu_1} \left(e^e_{\lambda,\lambda_1,k}\delta_{\nu,\nu_1} + e^h_{\nu,\nu_1,k}\delta_{\lambda,\lambda_1}\right) p_{\lambda_1,\nu_1,k} \\ -i\left(n^e_{\lambda,k} + n^h_{\nu,k} - 1\right)\Omega_{\lambda,\nu,k} + \Gamma^{deph}_{\lambda,\nu}.\end{aligned} \tag{3.35}$$

In this colossal expression, there are several dependencies that must be taken care of. To start, the population of carriers (denoted: "electrons(holes)") will follow:

$$\frac{\partial}{\partial t} n^{e(h)}_{\lambda(\nu),k} = -2Im\left(\Omega_{\lambda,\nu,k} * p^{*}_{\lambda,\nu,k}\right) + \Gamma^{e(h),scatt}_{\lambda(\nu)} + \Gamma^{e(h),fill}_{\lambda(\nu)}, \tag{3.36}$$

where the relaxation rates, $\Gamma^{e(h),scatt}_{\lambda(\nu)}$ and $\Gamma^{e(h),fill}_{\lambda(\nu)}$, account for carrier recovery from the pump, $\Gamma^{e(h),scatt}_{\lambda(\nu)} = -\frac{1}{\tau_{scatt}}\left(n^{e(h)}_{\lambda(\nu),k} - f^{e(h)}_{\lambda(\nu),k}\right)$, and scattering from other parts of the band structure, $\Gamma^{e(h),fill}_{\lambda(\nu)} = -\frac{1}{\tau_{fill}}\left(n^{e(h)}_{\lambda(\nu),k} - F^{e(h)}_{\lambda(\nu),k}\right)$. $F^{e(h)}_{\lambda(\nu),k}$ and $f^{e(h)}_{\lambda(\nu),k}$ here represent quasi-equilibrium states immediately after the pulse has exited ($F^{e(h)}_{\lambda(\nu),k}$) and after recovery to the background Fermi distribution of carriers for a given pump level ($f^{e(h)}_{\lambda(\nu),k}$). These rates are approximations; full simulation of the many body effects which give rise to these relaxation rates could be included, but would substantially increase the demand on computing resources for minimal improvements in simulation accuracy [52, 54].

The Ω term present in both Eq. 3.35 and Eq. 3.36 is the effective Rabi frequency:

$$\Omega_{\lambda,\nu,k} = \omega_R + \frac{1}{\hbar}\sum_{\lambda_1,\nu_1,q\neq k} V^{\lambda,\nu_1,\nu,\lambda_1}_{|k-q|} p_{\lambda_1,\nu_1,q}. \tag{3.37}$$

In this expression, ω_R, the bare Rabi frequency, is based on the field and the appropriate dipole matrix element, $\omega_R = \frac{d^{\lambda,\nu}_k E(z,t)}{\hbar}$. $V_{|k-q|}$ is the screened Coulomb potential, a form of the standard Coulomb potential with an exponential damping term that accounts for the presence of background media. The screened Coulomb potential also appears in the Hartree-Fock renormalized single-particle energies from Eq. 3.35:

$$\begin{aligned} e^{e}_{\lambda,\lambda_1,k} &= \varepsilon^{e}_{\lambda,k}\delta_{\lambda,\lambda_1} - \sum_{\lambda_2,q} V^{\lambda,\lambda_2,\lambda_1,\lambda_2}_{|k-q|} n^{e}_{\lambda_2,k}, \\ e^{h}_{\nu,\nu_1,k} &= \varepsilon^{h}_{\nu,k}\delta_{\nu,\nu_1} - \sum_{\nu_2,q} V^{\nu,\nu_2,\nu_1,\nu_2}_{|k-q|} n^{h}_{\nu_2,k}, \end{aligned} \tag{3.38}$$

which describe, essentially, the band structure of the quantum well after Coulomb screening is taken into account. The final term from Eq. 3.35 is the polarization dephasing, $\Gamma^{deph}_{\lambda,\nu} = -\frac{1}{\tau_{deph}} p_{\lambda,\nu,k}$. This rate is also an approximation for more complex many-body dynamics in semiconductor gain media [54, 55]. A term describing spontaneous emission can be added to the description of the microscopic polarization (and is necessary to generate a pulse from noise), but is calculation intensive and has been carefully shown to produce the same results as simulations without spontaneous emission that are started with a seed pulse instead of noise [52].

Fig. 3.8 is an attempt at visualizing the dependency between the different parts of the model to hopefully aid in understanding the overall process. Again, as this is a numerical model, each term is updated based on parameter values from the previous time step rather than attaining an analytical solution. Fig. 3.8 attempts to show the order of this process where red terms are values from the previous timestep and the color gradient from red → blue shows the general order in which values are updated from the previous time step to attain the new field value.

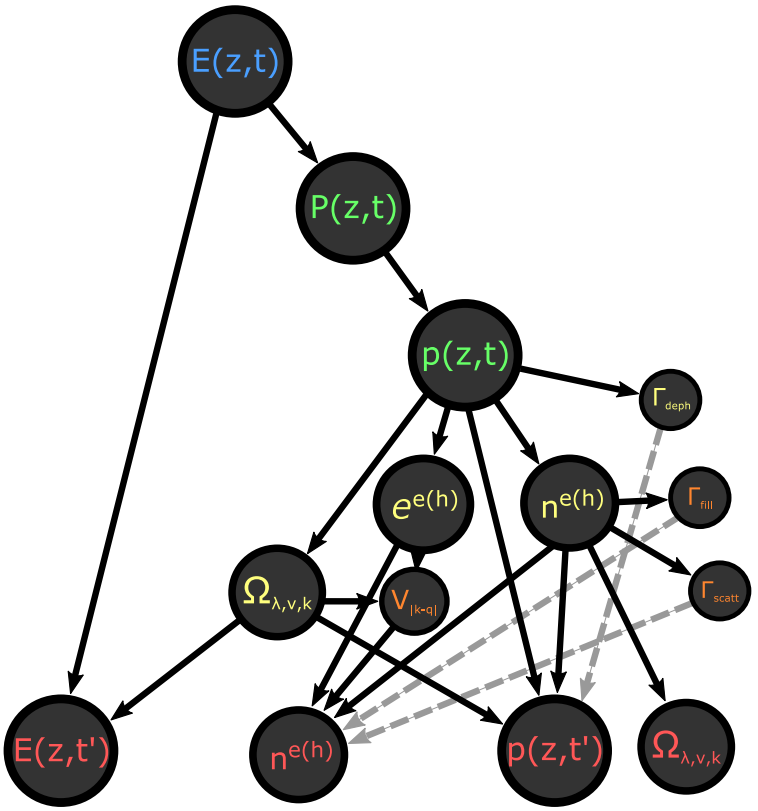

Figure 3.8: Dependence structure of many-body model. Red → Blue transition roughly indicates update sequence per time step. Grey arrows indicate rate approximations which could be expanded into full many-body dependence at the cost of lengthening simulation times by several orders of magnitude.

The above calculations are performed for every point in the cavity per time step and different device structures are investigated by altering the domain of the simulation and changing parameters such as the pump level (done by altering the background fermi distribution, $f^{e(h)}_{\lambda(\nu),k}$) and output coupling (a simple percentage loss applied to the field at a given point in the cavity). The pulse is tracked both in a local time window to display its temporal distribution and also a longer time window to track it for up to several tens of thousands of round trips.

The many body model has shown good experimental fit to pulses even in the sub-100fs regime where the accuracy of the rate equation model begins to become problematic [51,52] at the trade off of requiring substantially more computing power. Where the rate equation model can be run on a laptop computer, the full many-body simulation requires super computing resources and simulations can still take several hours to days, even with all above approximations included and a seed pulse.

3.6 Chapter Summary

This chapter concludes the background portion of this thesis by presenting the theoretical foundation for modelocking of lasers. Starting from the intuitive picture of frequency-domain active modelocking, several alterations were made to the master modelocking equation to eventually arrive at soliton regime modelocking, which accounts for realistic dispersion parameters in pulse formation. Finally, two models were briefly presented which numerically model modelocking in VECSEL devices specifically. The following chapters will cover the practical engineering concerns of modelocking in VECSELs specifically, addressing the design of the saturable absorber (Chapter 4), management of dispersion (Chapter 5), and design of the active region (Chapter 6).

Chapter 4

Saturable Absorbers

In this chapter, we discuss the practical design of saturable absorbers. We will motivate certain design choices by referencing the theoretical framework presented in the previous chapter and then present measurements regarding the dynamics of different absorbing region designs. Finally, a discussion of the supporting structure of the device (such as the DBR and spacer layer) will conclude the chapter.

With few exceptions, VECSELs are generally modelocked by means of a saturable absorber. From the body of theory in Chapter 3, we know that certain requirements will be placed on a potential absorber candidate. Broadly, an absorber must saturate quickly (more quickly than the gain), relax quickly, and introduce an acceptable level of total loss and dispersion to the cavity. This chapter will cover the practical design considerations for implementing these devices to modelock VECSELs.

In addition to dispersion, there are three specific, interconnected aspects to an absorber that are of particular importance when considering its practical viability:

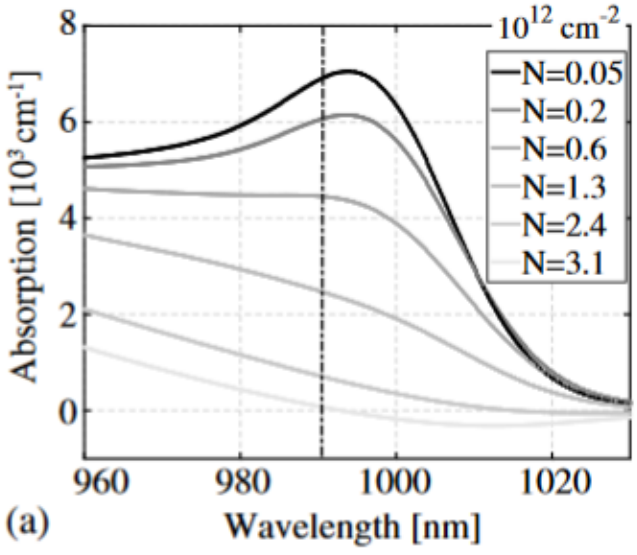

Figure 4.1: Simulated absorption curves for a quantum well SESAM designed for operation at 980nm. Maximum absorption corresponds to around 2% loss to the cavity. From [56].

the saturation fluence, which is the energy per unit area required to saturate the absorber to near-0 loss (the strict definition being to the 1/e point of the unsaturated absorption); the modulation depth, which is the difference in loss between saturated and unsaturated states; and the relaxation time, which is the time required for a saturated absorber to relax to its equilibrium state and be capable of absorbing at its full modulation depth again.

The former two principles are both displayed nicely in Fig. 4.1, which displays simulated curves for the absorption of a quantum-well-based Semiconductor Saturable Absorber Mirror (SESAM) at different carrier densities. The saturation fluence, here, is connected to the population density required to reduce the absorption to near 0, while the modulation depth can be seen as the difference in absorption values between the unsaturated (N = 0.05) and saturated (A $\simeq$ 0) curves. The maximum absorption values here correspond to about 2% loss to the cavity [56], which is consistent with the loss tolerance of the low-gain chips we use to achieve $\simeq$100 fs pulse durations (we

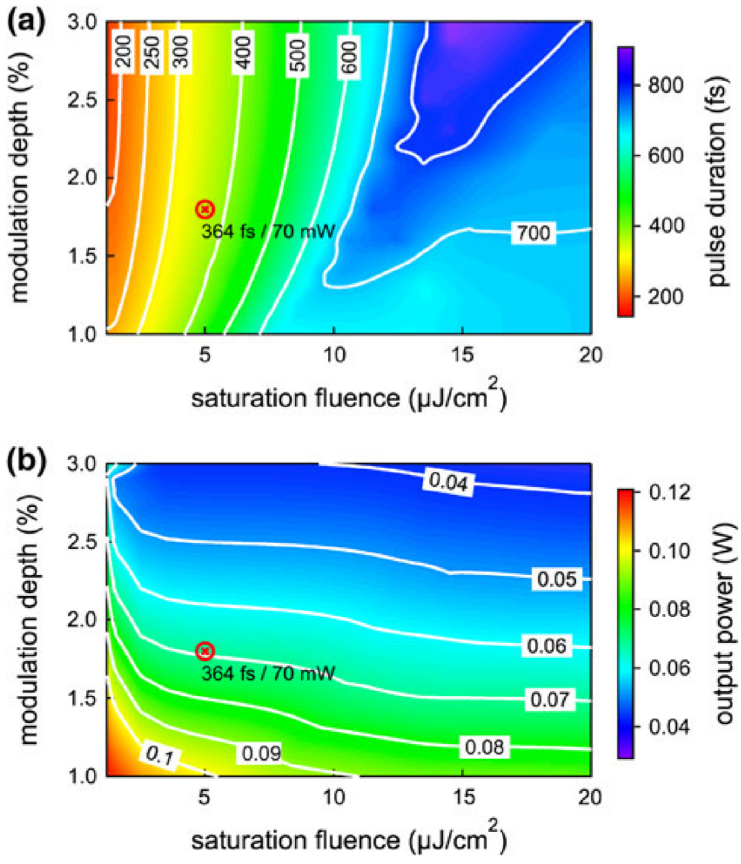

Figure 4.2: Simulated pulse durations (top) and output powers (bottom) dependent on the SAM's saturation fluence and modulation depth. From [50].

typically aim for 1-3% absorption in our SESAMs).

Results from the phenomenological rate equation model discussed in Chapter 3 indicate that the ideal conditions for short pulses in a modelocked VECSEL are a high modulation depth with a low saturation fluence (shown in Fig. 4.2) [50]. This is in general agreement with the many body model with one complication: while a low saturation fluence does correlate with shorter pulses, this effect becomes complicated when pulse durations approach the time scales of intra-band carrier scattering [52].

When pulse durations approach the timescales of intra-band scattering ($\simeq$100 fs), an additional level of detail is needed to describe the pulse shaping dynamics. In Fig 4.1 one can see that, spectrally, the saturation fluence is minimized and the modulation depth maximized near the excitonic resonance. If the lasing bandwidth is at higher energies, carriers can scatter out of the occupied states in the band at the lasing wavelength to lower energy states. For longer pulses, this would effectively serve as simply an increase to the saturation fluence. At first glance, then, one could conclude that operation near the excitonic resonance is desirable. When pulse durations approach the timescales at which intraband scattering happens, however, the quick recovery of the absorption due to this scattering is no longer a simple increase to the saturation fluence and can serve as an additional pulse-shaping mechanism. Efforts are still ongoing to conclusively determine where on the absorption spectrum of an absorber's quantum well is completely ideal for modelocked operation.

Recalling the fast and slow relaxation time scales discussed in Section 3.5, these scattering dynamics are in agreeance with the fast time scale discussed there. The longer relaxation time of the absorber is connected to the evacuation of excited carriers from the upper state of the absorber's band structure. To understand this process, we turn to a discussion of the design of the active portion of absorbers, as there are multiple approaches, each differing slightly in how they encourage carrier evacuation.

4.1 Ultrafast Pump/Probe Measurement

As a brief preface to discussion about active region designs, we must first familiarize ourselves with the pump/probe measurement. Time regimes faster than nanoseconds are difficult to measure; we do not presently have diodes or electronics fast enough directly measure phenomena that occur on the picosecond or femtosecond time scales. In order to conduct measurements on such timescales, then, we usually rely on measurement methods which are ostensibly conducted at speeds well-handled by existing instrumentation, but which can be used to calibrate a time axis on pico- or femto-second scales.

The principle example of such a measurement is optical cross-correlation, which dates back as far as the earliest pulsed lasers in the 1960s [57]. In this measurement, pulses are sent down two separate optical paths, one of which can have its length adjusted at will. The length adjustment of one path controls the timing of the two

pulses relative to each other, allowing well-controlled adjustment to "scan" one pulse through the other. Employing a physical process that depends on the two pulses being overlapped in time (such as Sum Frequency Generation) and measuring the signal relative to the path length difference, fine time scales can be calibrated (1 micron path difference corresponds to a time of flight difference of 3.3 fs in air) and the time dynamics of ultrafast phenomena can be reconstructed shot-to-shot as the pulse timing is varied. A very exaggerated display of this process is shown in Fig. 4.3

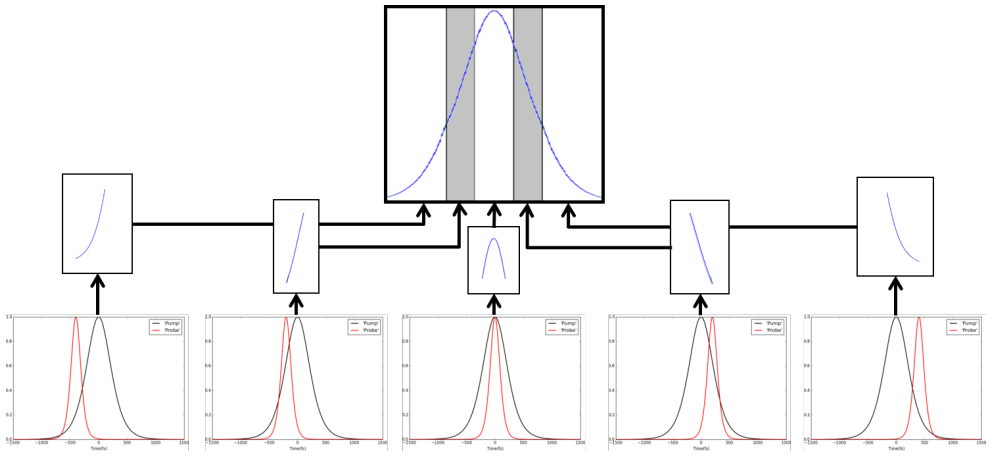

Figure 4.3: Exaggerated display of an optical cross-correlation, showing how different pulse timings can be used to reconstruct a time-domain signal from shot-to-shot measurements.

A pump/probe measurement is a measurement which is related to the optical cross correlation and designed to measure material responses on ultrafast time scales. This is generally done by taking a strong ultrafast pulse train and clipping a minor amount of power from it to be used as the probe beam. The strong pulse train (the "pump") is then used to excite the material of interest and the weaker probe, being too weak to significantly affect the material, is passed through the sample and measured to detect any changes due to the material's response to the pump. A schematic of this setup is shown in Fig. 4.4.

The data in the following sections will be reflectivity data, where the pump pulse is expected to saturate the absorption of a SESAM, yielding higher reflectivity when the probe pulse trails the pump in time, but it bears mentioning that pump/probe is not restricted to this type of measurement — it is suitable for *any* material response that will produce a measurable effect on the probe pulse.

The resolution limit in a correlation measurement is going to be the largest value of the following: the timing jitter between the pump and probe pulses, the resolution of the scan rate (i.e. the smallest stepper motion on a mechanical translation stage), or the duration of the pulses. In the measurements below, a single pulse train is split

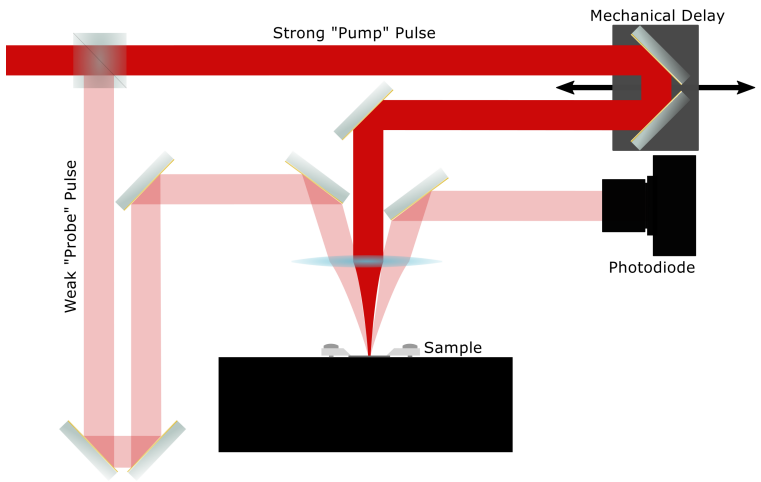

Figure 4.4: Schematic for Pump/Probe Measurement.

by amplitude with a beam sampler to form the probe pulse train, so the timing jitter will be negligible. The stage we use has a minimum increment of motion of .5 μm, corresponding to a total path length change of 1 μm and thus a fundamental time step of 3.3 fs. Therefore, the fundamental resolution limit is clearly the pulse duration of the pump and probe pulses, which is around 100 fs FWHM.

For our specific pump/probe setup we employ a commercial fiber laser system (Menlo Orange) that emits around 100 fs pulses with a center wavelength of around 1040 nm at an average power of slightly over 1 W and a repetition rate of 80 MHz. In addition to the timing resolution this laser provides (discussed above), the average power of 1 W provides more than enough power to saturate our samples — we usually have to attenuate the pump pulse to 100-200 mW of average power for measurements — and the repetition rate of 80 MHz gives a shot-to-shot timing of 12.5 ns, which is more than enough ensure complete relaxation of the device between pulses. We attenuate the probe to <10 mW to ensure that its influence on the carrier dynamics of the structure is negligible.

There are experimental improvements that will significantly aid in taking quality pump probe data, as well. A chopper can be inserted into either arm in order to use a lock-in amplifier to boost signal strength and reduce noise (or, indeed, in *both* arms). It is highly recommended that half waveplates be used to enforce orthogonal polarizations for the pump and probe beams, as that will minimize nonlinear measurement artefacts such as four wave mixing (FWM) or 2-photon absorption. The probe beam can also be further split into probe and reference beams for use with a differential photodiode. Finally, contrary to Fig. 4.4, the mechanical delay can be used in either the pump or the probe as convenience dictates and it is sometimes preferable to use separate lenses and more tightly focus the probe beam, making it less sensitive to any drift in beam pointing during the measurement.

Understanding the method for these measurements is important in deciphering the reflectivity data in the following sections as different active region designs are explored.

4.2 Practical Design (active region)

Active region designs tend to resemble scaled-down versions of VECSEL active regions, with the most widely used being Semiconductor Saturable Absorber Mirrors (SESAMs) based on single quantum wells or a single layer of quantum dots. There are a handful of exceptions where Graphene Saturable Absorber Mirrors (GSAMs) [58] and carbon nanotubes [59] have been used.

Other than targeting the correct lasing bandwidth for absorption, one of the primary goals of active region design is promoting carrier evacuation as quickly as possible. Although referred to as the "long" recovery time constant when compared to the $\simeq$100 fs time scales of intraband scattering, it is still advantageous to evacuate carriers from the upper state of the energy band on as fast a time scale as possible.

There have been multiple strategies for accomplishing this fast carrier relaxation:

4.2.1 Low Temperature Quantum Well Growth

When grown at low temperature ($\simeq$250-350° C), GaAs-based material systems will succumb to the formation of defects in the crystal lattice [60]. This is not generally advantageous, as it produces scattering sites where excited carriers can escape the upper portion of the band structure without recombining via stimulated or sponta-

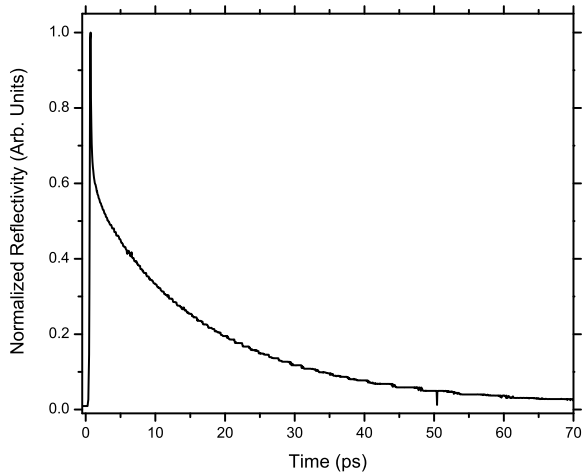

Figure 4.5: Pump probe measurement of LT GaAs SESAM. The spike in reflectivity can be understood as saturation of the absorbing quantum well due to the strong pump pulse. Relaxation occurs via two visibly different time constants.

neous emission. Such defects in a quantum well intended for gain can actually be disastrous to the modelocking performance of the device.

For an absorber, however, these defect sites present an opportunity for the exact sort of fast, non-radiative carrier relaxation discussed above. Thus, one strategy for achieving fast relaxation times in a SESAM is to grow a device with standard DBR and barrier regions, but with a single quantum well grown at low temperature. As shown with an ultrafast pump/probe measurement we conducted of a SESAM designed in this manner in Fig 4.5, (long) time constants of slightly under 100 picoseconds can be attained from such a device.

The unfortunate downside to using low temperature growth is that the distribution of defects is statistical in nature, with control over the temperature giving a very coarse adjustment of defect density, but not ensuring precision and uniformity for consistent absorber quality.

4.2.2 Surface Quantum Well

Another approach is to grow a single quantum well in the standard manner, but place it a few nanometers below the surface-air interface. The core principle is that with such a thin barrier between the well and a surface interface, excited carriers are capable of tunneling across the barrier and recombining on the surface. In this manner, picosecond relaxation times are possible.

Naively, it would seem that the barrier between the well and the surface should be as thin as possible, but we have consistently observed unrecoverable device degradation when the cap layer is grown too thin. Our initial investigation of surface QW absorbers, in collaboration with a fabrication team at University of New Mexico, consisted of designs with wells placed 3 nm, 5 nm, and 7 nm below the surface interface. The 3 nm devices gave acceptable modelocking performance, but degraded almost immediately during operation and even degraded in open air over the course of several days. The 5 nm and 7 nm samples more effectively displayed the careful balance required when designing surface quantum well (SQW) SESAMs.

Relaxation rate measurements we conducted for 5 nm and 7 nm devices are shown in Fig 4.6. The difference in longer time constants here is slight, but produces a noticeable effect on pulse duration, with 5 nm cap SESAMs being capable of generating pulses as short as $\simeq$100 fs and 7 nm cap devices with the same VECSEL and similar GDD profile being unable to generate pulses significantly shorter than 1 ps. That said, Fig 4.7 shows the time-dependent power of a VECSEL modelocked using a 5 nm SESAM. A clear degradation in performance occurs over tens of minutes which, after careful elimination of other factors that could cause power loss, can be attributed to the degradation of the absorber. In our investigations, devices with a quantum well 7 nm below the surface did not show such degradation.

Explanation for this degradation is, at the time of this thesis, still ongoing. Initially, oxidation in the quantum well due to oxygen diffusion across the small cap layer

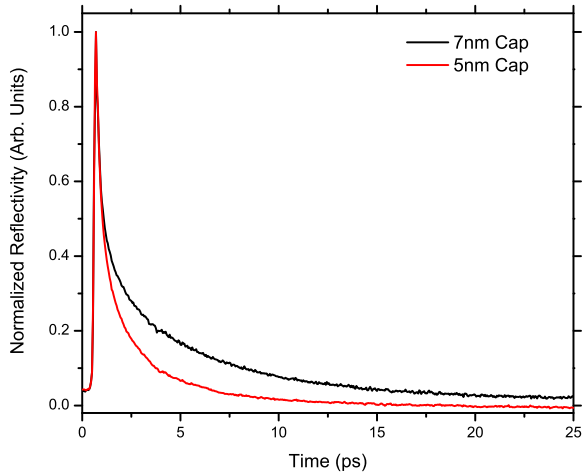

Figure 4.6: Pump/probe measurements for 5nm and 7nm cap SQW SESAMs.

was suspected. However, transmission electron microscopy (TEM) and Energy Dispersive Spectroscopy (EDS) linescan measurements conducted by our collaborators at UNM have suggested that oxidation occurring in the AlAs portion of the DBR is a root cause. An investigation into whether the damage is field-related and/or exacerbated by the stronger field values in short (100-200 fs) pulses as opposed to longer ($\geq$ 500 fs) will also be important to fully characterizing the process.

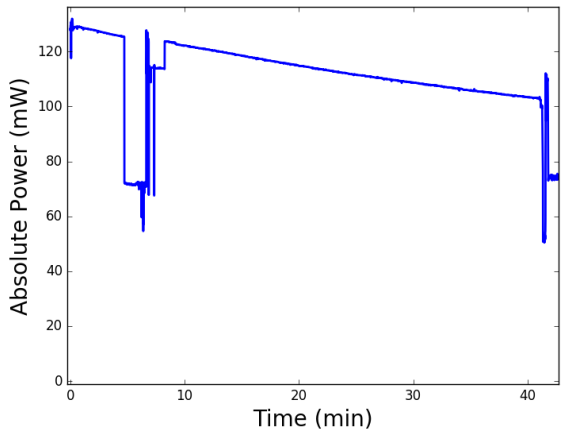

Figure 4.7: Power decay in 5nm cap Surface Quantum Well SESAM. The power decay here is indicative of degradation in the SESAM. Points of dramatic power reduction are due to a break in the modelocking state of the laser; modelocking was recovered after breaking at around 10 minutes, but was not recoverable upon breaking again after 40 minutes. This device emitted 250 fs pulses

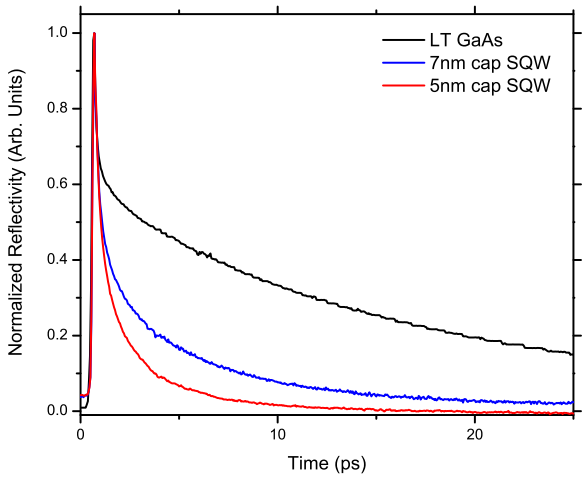

Figure 4.8: Pump/probe measurements showing a comparison of relaxation rates for LT-GaAs and SQW SESAMs with 5nm and 7nm cap layers, respectively.

Fig. 4.8 shows a comparison of the relaxation data presented above for LT-GaAs and SQW structures; our surface quantum well SESAMs have produced much more desirable relaxation behavior than the commercial Low-temp GaAs SESAMs which we have used. This relationship, however, is not necessarily intrinsic to the LT GaAs material system and could be the result of simply poorly-controlled LT GaAs growth. It is important to note, then, that while the tighter design constraints placed on SQW SESAMs can lead to device degradation if the design does not insulate vulnerable material layers appropriately, the more uncontrolled nature of defect distribution in LT-GaAs growth could present more headache in its inconsistency.

4.2.3 Quantum Dots

Quantum Dots (QD), the colloquial name for nanometer-scale structures with 3D quantum confinement (compared to the 1D confinement of a quantum well), have been shown to be effective saturable absorbers [61].

The QD structures used in our lab, also in collaboration with UNM, are based on InAs/GaAs interfaces for operation around 1 μm. The lattice constant mismatch between the two materials causes disruptions in the uniformity of the crystal growth, growing small "islands" of material which can then be buried by substrate material to form the quantum dots (referred to as "Stranski–Krastanov growth" [62]).

The carrier relaxation mechanism is ostensibly the same as in an SQW SESAM — carrier scattering at the surface interface of the device. However, the increased thickness of quantum dots (16 nm at 1040 nm compared to a 1040 QW's 8 nm thickness), potentially combined with the different material properties of InAs, seem to help prevent degradation while not sacrificing other aspects of the absorber's perfor-

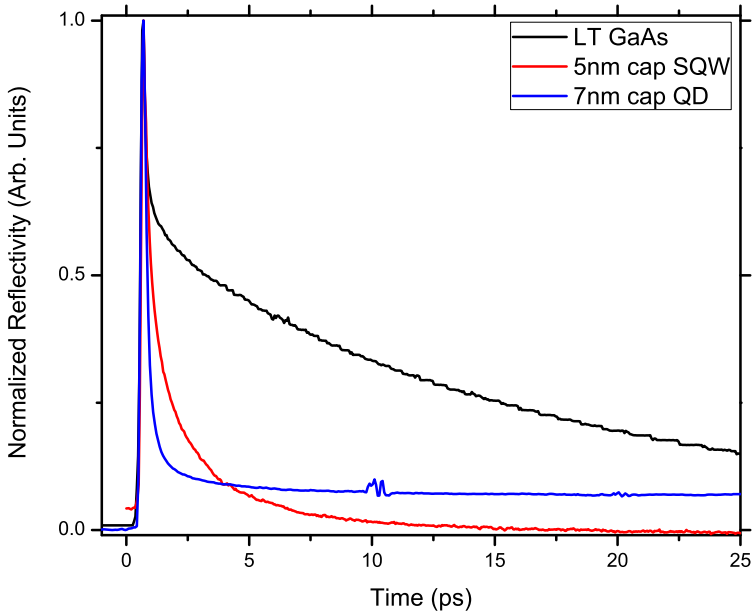

Figure 4.9: Pump/probe measurements showing a comparison of relaxation rates for LT-GaAs, SQW, and QD SESAMs.

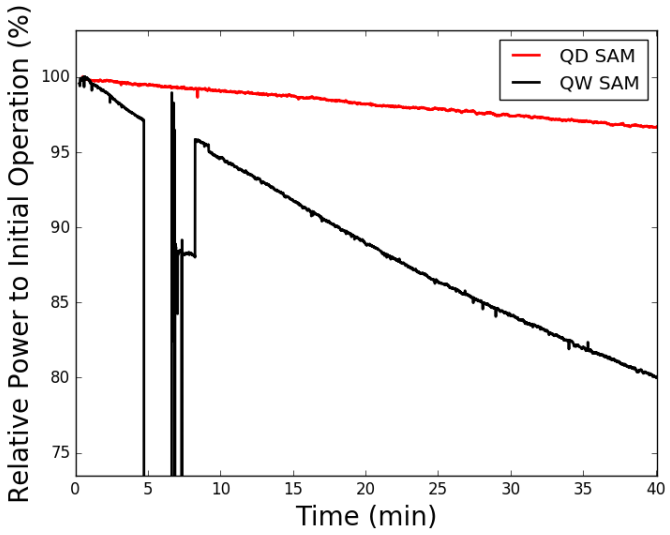

Figure 4.10: Power decay comparison between Quantum Dot and Surface Quantum Well SESAMs. The SQW SESAM here is the same as shown in Fig. 4.7.

mance. Relaxation rates we measured for QD and surface QW devices at 1040 nm are shown in Fig. 4.9. The QD SESAM measurement seems to display a dominant fast time constant (with almost 80% of the modulation having recovered in less than 1

ps), but a long time constant that is considerably longer than quantum well SESAMs. The egregiously long (˜ns) longer time scale could indicate that for this particular QD SESAM, surface tunneling was not occurring as expected, leaving carriers with recombination as the only relaxation mechanism from the conduction band (this is supported by the nanosecond relaxation rate, which agrees with the expected time scale of recombination processes). We have still experienced some long-time degradation in QD SESAMs, but as seen in another power plot in Fig. 4.10, operation (in this case, with the same VECSEL structure as 4.7) can be far more stable than an SQW SESAM vulnerable to degradation (this could potentially be to QD being more easily saturated, allowing less tightly focused beams and lower field intensities).

4.2.4 Graphene

The band structure of Graphene is in a lot of ways highly ideal for use as a saturable absorber. With zero band gap and nearly linear density of states near it, carriers can easily scatter out of the portion of the energy band near the lasing bandwidth and subsequently recombine non-radiatively across the band gap [52]. A single layer of Graphene can also provide modulation depth of up to 5% [58].

A Graphene Saturable Absorber Mirror (GSAM) has been demonstrated to produce pulses on the order of hundreds of fs at low power levels [58]. The primary drawback of GSAM use is in its burn threshold, exacerbated by the desire to focus tightly onto the SAM inside a cavity for faster saturation of the device, which can easily be reached in the 10s of mW regime of output powers. This is also problematic given that a GSAM will have a high saturation fluence, meaning that if short enough pulses are not supported by the other parameters of the cavity (notably the gain bandwidth and GDD), it may not be possible to saturate the absorption at power levels below the Graphene burn threshold.

4.3 Practical Design (Supporting Structure)

Although the overall structural design of a SESAM is much simpler than that of a VECSEL, it is still nontrivial and highly important to the ultimate modelocking performance of the device.

The DBR is generally of the same design as the VECSEL structure, although if the VECSEL DBR is hybridized with Au/Ti patterning, it is generally permissible to forgo that process and simply grow a thicker DBR on the SESAM to achieve acceptable reflectivity values (>99.9%). As VECSELs push towards pulse durations of 100 fs and shorter, the accuracy of this DBR becomes increasingly important. A transform-limited, 100 fs pulse at 1 micron requires a >10 nm bandwidth, across which the DBR needs to be nicely uniform in its reflectivity. This is achievable, but problems (particularly with dispersion) can arise if the DBR is offset spectrally such that the lasing bandwidth approaches either edge of the stopband.

Fig. 4.11 shows an example SESAM design for 1040nm operation (in this case, an SQW-based SESAM). The first thing to note is the spacing layer employed to control the placement of the active portion (be that a quantum well or a layer of dots) relative to the intracavity field. Depending on this placement, the absorption of the SESAM can range from 0% (at a field node) to levels high enough that it would prevent lasing entirely if the device does not have an anti-reflection coating. We typically aim for 1-3% absorption. The thin cap layer across which carriers are expected to tunnel before scattering away at the surface interface can also be seen. Finally, an anti-reflection coating is present on the surface; AR coatings are quite often employed on both VECSEL and SESAM devices to help manage dispersion. Coating design will be discussed in much greater detail in the following chapter.

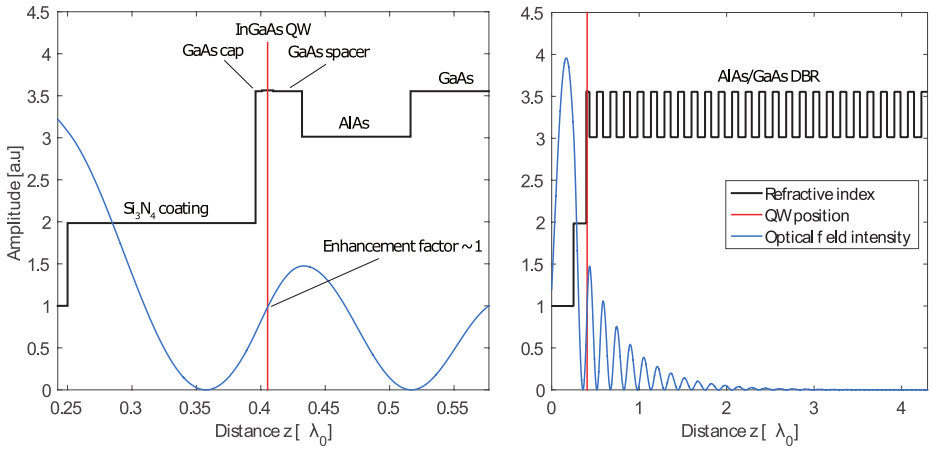

Figure 4.11: Design of a surface quantum well based SESAM.

4.4 Chapter Summary

This chapter has discussed the design of saturable absorbers, starting with the general goals when designing an absorber and proceeding to the strategies for designing the active portion and supporting structure of an absorber. Chronologically, this issue was one of the first we faced when beginning to work on modelocking VECSELs in the ultrafast regime, with the first consistent sub-ps results attained in our lab coming from custom-designed absorbers from collaborators at University of New Mexico. Iterations on absorber designs then followed as improvements to other aspects of our lasers were made.

The next chapter will discuss the next major issue (and persistent, recurring headache) we encountered in dispersion management.

Chapter 5

Dispersion Management

In this chapter, we will discuss the management of dispersion in ultrafast VECSEL devices. Like in Chapter 4, a brief link to the underlying theory will serve as motivation and then work and measurements we have conducted in dispersion management will be presented.

As discussed in Chapter 3, managing dispersion is a fundamental challenge in the operation of any modelocked laser system. Mitigating the temporal drift of a pulse's constituent modes relative to each other per round trip is vital to the formation of a stable modelocking state and a significant factor in the ultimate pulse duration emitted.

If the emitted pulse is simply stretched in time ("chirped"), but still possessive of the spectral bandwidth for the desired pulse duration, techniques to compress it to its transform limit have existed since as far back as 1969 [63]. In some laser systems (notably rare earth fiber and Ti:Sapphire), it is actually commonplace to put a prism or grating compressor inside the cavity itself, but this practice has numerous challenges in VECSELs due to their lower gain and lifetime-imposed restrictions on cavity length. Use of an external compressor, however, is not always ideal and in some cases not even possible, as poor dispersion properties or other factors will prohibit a compressible modelocked state from arising. Thus, we often seek to manage dispersion within the laser itself as much as possible before relying on external compressors.

The rate equation model presented in Section 3.5 predicts significant pulse lengthening for Group Delay Dispersion (GDD) values of less than 100 fs^2 [50]. Fig. 5.1 shows the level of significance that minor amounts of GDD can have by displaying

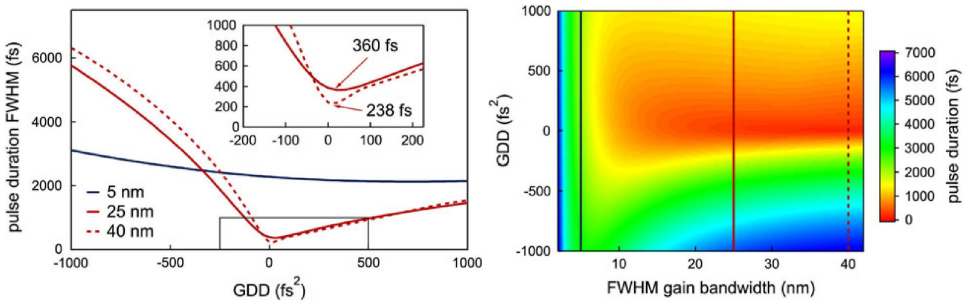

Figure 5.1: GDD dependence in rate model simulations. (Left) For two specific gain bandwidths denoted as lines in (Right) the full parameter space. From [50].

simulated pulse durations against VECSEL gain bandwidths—which would be supportive of certain fourier-conjugate pulse shapes under ideal conditions— and net cavity GDD. Results from the Many-Body model discussed in Section 3.5 are in agreement, with the shortest pulse durations nearly always coming from structures simulated to have flat, near-0 GDD profiles across the lasing bandwidth.

This requirement for flat GDD profiles near 0 fs^2 presents a challenge when utilizing VECSELs that have traditionally been grown for high power CW emission. A dispersion profile of one such device — a 10 QW RPG structure designed for operation near 1 micron, which we have measured using a commercial white light interferometer (KM Labs "Chromatis") — is shown in Fig 5.2. As can be seen, there is a dramatic resonance feature near the quantum well absorption line - where desired GDD values are less than even 100 fs^2, here the peak GDD values exceed $\pm$10,000 fs^2. While this is an unpumped device, a similar resonance feature would be expected during lasing operation as the primary causes can be understood to be the microcavity GTI resonance and the strong interaction between the quantum wells and the intracavity field. The microcavity is present regardless of the pump rate and both strong absorption and strong gain yield significant changes to the dispersion via the Kramers-Kronig relations.

There are several strategies that can be employed to either correct the dispersion

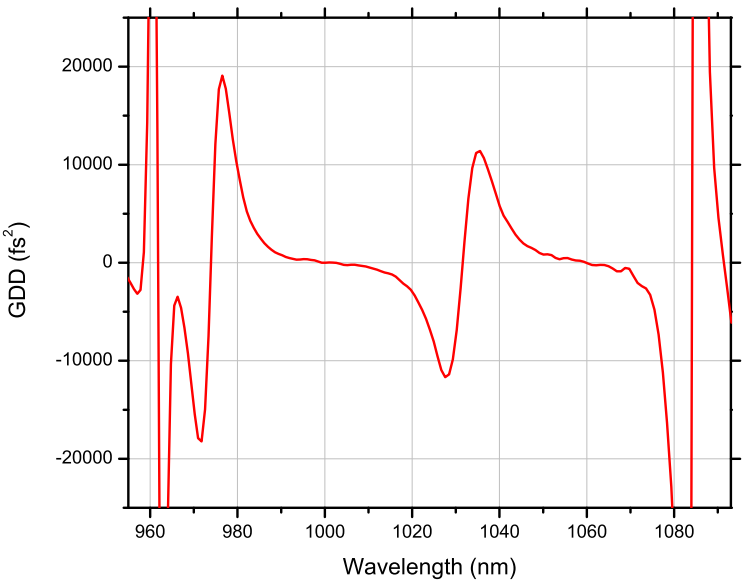

Figure 5.2: Group Delay Dispersion for an RPG-design VECSEL with no AR coating. Intended for use at around 1020-1040 nm; the severe GDD fringes on the edges indicate the edge of the DBR stopband; the resonance feature near the intended operation wavelength is understood to be caused by the device's microcavity.

of the device directly or else manage it during operation. As one of the primary causes of the strong GDD resonance feature in Fig. 5.2 is the strong interaction between the wells and the field, one strategy is to shift the wells themselves away from the antinodes of the design wavelength's standing wave. This has the effect of broadening and flattening the overall gain bandwidth of the device, reducing the wavelength selectivity of the gain and the contribution to the GDD from the Kramers-Kronig relationship. Such designs will be discussed more in depth in the next chapter, as they are important when considering a unique aspect of semiconductor band structures on modelocking operation, but also merit a mention when discussing the dispersion properties of VECSELs. A discussion of the more GDD-specific strategies is presented here.

5.1 Element Compensation

Although the level of precision by which we can engineer nanoscopic structures is impressive and ever increasing, practical reality dictates that controlling the dispersion of a semiconductor mirror to within 5 to 10 fs^2 is still an extremely tall order. Practical concerns such as variance in material thicknesses and composition will always yield some variance in the resulting dispersion of a device. A simple solution, often used in other ultrafast systems, is — rather than exhausting resources on controlling the manufacture of one sample — to have a collection of samples (and laser mirrors) with well-characterized properties from which you can select samples which nicely work together.

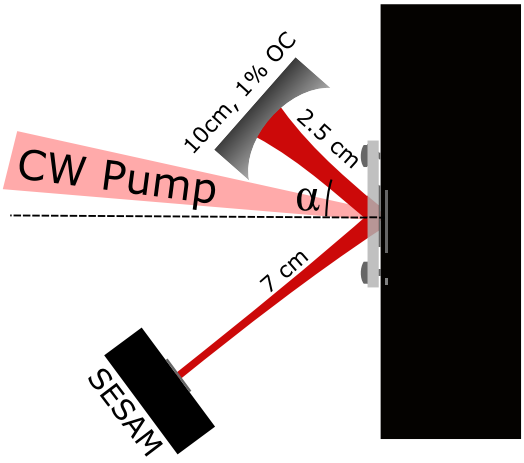

Figure 5.3: V-cavity used in element compensation and angle tuning experiments. The angle, α, was held at a constant 15°while different SESAMs were tested (shown in Fig. 5.4). In Fig. 5.6, the SESAM was held constant and the output pulse duration was recorded for different angles.

In VECSELs this is most commonly done by balancing the GDD of the VECSEL with the GDD of the SESAM used to modelock it. An example of this, which Dr. Maik Scheller and I contributed to [52], is shown in Fig. 5.4, which shows the GDD profiles of a VECSEL and several SESAMs, with the corresponding pulse durations emitted by each pairing. The net round-trip GDD, assuming VECSEL and SESAM to be the only significant contributors, would consist of the sum between the SESAM curve and double the VECSEL GDD (corresponding to two passes through the gain). As the VECSEL used here possesses strongly positive GDD, one expects the desired net GDD profile of flat and close to 0 fs^2 to be approached (in reality, it is unlikely to attain a perfect net GDD profile with these samples) when the SESAM possesses a more strongly negative GDD. That SESAM, as expected, produced the shortest pulse durations when modelocking the VECSEL.

We took the data here using one of our earlier MQW devices which was optimized for modelocked operation near 990 nm. This design featured eight quantum wells stacked in pairs near the field anti-nodes (i.e. two wells per antinode on either side of the field peak). The setup, shown in Fig. 5.3, is a simple V-cavity with a 1% output coupler and it emitted average powers in the range of hundreds of milliwatts. The pulse durations varied depending on the experimental parameters, shown respectively in Figures 5.4 and 5.6.

We have not employed this technique much in our lab, but it should also be noted that other intracavity elements can be used as dispersion compensation. For example, a 100 μm thick parallel plate of fused silica will add about 2 fs^2 of material dispersion. If used uncoated and at normal incidence, this would not be an attractive intra-cavity element due to the spectral filtering of the Fabry-Perot resonance of the parallel plate, but if placed at Brewster's Angle or used with a suitable anti-reflection coating, the spectral filtering and loss will be negligible, leaving only the material dispersion. One can easily see, then, how different thicknesses of glass or fused silica can be used to add minor amounts of positive dispersion to a cavity. In other laser systems, dispersion compensating mirrors are often employed, but since we generally balance dispersion between the VECSEL and SESAM, we employ mirrors with as flat of GDD profiles as possible.

It is interesting to note that in our collaboration with University of New Mexico for SESAM growth, we encountered the issue that the MBE reactor used to grow samples had a lateral drift in certain growth parameters from the center of the wafer. This resulted in the spectral shift of the dispersion profile by a couple nanometers per centimeter in spatial distance from the center of the wafer. After careful consideration, we actually encouraged them to not completely correct this drift, as (within reason) the relatively smooth spectral drift in the dispersion profile provided a degree of freedom to finely control the dispersion compensation between VECSEL and SESAM through the selection of SESAM devices cleaved from different radial positions of the wafer.

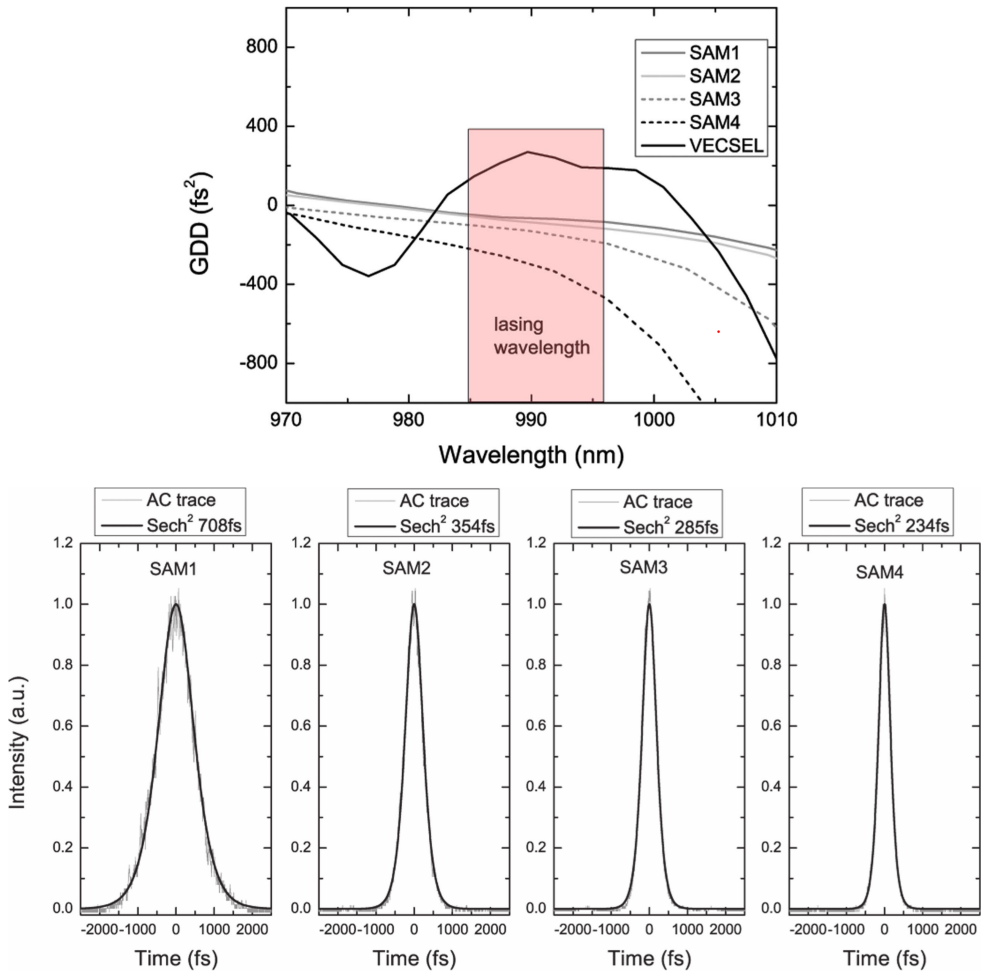

Figure 5.4: Demonstration of GDD compensation between multiple elements of a cavity. (Top) GDD profiles for a VECSEL and four SESAMs with overlaid lasing bandwidth. (Bottom) Pulse durations resulting from the VECSEL in use with each SESAM above shorter pulses result when the sum of the elements' GDD is closer to 0. From [52].

5.2 Anti-Reflection Coatings

Anti-Reflection (AR) coatings are a standard method of reducing dispersion of optical elements used in the construction and use of modelocked lasers. In VECSELs, precise AR coatings can be employed on both VECSEL and SESAM device surfaces to significantly help manage the net dispersion of the cavity.

The primary function of the anti-reflection coating is to reduce the influence of the Gires-Tournois Interferometer (GTI) resonance between the DBR and the cap layer of the device. This is most often accomplished by introducing an intermediary layer that "steps down" the index of refraction across multiple interfaces in lieu of a sharp index contrast at a single interface. As the fresnel reflection coefficient of an interface strongly depends on the index contrast between materials and the effect of the GTI resonance depends strongly on the magnitude of the reflection as light exits the device, the Q-factor of the device's microcavity is thus significantly lowered. As a secondary effect, the change in effective thickness due to an AR-coating will also spectrally shift the field enhancement caused by leftover GTI effects. This can be exploited to raise or lower the field enhancement factor, depending on the design of the device's active region - devices with a design wavelength in resonance with its microcavity will seek to lower the field enhancement (thus removing spectral filtering) while anti-resonant designs will sometimes seek to raise the enhancement (thus providing more gain).

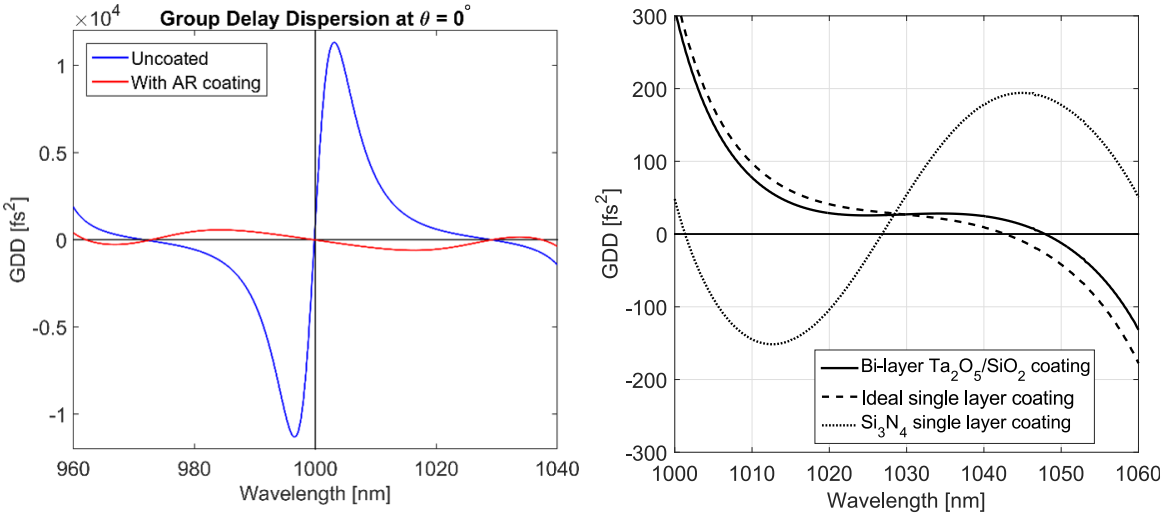

(a) (b)

Figure 5.5: Simulated GDD profiles for (a) an RPG device designed for high power, showing the significant reduction in GDD around the microcavity resonance when an AR coating is applied and (b) an MQW structure designed for modelocking, showing the results of three different AR coating designs.

Thus, a precisely designed AR coating can be instrumental in both attaining GDD values near 0 fs^2 and in flattening out resonance features due to field interactions with the active portions of a device. Fig 5.5 shows the effect a well-designed AR coating can have on a VECSEL designed for modelocking and even on a high-power RPG structure with a pronounced resonance feature in its GDD profile. The GDD curves here are from a matrix transfer simulation written by Dr. Alexandre Laurain. After measuring a device's actual GDD at room temperature and 0°angle of incidence, the simulated structure is then modified to fit the experimental measurement. Simulated anti-reflection coatings can then be added to the device.

The design strategy for these coatings is heavily dependent on the structures they are intended for. It is interesting to note that an ideal coating would consist of a gradient-index material that smoothly transitions from the index of the material to 1, but this is not currently possible with available material science. The next best option is to employ multi-layer coatings, which can range from simpler, higher-tolerance bi-layer coatings (e.g. [56]) to complex, 7-8 layer designs involving both semiconductor and dielectric layers (e.g. [12]). Our lab tends to employ bi-layer coatings on VECSELs and single-layer coatings on SESAMs, but this is not a strict rule and is always subject to optimization using simulated structures before deposition. A general procedure is:

- Measure GDD and linear reflectivity of the VECSEL.
- Fit matrix transfer simulation to the measured data.
- Add AR coating to the simulated structure and optimize layer design (thickness; material composition) using simulated GDD and gain curves.
- Deposit VECSEL coating.
  - Re-measure GDD and reflectivity after first layer and re-optimize second layer before deposition if necessary.
- Measure GDD for the coated VECSEL and the SESAM intended to be used with it.
- Fit simulation of the SESAM to the measurement.
- Optimize simulated SESAM AR coating to compensate for coated VECSEL's GDD profile.
- Deposit SESAM coating.
  - Re-optimize between first and second layers if necessary.

Although the final parameters for coating deposition are almost always the result of fine tuning via trial and error in a simulated environment, there are a handful of useful "ballpark guess" starting points. AR coatings tend to be close to $\frac{\lambda}{4}$ in total thickness (where λ is the center wavelength of the desired lasing bandwidth), with $\frac{\lambda}{8}$ being a good starting guess for each layer of a bi-layer coating.

Our lab generally uses a combination of Silicon Dioxide (SiO_2; $n = 1.45$), Silicon Nitride (Si_3N_4; $n = 2.016$), and Tantalum Oxide (Ta_2O_5; $n = 2.275$)). This is simply due to availability of these materials in our facilities, however, and not a commentary on the ideal nature of any materials. A fused silica layer (given as "SiO_xN_y"; $n \simeq 1.5$) was used in [12].

5.3 Cavity Angle Tuning

Simple geometry dictates that interacting with a material stack under differing angles of incidence will alter the effective thickness of the material layers. In VECSELs, this will generally serve to blue shift the optical properties of the device, including the favored emission wavelength and the dispersion profile. Thus it is possible, as Fig. 5.6 shows, to employ a V-cavity geometry (in this specific case, the one found in Fig. 5.3) and alter the emitted pulse durations of a single VECSEL/SESAM pair by changing the fold angle of the cavity. In this case, the same VECSEL from Fig. 5.4 above is used with the SESAM "SAM2."

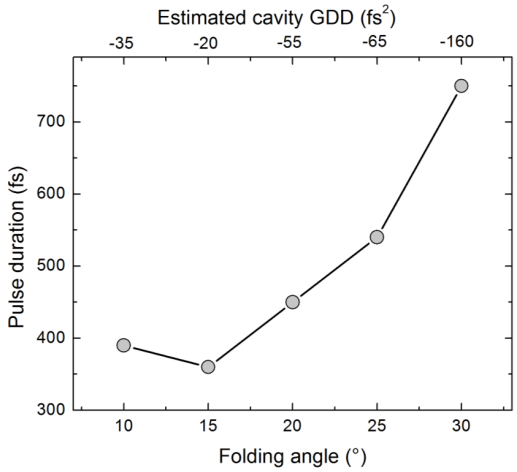

Figure 5.6: Angle compensation of GDD in a V-cavity. "SAM 2" is used in conjunction with the VECSEL from Fig. 5.4 under different angles. The GDD is estimated based on two passes through the VECSEL at a given angle and one pass through the SESAM.

Ultimately, the plot in Fig. 5.6 is still essentially GDD management through element compensation. The angle tuning shifts the GDD profile of the cavity element

which is at an angle (in this case, the VECSEL) relative to the static GDD profile of a consistent element (in this case, the SESAM, which is used as an end-mirror and is thus always at normal incidence). It is this relative spectral shift of the two GDD profiles that alters the net GDD with the cavity angle; if there were no static element, it is likely that the lasing wavelength would also blueshift with angle and effect relied upon to balance the dispersion would be greatly diminished.

That said, the dispersion profile can be shifted spectrally by tens of nanometers from its distribution at normal incidence if wider fold angles exceeding 30-40° are employed. This creates the possibility displayed with simulated GDD curves in Fig. 5.7: that the GDD resonance feature familiar to RPG structures (in this case, with a single-layer AR coating applied) can be shifted so far that it can at least partially compensate itself.

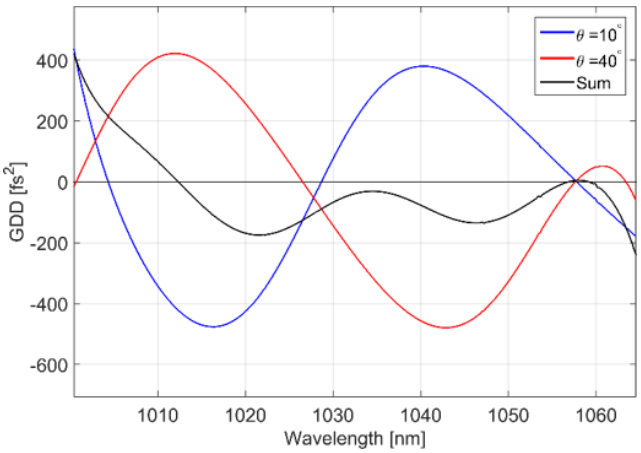

Figure 5.7: Simulated GDD of a VECSEL under narrow (10°) and wide (40°) angles. The sum of the two curves is also plotted to demonstrate that the could serve to compensate each other.

This could be done with multiple VECSEL devices or potentially by a single VECSEL under multiple angles per round trip of the cavity.

5.3.1 The F-cavity

The double-fold cavity shown in Fig. 5.8, which we refer to as the "F-cavity," is a cavity design we created to utilize this angle tuning effect to balance a VECSEL's GDD with itself.

The structure used to investigate this idea was an RPG design (with ten InGaAs QWs) with a single-layer AR coating. This selection was made in an attempt to balance the gain of the device against the spectral filtering of the microcavity and the aberrant GDD resonance of RPG designs. The peak GDD values of an uncoated RPG are simply too high and the spectral filter imposed by the microcavity of such a device

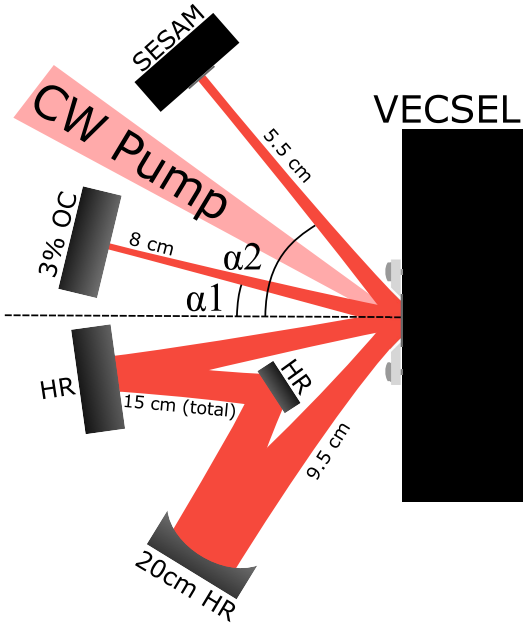

Figure 5.8: Schematic for a double-fold VECSEL cavity (F-cavity). Angles of $\alpha_1 = 10°$ and $\alpha_2 = 40°$ were used in the results from [37] discussed here.

is too limiting. On the other hand, a perfect AR coating and/or the use of a chip designed for ultrafast operation would yield significantly reduced microcavity field enhancement (and thus gain) and leave no resonance feature with which to test the principle of GDD self-compensation. An RPG design with a single-layer AR coating still possesses the expected GDD resonance feature and will still feature relatively high gain, but the extreme GDD peaks seen in Fig. 5.2 will be heavily damped to a

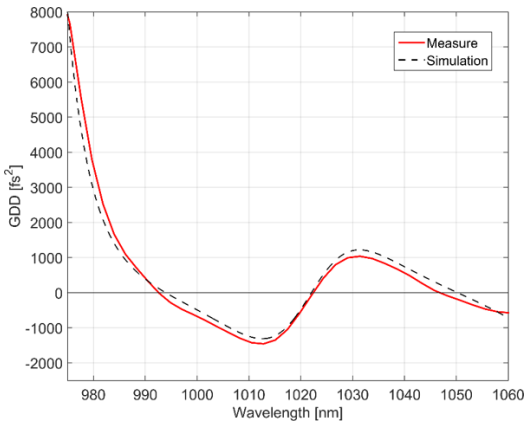

Figure 5.9: Measured GDD profile of the VECSEL device used in [37] plotted alongside simulation results for the same structure.

level where management via angle tuning is possible. Fig. 5.9 shows the measured GDD profile of the VECSEL chip used; the peak values of the GDD resonance are around ± 1000 fs^2, which is a factor of 10 lower than what is shown in Fig. 5.2.

This VECSEL, when used in a V-cavity, was never able to generate pulses below picosecond durations, despite careful optimization. When placed in an F-cavity configuration, however, it produced pulses as short as 410 fs. The autocorrelation and RF frequency data we recorded are shown in Fig. 5.10. The RF shows a narrow linewidth, limited by the resolution bandwidth of 1kHz, and equal power levels of the fundamental and second harmonic repetition rates of the cavity. Both are evidence of stable, single-pulse modelocked operation of the device. The 3.2 nm spectrum would have a transform limit of around 350 fs, meaning the 410 fs measured pulse duration is not too far off of transform limited.

As a brief aside from the discussion of the dispersion, the fact that the cavity used here is a multi-pass cavity has great benefits for the output power of the laser. More interactions with the gain yields better power scaling, while also allowing for the use of a longer cavity without encountering problems involving the nanosecond-scale lifetime of VECSELs' upperstate, which frustrates many attempts at making stable, MHz-regime pulsed VECSELs. Using a 3% output coupler, an average power of 1.14 W was attained at a repetition rate of 390 MHz. With a 410 fs pulse (and assuming a sech^2) pulse shape), that yields a peak power of 6.3 kW, which at the time of this

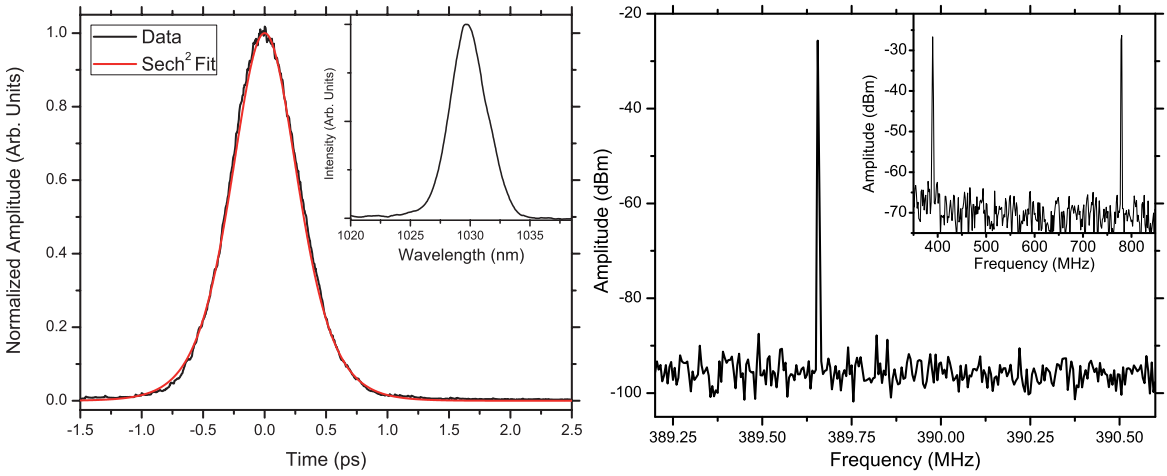

(a) (b)

Figure 5.10: (a) Autocorrelation trace of a 410 fs pulse with (inset) corresponding optical spectrum emitted from a VECSEL in F-cavity configuration. (b) RF spectrum measured at the same time as the autocorrelation trace, displaying narrow linewidth and (inset) equal amplitude between fundamental and first harmonic RF signatures. Both from [37].

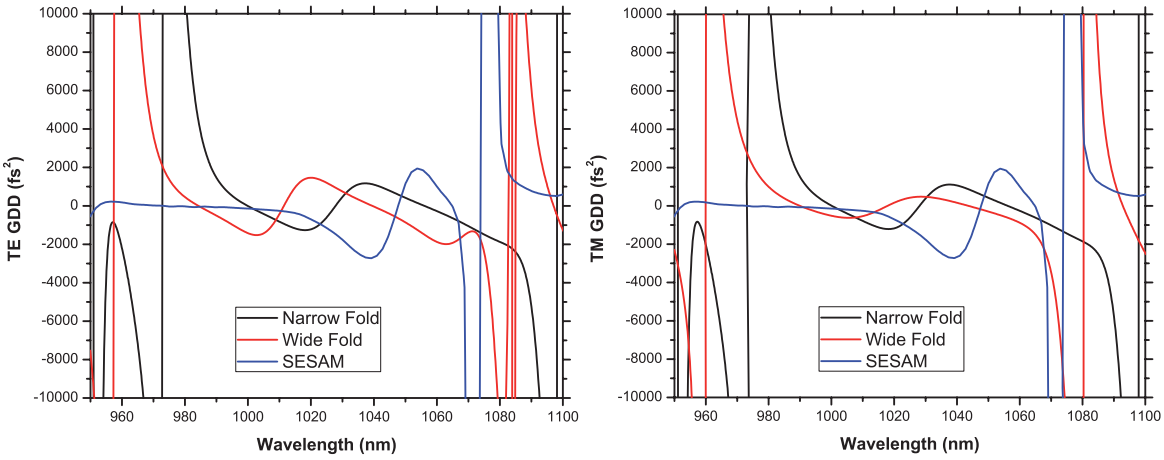

Figure 5.11: Simulated GDD profiles for the RPG device in Fig. 5.9 under different angles, with the GDD of the SESAM overlaid. (a)For the TE mode of the cavity and (b) for the TM mode of the cavity. From [37].

thesis is still the peak power record for a VECSEL device.

To fully understand how GDD compensation functioned to help produce this result, parameters of the simulated structure, which were carefully fitted to measured data at room temperature and 0° angle of incidence (as shown in Fig. 5.9), were adjusted to more closely resemble the VECSEL in operation (operational temperature and carrier densities; appropriate angles of incidence). The resulting simulated GDD is shown in Fig. 5.11 for the cavity's TE and TM modes. Of interest here is the difference in the wide-angle GDD profile between the cavity TE and TM modes; this will be exploited later experimentally as a confirmation test that it is indeed the dispersion-compensating aspect of the F-cavity which allows the formation of sub-ps pulses.

Additionally, the net GDD — accounting for 2 passes through the gain at narrow angle, 2 passes at wide angle, and one pass through the SESAM — is plotted for both cavity modes in Fig. 5.12. Of primary importance is that the net GDD for the cavity TE mode is relatively flat and close to 0. It is not completely ideal, but compared to the ± 1000 fs^2 peaks shown in Fig. 5.9, the values are well below 500 fs^2 and the slope is minimal. By contrast, the GDD for the TM mode is strongly negative within the lasing bandwidth and has a strong slope. It would be expected, then, that if the cavity could be forced to operate on its TE mode that much shorter pulses would be produced than if it were forced to operate on its TM mode.

This polarization enforcement is experimentally possible (and actually fairly easy, at that) by placing a thin glass etalon in the cavity at Brewster's Angle. When placed at Brewster's Angle for a given polarization, a Fabry-Perot etalon will have

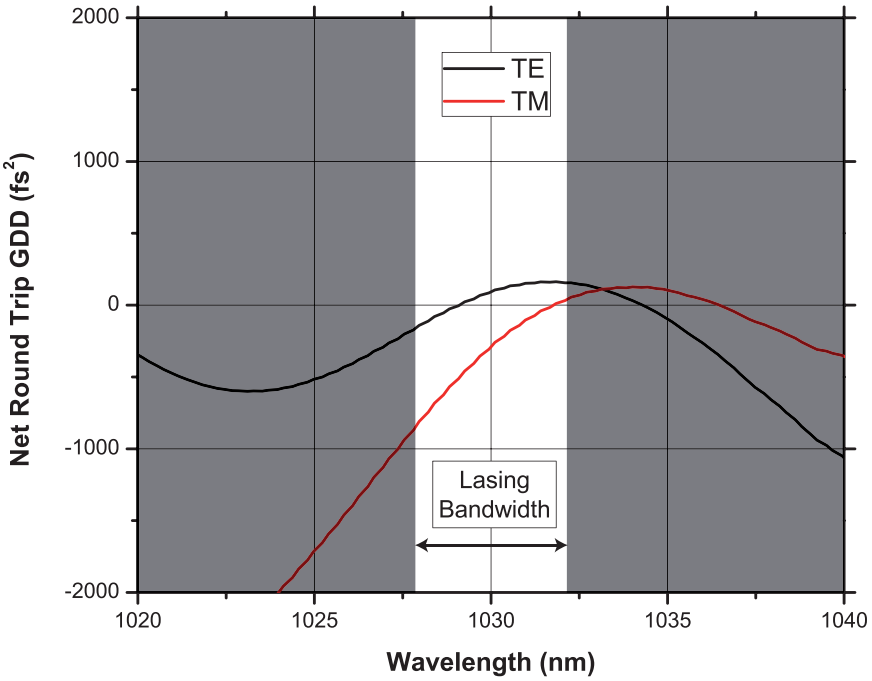

Figure 5.12: Simulated net GDD for cavity TE and TM modes calculated from the curves in Fig. 5.11. From [37].

minimal reflection at its surfaces, yielding negligible loss and spectral filtering for that polarization. The other polarization, however, will feel substantial loss and likely be prohibited from lasing. A thin glass piece is therefore a very nice polarization enforcer as, when great care is taken to ensure it is precisely at Brewster's Angle, it can prohibit an undesired polarization from lasing for only the cost of its added material dispersion (which for a 100 μm thick glass piece is a meager 2 fs^2).

A glass piece was used to enforce the cavity's TE and TM modes for the F-cavity discussed here and when the TE mode was enforced, corresponding to the well-behaved net GDD profile in Fig. 5.12, the 410 fs result presented above was obtained. When the TM mode was enforced, however, the VECSEL emitted pulses as long as 16 ps, with all other factors held constant except the polarization which was enforced. This is strong evidence that the precise blue-shifting of the VECSEL's GDD profile was a dominant factor in determining the durations of the resulting pulses emitted by the VECSEL.

The astute observer will note from Fig. 5.11 that this dispersion compensation scheme is still in some ways simple element compensation as both passes through the gain have positive GDD values and thus they serve to balance the negative GDD from the SESAM. That said, it is also clear that the *slopes* of the narrow- and wide-angle GDD profiles are strongly opposed, indicating that while second-order dispersion was controlled via compensation between the VECSEL and SESAM, third

and higher-order dispersion was indeed partially self-compensated in the VECSEL by the multiple angles of interaction.

Overall, the F-cavity design is unlikely to be used to generate truly ultrashort (100 fs and below) pulses. It is, however, a very nice design for balancing short pulse durations with high peak power yield by allowing for dispersion management without completely sacrificing field enhancements from microcavity effects (and also having a multi-pass design, which allows for longer VECSEL cavities). This is nicely displayed by the current record for VECSEL peak powers being from an F-cavity design.

5.4 Chapter Summary

This chapter has covered the requirement for, and the different techniques to attain, well-managed dispersion profiles for ultrafast VECSELs. After a brief motivation from the theory presented in Chapter 3, the broad categories of techniques we employ to manage dispersion, including element compensation, anti-reflection coatings, and angle tuning, have been presented.

Dispersion management has been a recurring challenge for us, with each new advance in other parts of ultrafast VECSELs yielding new requirements for precision in dispersion management. By example, if one is interested in picosecond pulses, dispersion aberrations even on the order of a few hundred fs^2 are not very worrisome, but if one desires pulses below 500 fs, GDD features on the order of 50-100 fs^2 become problematic and more precise dispersion control is required. As advancements are made in active region designs, understanding of modelocking techniques, and other aspects of ultrafast VECSEL design, it is likely that new dispersion management challenges will continue to present themselves.

Chapter 6

Gain Design and Non-Equilibrium Dynamics

In this chapter we will discuss the design of the gain structure. Much of the gain structure design is done in the theoretical arm of our group with a complete many-body quantum treatment (discussed in Section 3.5.2) and the significant experimental contributions here are in experimentally measuring the gain dynamics in-situ to provide verification for this theoretical effort. Nevertheless, I think it is important, in order to present a full picture of ultrafast VECSEL development, to discuss the ways in which active region design has progressed. I will be doing this mostly from the perspective of the linear gain, which is a useful perspective phenomenologically, but should be understood to arise from the full many body dynamics. It should be kept in mind while reading this, then, that much of the linear gain picture provided here is emergent from the behavior of the carrier distributions inside the gain, which will not be extensively described. For that complete quantum picture, please refer to the many papers our group has published on the topic, such as [51] and [52].

I have placed the chapter about gain last in the chapters describing modelocked VECSEL design because, chronologically during my Ph.D. progress, the design of the gain region was the last major category of challenge that we as a group undertook in the push to get closer to the 100 fs benchmark for modelocked VECSELs (although as I said in chapter 5, ever more precise dispersion management is a consistently recurring challenge). That said, this ordering is bit ironic considering that, if one is setting out to build an ultrafast VECSEL device, the gain is actually the first element that should be designed. Being cheaper and quicker to manufacture, the absorber, AR coatings, etc. should all ultimately follow from the result of the gain's design and growth.

The main improvement to active region design which has allowed VECSEL devices to approach the 100 fs benchmark can be seen broadly as the widening and flattening of supported gain bandwidths. This is assisted in part by practices discussed in Chapter 5, as reducing the effect of a VECSEL's microcavity will reduce the spectral filter it imposes, but equally important is the placement of the quantum wells inside the active region. Fig. 6.1 shows simulated gain profiles for a standard RPG device designed for high power emission and one of the stacked-quantum-well designs our group has developed for ultrafast emission. The contrast is stark between the sharp peak of the RPG device and the smoother, lower gain curve of the device intended for modelocked operation.

The core principle is how the placement of the wells affects the enhancement of

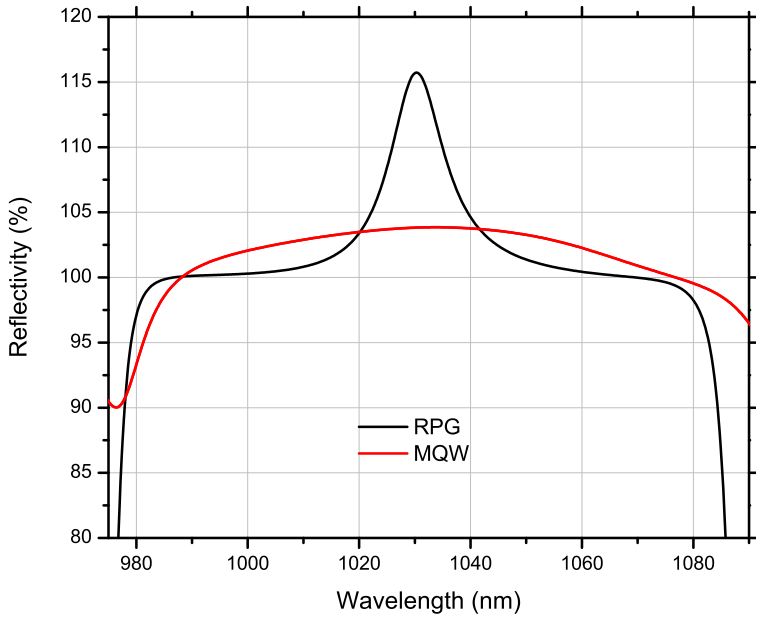

Figure 6.1: Comparison of simulated reflectivity curves (R > 100% indicates gain) for a high-gain RPG device and an ultrafast stacked quantum well device.

the interaction between the field and the quantum wells. An RPG structure, shown in Fig. 6.2a, maximizes the interaction for a specific wavelength by placing the wells where that field is strongest. By contrast, the stacked quantum well design, shown in Fig. 6.2b, still places the wells somewhat periodically, but now moves them off-center of the design wavelength's anti-nodes and stacks multiple wells per anti-node. The intention of the second structure is to produce the broader, flatter sort of gain bandwidth shown in Fig. 6.1 by spectrally diluting the field enhancement of the active region - that is, by allowing the active region to give low to moderate field enhancement to a broad bandwidth rather than strong enhancement to a narrow window around one line.

This is important for multiple reasons. Generally speaking, a strongly peaked gain is challenging to modelock with a broadband spectrum as it may prevent the occurence of a lasing state where a broadband pulse sees more gain than a narrowband one - if for every level of pump power, the narrowband state defeats the modelocking-enforcing loss mechanism just as easily or more easily than a broadband state, the laser is not incentivized to run in a broadband state. This is possible if the disparity is too high between the gain maximum and the wings of the desired lasing bandwidth.

Compounding this is the limited power scaling of a stable modelocked state. The spectral width of the lasing bandwidth tends to increase as pump powers are increased, which is often exploited (assuming the GDD is agreeable) to obtain slightly broader lasing spectra (and thus slightly shorter transform limits for the emitted

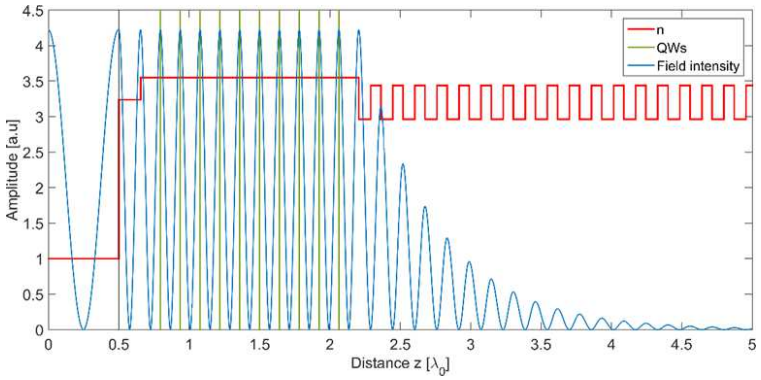

(a)

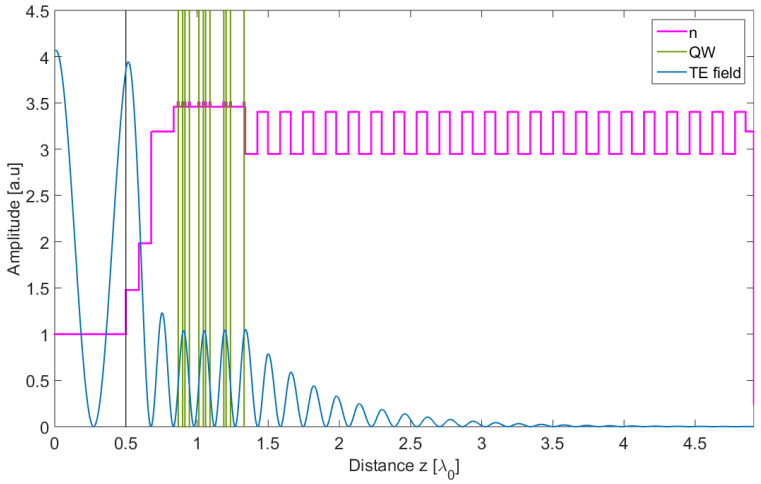

(b)

Figure 6.2: Design Schematics for (a) a high-gain RPG device and (b) an ultrafast stacked quantum well device. [9]

pulses). However, power scaling cannot be continued indefinitely as excessive intracavity power levels will become susceptible to instabilities in the modelocking state such as harmonic multi-pulsing, formation of parasitic side pulses, or the complete destruction of the modelocking state [56]. The usual solution when such instabilities arise is to reduce the pump power or increase the output coupling. Doing so will allow the return to a stable modelocking regime, but will negate the spectral broadening of higher pump levels. Thus, it is desirable to begin with a structure that supports a broad, flat gain bandwidth at reasonable pump levels.

And finally, having a broad, flat gain bandwidth reduces the ability of the intracavity field to saturate the gain. This is advantageous in general, but especially true in semiconductor media because of carrier dynamics unique to media with a band structure. Fig. 6.3 shows the timescales of different carrier dynamics that can

influence a VECSEL during lasing operation [52]. Of particular importance in this discussion is the fastest, intraband scattering, and it's ability to fill spectral holes burned in the gain's excited carrier population.

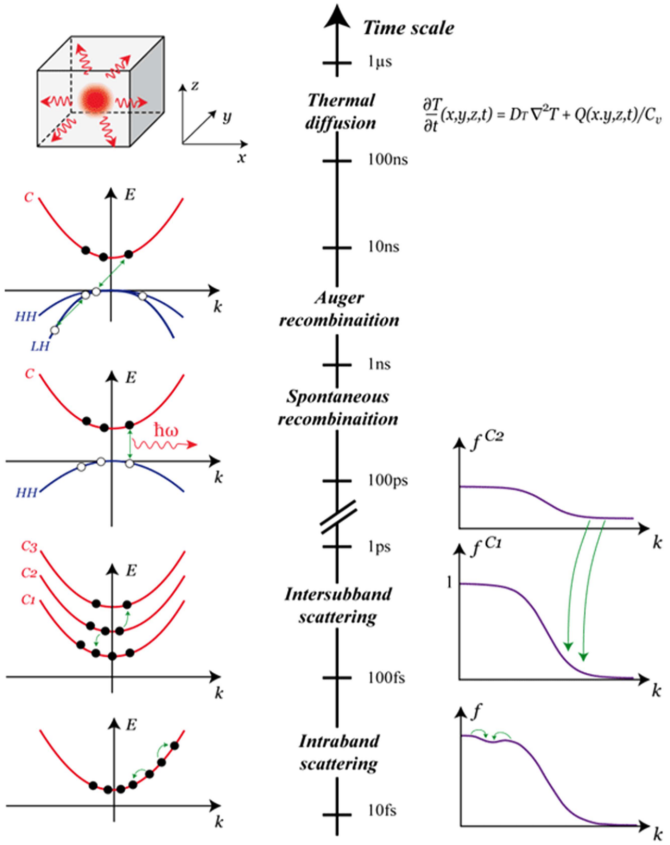

Figure 6.3: Time scales for various classes of carrier dynamics inside a semiconductor band structure. Along with schematic depictions of, on the left, the general process and, on the right, the corresponding change to the band's fermi distribution. From [52].

Because intraband scattering can kinetically fill holes burned in the carrier population on the order of 10 to 100 fs, it is possible that this will affect the formation of the modelocking state [51]. This is particularly true for a pulse in a device that supports a high gain, but a narrow spectrum. Such a pulse will burn a narrow, deep hole in the carrier population. The effect of kinetic hole filling is strong for this sort of deep, narrow hole when compared to the shallower, broader hole burned by a pulse of the same energy, but broader spectrum, as might be seen in a device that supports a lower gain but broader bandwidth [52].

It is because of these non-equilibrium dynamics that our group has turned to a full, many-body quantum simulation approach to optimizing the design of our

active regions. As with any theoretical model, it is important to connect the model to observable behavior in experiment. Thus there was an impetus to experimentally measure the dynamics of VECSEL gain media and to do so in a manner that comes as close as possible to the conditions during actual modelocked operation. Our response was to simply measure those conditions directly, leading to the idea of probing the gain of modelocked devices in-situ.

6.1 In-Situ Measurements of VECSEL Gain

The general idea of the in-situ measurements we performed is very similar to the ultrafast pump/probe setup discussed back in Section 4.1. There are two primary differences: first, instead of using a single ultrafast laser, clipping a small amount of power from its pulse train, and thus generating pump and probe beams derived from the same laser, here we attenuate the ultrafast laser's power and employ it solely as probe. The "pump" is taken to be the intracavity pulse of a lasing VECSEL in stable, modelocked operation.

The purpose of this change is clear: the "pump" pulse in the pump/probe measurement, generally, is the mechanism that prompts a response in the material being measured. Because we are using the intracavity pulse of an actual lasing, modelocked VECSEL as the "pump," whatever material response we measure (and thus, whatever carrier dynamics that measurement indicates) *must* be a result of the actual dynamics inside a working VECSEL, with no post hoc assumptions necessary about real carrier distributions, temperature levels, etc. to connect the measurement to the state of the gain during operation (as would be necessary with a traditional pump/probe measurement of the gain).

The second primary difference is enabled by the fact that we are now conducting this measurement with different lasers. Where in a standard pump/probe, the timing of the two pulses is controlled solely by the mechanical delay as both beams originate from the same laser, in our in-situ experiment we are utilizing two separate lasers with potentially different repetition rates. On one hand, that merits caution and care taken to control the timing jitter between the two pulse trains. On the other, it enables the use of an asynchronous scan — setting the repetition rates to be slightly different such that, shot-to-shot, the timing shifts slightly, eventually "scanning" the entire time domain as the timing difference slowly accumulates across hundreds or thousands of pulses.

This design is nice as it allows for very fine scan resolutions, easy tunability between scan resolution and scan window (by controlling the difference in repetition rates), and very fast data acquisition. We can take nanosecond-long scans with femtosecond resolution in milliseconds of real time; a comparable measurement with a mechanical delay line would take several minutes and be prone to beam pointing errors due to the required translation distances involved.

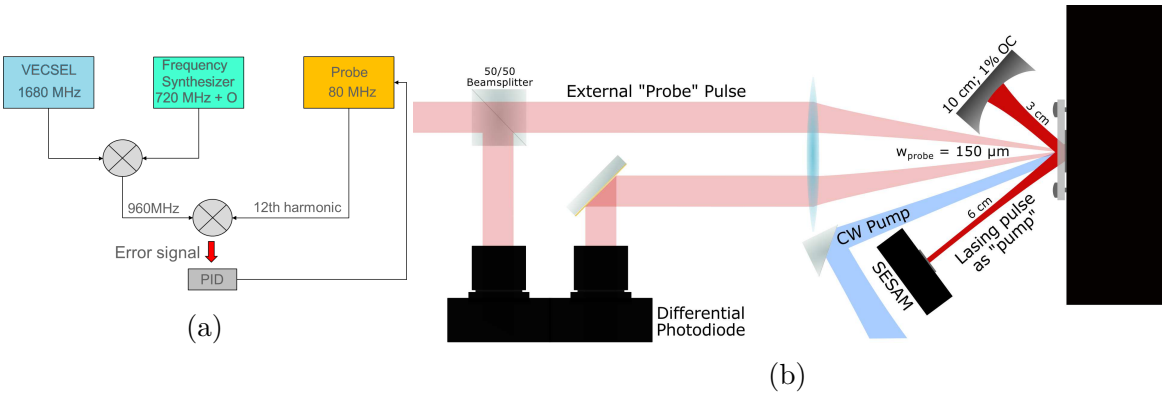

Figure 6.4: In-situ probing setup featuring (a)the electronic design and (b) the optical assembly. The waist given for the probe beam is on the surface of the gain.

Our specific setup is shown in Fig. 6.4. The same fiber laser described in Section 4.1 is used as probe (attenuated to < 10 mW of average power as in the probe arm before), with the addition of commercial locking electronics used to control the timing jitter (in this case, by ultimately locking the probe laser to the VECSEL) and provide an adjustable offset between the repetition rates of the two lasers (which gives rise to the asynchronous scan). Specifically, the VECSEL's 1.68 GHz repetition rate is mixed with a frequency synthesizer at 720 MHz plus the controllable offset mentioned above. The result is a signal at 960 MHz, which nicely mixes with the 12th harmonic of the probe laser's 80 MHz repetition rate (selected with a bandpass filter) to produce a DC error signal that can be used to lock the probe laser's repetition rate to the VECSEL, with a slight, controllable offset provided by the frequency synthesizer. The frequency offset can be set as low as Hz-range, giving scan resolution on the order of femtoseconds. Additionally, the probe laser has an ultimate pulse duration of 20 fs, leaving the limiting factor of our temporal resolution to be the remaining timing jitter between the two lasers, which we estimate to be around 100 fs.

Fig. 6.5 shows an example of what an in-situ scan of a modelocked VECSEL might look like. Fig. 6.5a shows the result of mixing the VECSEL repetition rate directly with the 21st harmonic of the probe laser's 80 MHz. This is incredibly important because it allows us to calibrate the time axis of the actual measurement, shown in Fig. 6.5b. One might suggest that the calibration could be attained by taking two points that indicate the round-trip time (which would be well known from the repetition rate of the laser), but such a method might fall prey to small perturbations in the pulse train timing, unlike a measurement calibrated with a complete error signal depicting the RF phase between the VECSEL and probe pulses.

Some important results from these pump/probe measurements are showcased in Fig. 6.6. In-situ probe measurements were conducted principally on two specific VECSELs — one a 10 well RPG structure that emitted 800 fs pulses at 1030 nm and

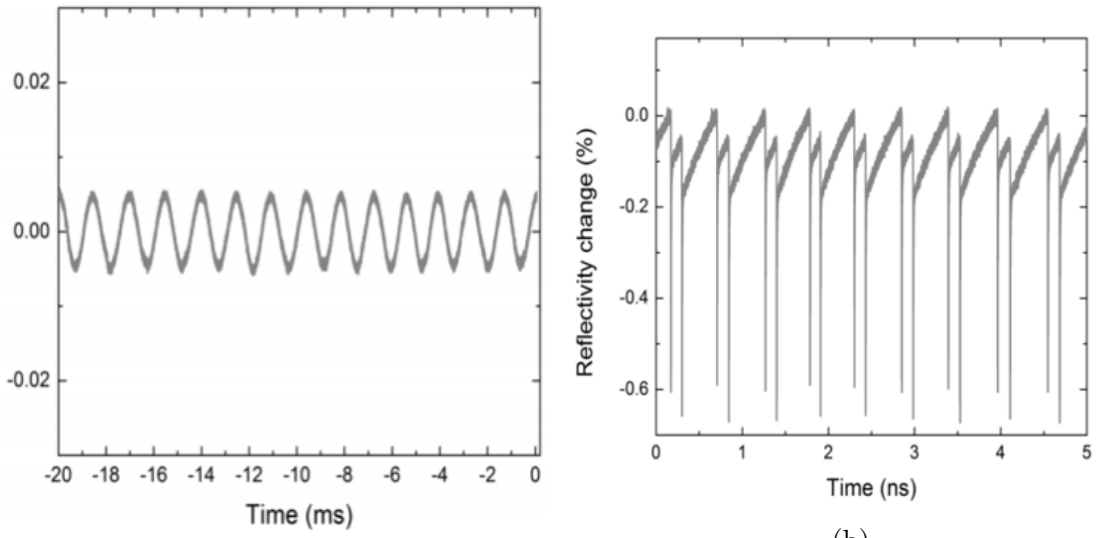

Figure 6.5: An example of an In-situ probe data of a VECSEL. (a) The error signal tracking the phase of the probe and the intracavity pulse; used to calibrate the time axis of the signal. (b) The scan with calibrated time axis. From [64].

the other an 8 well MQW structure that emitted 350 fs pulses at near 980 nm — which were modelocked in the same cavity design (shown in Fig. 6.4b). An example of a longer time window scan is presenting in Fig. 6.6a for the 800 fs device. The asymmetric signature here is indicative of the cavity geometry; the sharp dips are where the pump pulse interacts with the gain media, depleting the gain and thus the measured amplitude of the probe after it has passed through the gain structure. The asymmetry of these dips relative to a round trip of the cavity corresponds perfectly to a V-cavity design where the VECSEL is used as fold mirror (and thus incurs two dips for the two passes of the intracavity pulse through it) and the V-cavity's arms are not equal (a common design used to attain a mode size ratio of 1/3 to 1/2 between the gain and absorber). The important physics to note is that, if we define the higher-level dip as the first round-trip pass of the pulse, the gain has not recovered by the second pass of the pulse to the same level that it will recover to at the beginning of the next round-trip. This means that for each pass through the gain, the pulse will see a different gain level.

Fig. 6.6b is the most important plot. It shows a zoom-in of a single gain interaction for both the 800 fs and 350 fs VECSELs, overlayed with the results of a similar interaction in our many body simulation. The in-situ data is originally presented in [64]. The simulation fit was presented in [52]. The two important things to note are that, first, the many body calculations agree nicely with the experimental results

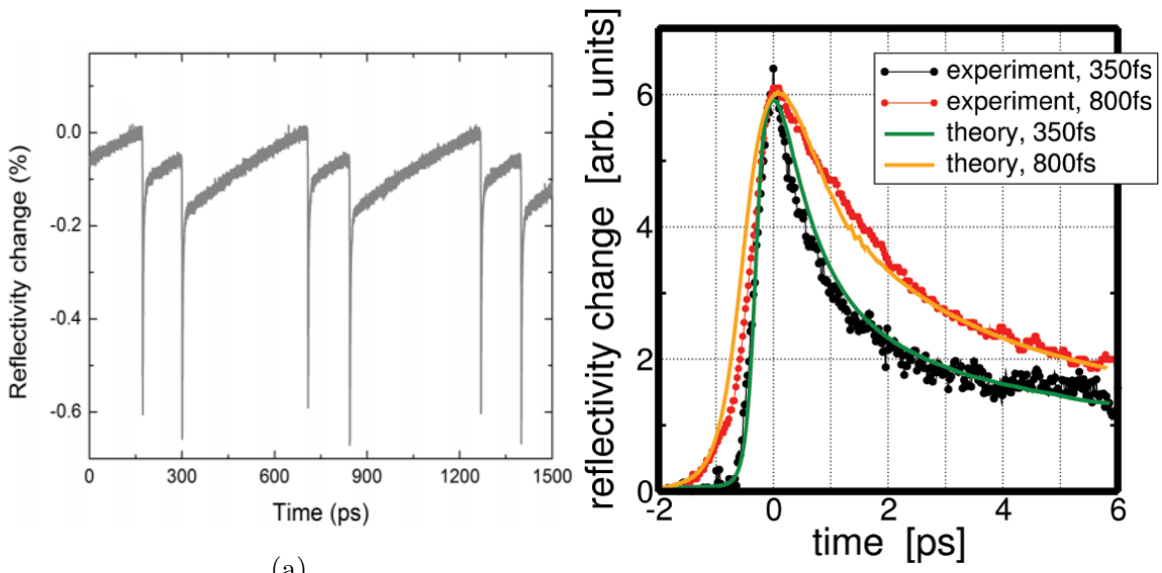

(a) (b)

Figure 6.6: In-situ scan of modelocked VECSELs. (a) A scan of the VECSEL emitting 800 fs pulses, depicting two full round-trips of the intracavity pulse. From [64]. (b) A zoom-in on a single interaction between a pulse and the gain, overlayed with simulation results, and depicted for two different VECSELs emitting different pulse durations. Plotted with a negative scaling factor; From [52].

which have been measured in-situ, and second, the experimental measurements depict gain recovery that visibly happens on at least two different time scales. The slow time scale is likely attributed to the CW pump. The faster portion of the recovery, while probably not directly showing kinetic hole filling due to intraband scattering — because the pulse lengths (350 fs and 800 fs, respectively) are long enough to screen any direct measurement of those effects in the 10 to 100 fs regime — *does* indicate that a significant portion of the gain recovers on the order of the duration of the pulses. This is further justification of the necessity of the many body calculations when considering VECSEL gain media.

6.2 Practical Limits on Active Region Design

From the above discussion, one may be tempted to try stacking as many quantum wells onto an anti-node as would be possible without allowing significant quantum tunneling between them. That would nicely fulfill all of the requirements for a broad, flat gain bandwidth discussed above and still maintain respectable levels of gain. And indeed, when simulated, structures with two or three repeats of several tightly-packed quantum wells near the intracavity field anti-nodes performed very well. Yet, our first

attempt at growing devices with such a design — four repeats of four stacked quantum wells, specifically — was quite disastrous. A picture of the photoluminescence emitted from the surface of one of those devices is shown in Fig. 6.7. The dark lines here are dislocations where the crystal lattice has actually fractured due to strain in the structure.

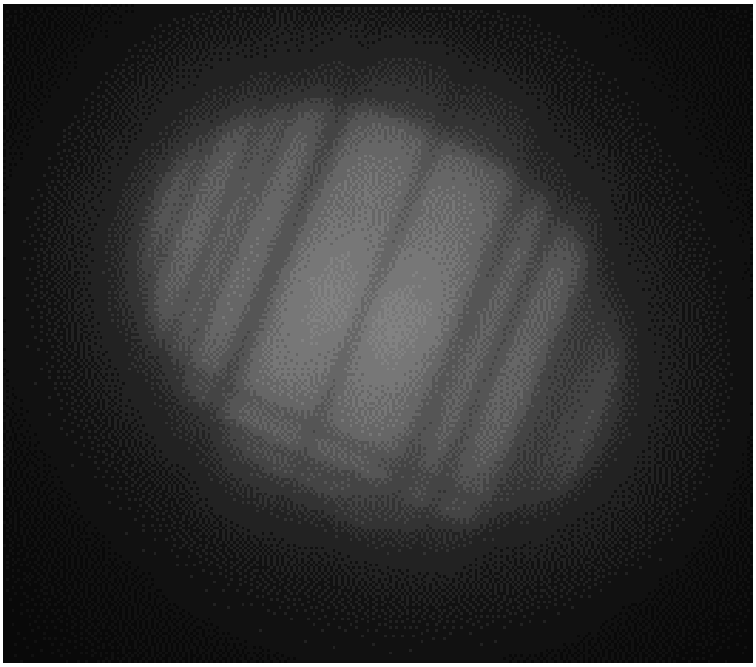

Figure 6.7: Surface photoluminescence measurement of a stacked quantum well VECSEL. The dark lines are indicative of dislocation defects.

The strain is created by the lattice mismatch between the barrier and well materials. Prior to the advent of stacking multiple wells in such close proximity, this strain wasn't that problematic, as the barriers were thick and designed to compensate for any strain introduced by the wells. When several wells are placed together, however, the strain in the structure accumulates and can cause the crystal lattice to fracture. This is problematic for multiple reasons, chiefly among them that the newly-formed dislocation provides a scattering site for carriers to escape confinement, significantly lowering both the gain of the structure and the carrier lifetime. It is important, then, that the strain be carefully considered when designing active regions for ultrafast VECSELs. One can see back in Fig. 6.2b that our design for a stacked well device, while it does still feature two anti-nodes with four wells, has a visibly increased spacing between the inner two and the outer two wells on each anti-node. This was a design choice heavily influenced by the need to relax the strain of the structure and prevent the crystal lattice from fracturing. To that end, it has been effective and our current generation of ultrafast VECSEL designs do not display any dislocation

defects in the crystal structure.

6.3 Chapter Summary

This chapter has explored the design strategy for the active region of VECSELs intended for modelocked operation in the ultrafast regime. The primary discussion is centered around the broadening and flattening of the gain bandwidth in order to support more broadband modelocking states.

A significant contribution on the experimental side is the measurement, in-situ, of the gain dynamics of a modelocked VECSEL. Probing the carrier dynamics of a device in-situ gives valuable insight to the behavior of the gain without necessitating any post hoc assumptions to adequately describe actual working behavior. The good agreement between experimental results and the results of our group's many body simulations is also important to note.

Finally, a very brief discussion brought up one of the potentially unforeseen practical limits to the design of these active regions, in the form of strain-induced dislocations.

Chapter 7

Multi-Wavelength and Spectral Broadening Applications

In this chapter, we will discuss some specific practical applications for VECSEL devices. As was mentioned in Chapter 2, the range of applications for which VECSEL devices are being investigated is impressively wide. Because of this, I have chosen to focus in this chapter on two applications which I have directly worked on myself: two-color emission for potential Difference Frequency Generation (DFG) of THz-regime frequencies and external amplification and spectral broadening for potential VECSEL-based frequency combs and spectroscopy.

7.1 Dual Wavelength VECSEL operation

One additional consequence of the unique carrier dynamics in VECSEL gain media discussed in previous chapters is that they lend themselves to the possible existence of multiple wavelengths within the cavity simultaneously [21,65–67]. This is unexpected for a homogeneously broadened gain medium and very interesting when considering difference frequency generation of THz frequencies, as a single source emitting on two wavelengths with a THz separation would alleviate the efficiency problems and additional mechanical effort that accompany attempting to beam combine separate lasers with a THz frequency spacing. This scheme has been demonstrated using VECSELs before [21, 68], but possessed significant stability challenges and (in the narrow linewidth case, limited power output).

There are two different perspectives from which Dr. Maik Scheller and I approached this idea of a dual-wavelength VECSEL for THz generation. Because DFG is a nonlinear process, high field intensities are desirable when attempting to employ a DFG scheme. This could be achievable by managing to operate the VECSEL in a state where it emits two pulses which are spectrally distinct (i.e. each pulse has its own spectral bandwidth) and overlapped in time such that they co-propagate (i.e. there would be no external delay required to overlap them temporally for DFG application) [40]. Another strategy would be to enforce especially stable dual wavelength CW emission such that utilizing intracavity fields for their increased intensity may be possible [41]. This section will cover the experimental setup for both of these modes of VECSEL operation.

7.1.1 Dual Wavelength Modelocked Operation

The VECSEL chip we used in this experiment was a 10 quantum well RPG structure optimized for emission at 1010 nm. It was placed in a standard V-cavity setup with a commercial LT-GaAs SESAM (nominally 0.5% modulation depth) to induce modelocking and a BK7 etalon to enforce 2-color operation. The setup is shown in Fig. 7.1.

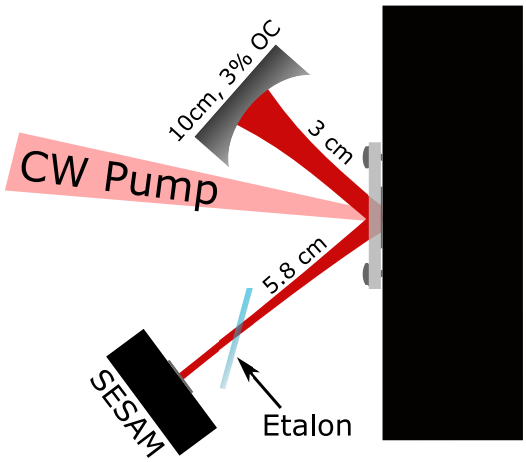

Figure 7.1: V-cavity used to induce 2-color modelocked operation in a VECSEL device. 100 μm and 180 μm thick etalons were used.

The characteristics of the etalon used to enforce two color operation are very important, as the free spectral range (FSR) of the etalon used is the determining factor for the wavelength separation of the two emission bandwidths. We employed, in separate trials, two etalons of 100 micron and 180 micron thicknesses, corresponding to FSRs of around 1 THz and 580 GHz, respectively, for the angle at which they were placed in the cavity. Because the effective thickness of an etalon changes with respect to the angle of incidence between it and the intracavity beam, the angle at which the etalon is placed in the cavity must be carefully controlled. The change in effective thickness can have a significant effect on the FSR and the supported wavelengths of the etalon. The 1 THz and 580 GHz values quoted above were for an etalon angle of around 60°. That angle was chosen because it corresponded to a selection of spectral peaks which, for our particular operational parameters, exhibited roughly equal power distribution between them.

The setup was particularly sensitive to both the pump level and the VECSEL's heat sink temperature, as both of these parameters will shift the microcavity filter function of the VECSEL, affecting the lasing wavelengths. We held the heat sink temperature and pump level at 16° C and 12 W, respectively. This yielded an output power of 820 mW.

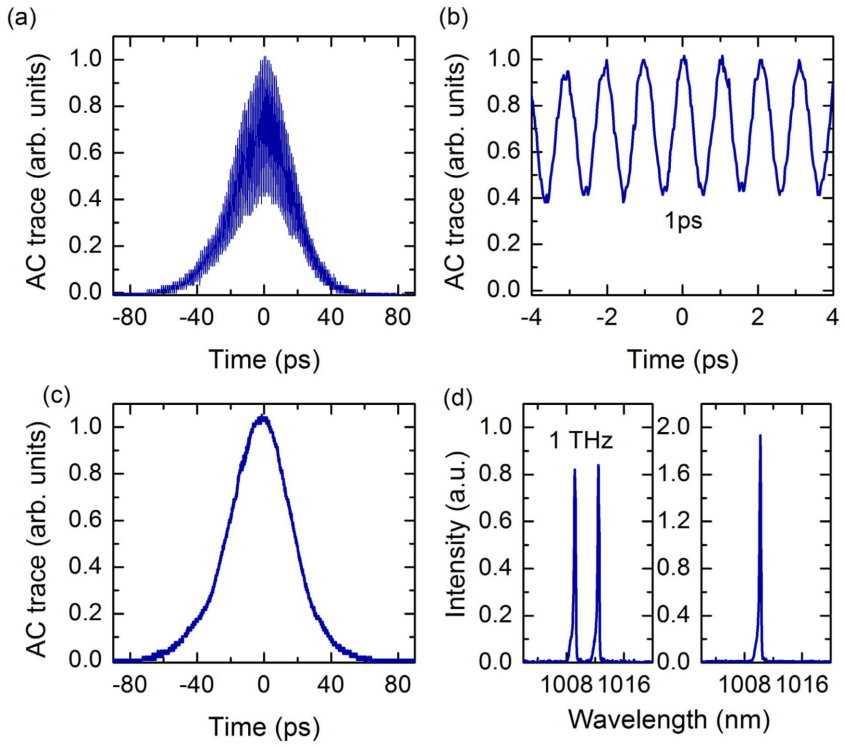

Figure 7.2: (a)Autocorrelation trace of a VECSEL in 2-color modelocked emission. (b)a zoom view to better display the beating signature. (c) Autocorrelation trace from the same VECSEL under single-color operation. (d) Optical spectra for (a) and (c), respectively. From [40].

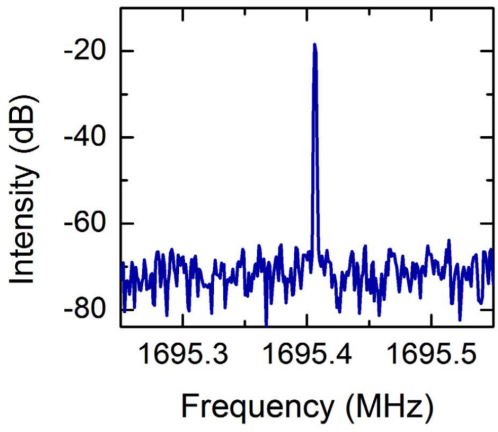

Figure 7.3: RF trace of VECSEL in 2-color modelocked operation. The width is limited by the resolution bandwidth of 10kHz. From [40].

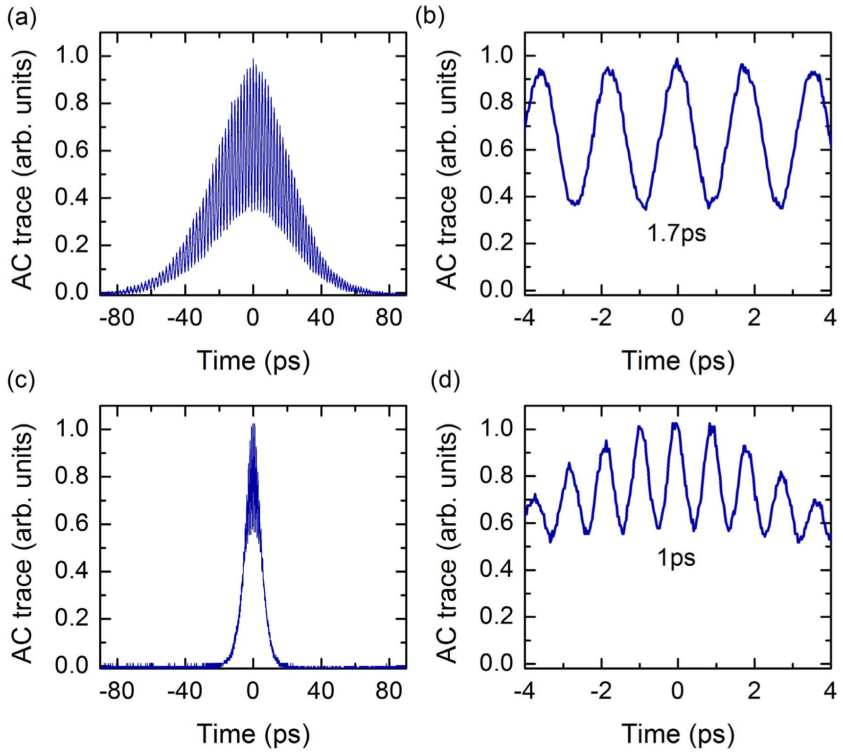

Figure 7.4: (a)Autocorrelation trace for 2-color modelocking enforced by a 180 micron thick etalon. (c) Comparison with similar state enforced by a 100 micron thick etalon. (b);(d) zoom views for (a) and (c), respectively. From [40].

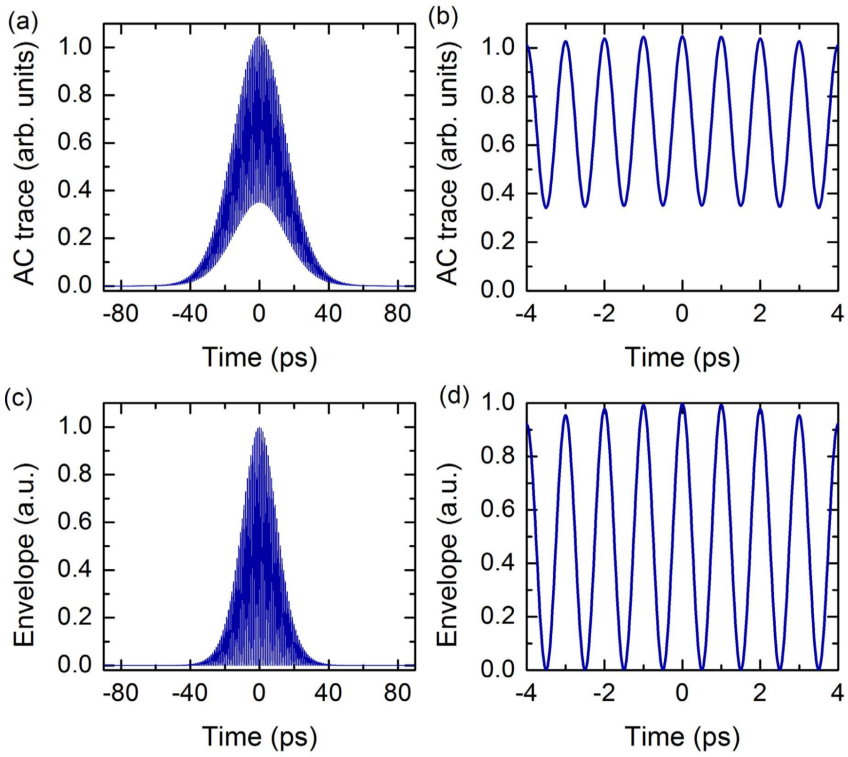

Figure 7.5: Calculated results featuring (a) the autocorrelation and (c) the field envelope for two pulses of sech^2 shape which are spaced in frequency by approx. 1 THz and overlapped in time. From [40].

Fig. 7.2 shows an autocorrelation trace of emission with the 100 micron etalon in the cavity and compares to an autocorrelation without the etalon present to enforce dual wavelength emission. The corresponding optical spectra are shown in Fig. 7.2d. Both dual-wavelength and single wavelength pulses display a FWHM duration of around 40 ps. However, there is a marked interference signature on the measured pulse envelope in the dual-wavelength case. The beating here occurs with a periodicity of 1 THz, which agrees nicely with the frequency separation of the two lasing bandwidths. Further, there are no side pulses found in the autocorrelation and Fig. 7.3 shows an RF trace with a 10 kHz linewidth, demonstrating stable mode locking at the fundamental repetition rate of the cavity.

A comparison of the pulses emitted by this setup with both etalons is shown in Fig. 7.4. When the 180 micron thick etalon was present, the pulse durations were significantly longer, with a FWHM of roughly 50 ps. The period of the beating here is about 1.7 ps, which corresponds to a frequency of 580 GHz.

To better understand what is occurring here, we use simple field calculations to simulate two pulses with perfect sech^2 pulse shapes and spectral bandwidths around 1010 nm separated by 1 THz. Then, we compute the autocorrelation of these pulses. The simulated results are shown in Fig. 7.5 for pulses that, individually, each have a FWHM of 24 ps. When the two pulses are assumed to be overlapped in time, the resulting autocorrelation results in a 40 ps FWHM and a beating signature with a period of 1 ps, which is in excellent agreement with the experimentally measured autocorrelation in 7.2. These results were all reported in [40].

7.1.2 Stable CW Dual Wavelength Operation in an F-Cavity

The idea to use the F-cavity described in Chapter 5 to attain dual-wavelength operation from a VECSEL arose during the proof of concept for its dispersion managing properties. Recall from Chapter 5 that a glass etalon was employed at Brewster's Angle as a polarization filter to enforce cavity TE and TM modes. While conducting that particular portion of the experiment, the VECSEL so strongly preferred 2-color operation that the only time we obtained a broadband modelocking spectrum instead of a two-color spectrum was when the glass piece was exactly at Brewster's Angle and imposing negligible spectral filtering on the intracavity field.

Upon further investigation, we discovered that the F-cavity design supports multi-wavelength modes of operation which are far more stable than any multi-wavelength emission in simpler cavity designs. The setup we used to demonstrate this is practically identical to the setup discussed in section 5.3.1 with the slight modifications being that we replaced the SESAM with an HR mirror, adjusted the angles slightly to 8° and 45° and placed the etalon not at Brewster's Angle. In addition, we constructed a single-pass V-cavity for comparison. That V-cavity, shown next to the F-cavity in Fig. 7.6, was selected as it allowed us to have both cavities aligned simultaneously to the same pump spot and select which cavity was operational simply by placing a

beam block in the other.

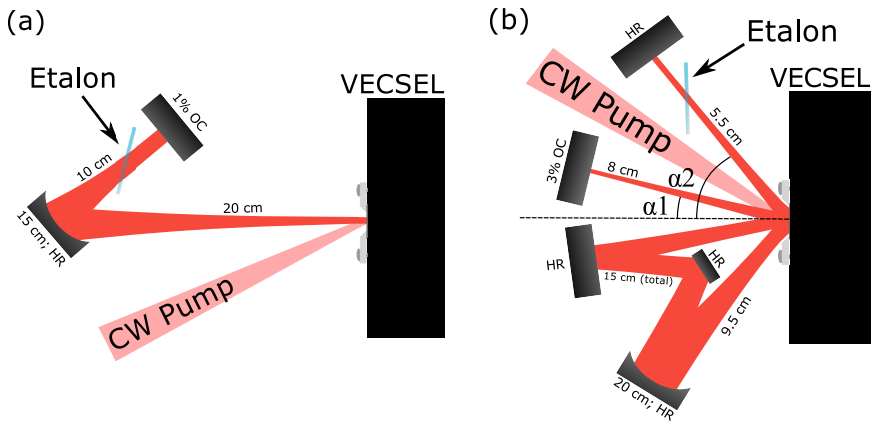

Figure 7.6: (a) V-cavity and (b) F-cavity used to induce 2-color CW operation in a VECSEL. The F-cavity angles were 8° and 45° .

To extract power from the F-cavity effectively, we employed a 3% output coupler and to keep the intracavity field intensities comparable, we employed a 1% output coupler in the V-cavity despite it being operational with higher output coupling. This is an important note because the F-cavity features multiple passes through the gain per round trip while the V-cavity features one; if the same output coupling were used for both cavities, the F-cavity would have a significantly higher intracavity field. This would confuse the investigation into the properties of the double-angle design as higher intracavity field is linked to more stable multi-wavelength operation inherently [66], rendering the comparison V-cavity meaningless. With comparable

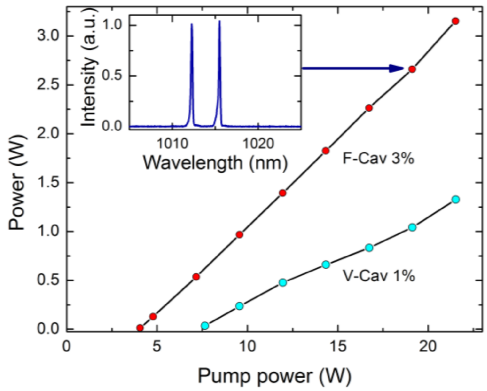

Figure 7.7: Power curves for F-cavity (red) and V-cavity (cyan) under 2-color emission. Spectrum displayed for F-cavity output power of 2.4 W, where the two wavelengths shared approximately equal spectral intensity distribution. From [41].

intracavity fields, though, the comparison of the geometrical differences is still valid.

Fig. 7.7 shows power curves for both cavities. Although the output power of the F-cavity exceeded 3 W while maintaining stable two-color operation, we choose to focus on the 2.6 W level as the spectral distribution of the power was roughly equal at this point. We employ the same 100 micron thick etalon from section 7.1.1 to enforce the two color operation and thus the emission wavelengths are separated by the same 1 THz frequency difference as was measured in the two-color modelocking experiments. It should be noted that for the power levels shown in 7.7 and the output coupling we employed, neither cavity exceeded 100 W of intracavity power, which is much lower than the range of relatively stable two-color operation seen in [66].

In order to assess the stability of this multi-color operation, we employ a grating to separate the two wavelengths and focus each onto separate fast photodiodes. This

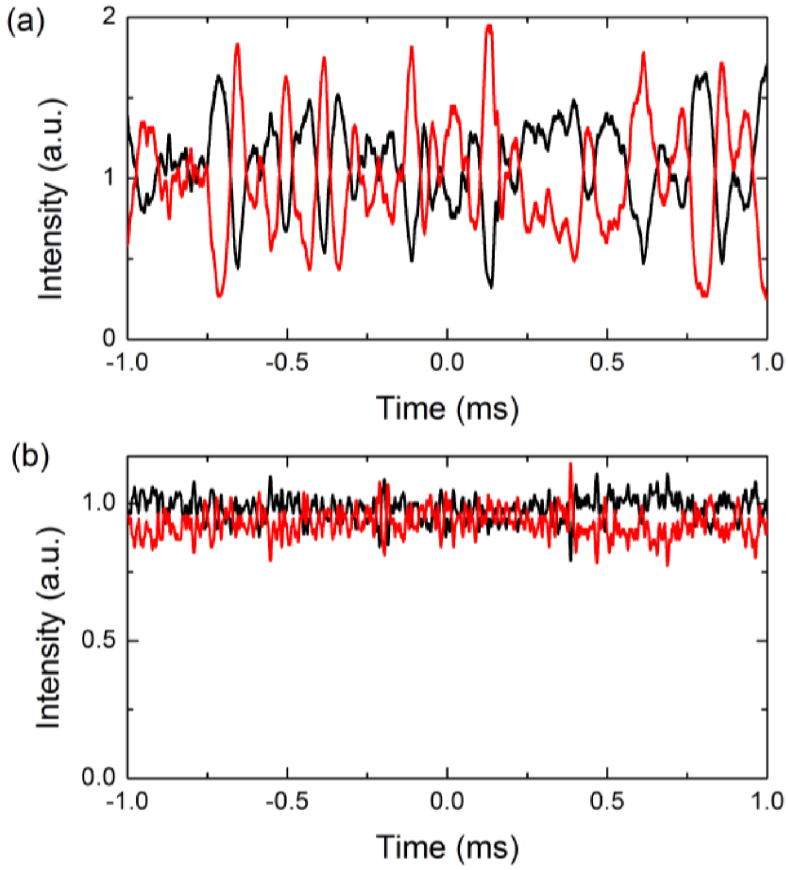

Figure 7.8: Intensity measurements for (a) V-cavity and (b) F-cavity using a grating and two fast photodiodes. Gain competition can be seen from the anti-correlated noise fluctuations between the two curves, which represent the intensity of each spectral line. From [41].

is necessary as the integration time of the standard spectrometers we had available exceeded 100 μs, leaving the possibility that rapid mode-hopping could appear as "stable" multi-wavelength operation. By using two photodiodes each with a bandwidth exceeding 100 MHz, fast fluctuations on the order of ns can be detected in both wavelength components simultaneously.

Fig. 7.8 shows relative intensity measurements, conducted with the two-photodiode measurement setup described above, for the two different cavities shown in Fig. 7.6. For the V-cavity, a strong anti-correlated noise is seen between the two emitted wavelengths, with the intensity of the noise fluctuations being almost on the order of the average intensity signal itself. By contrast, while the F-cavity demonstrates similar anti-correlated noise between the two emitted wavelengths, the fluctuations are no more than 10-15% of the average intensity level.

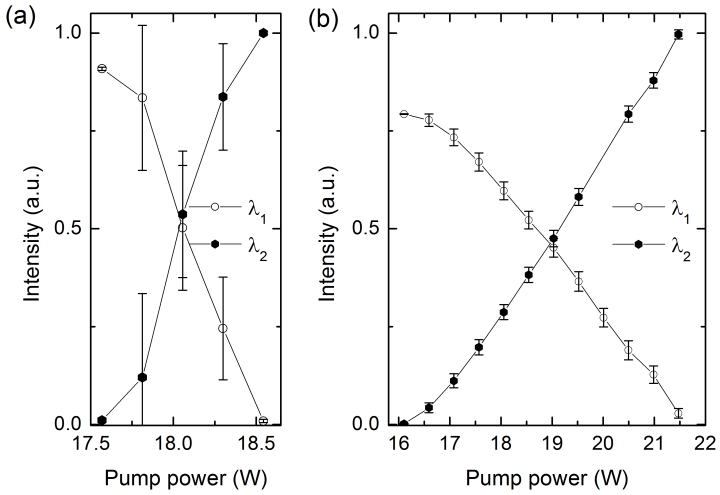

Figure 7.9: Comparison of 2-color stability in (a) V-cavity and (b) F-cavity relative to the level of pump power. Errors may be understood as the magnitude of the anti-correlated noise fluctuations seen in Fig. 7.8.

Another metric by which to test the stability of this two-color state is shown in Fig. 7.9. Here, the range of stable 2-color emission relative to the pump power level is plotted. The error bars a the amplitude of the anti-correlated intensity noise shown in Fig. 7.8. The V-cavity displays a relatively narrow range of pump powers over which two-color emission exists at all and the intensity noise is significant at every point. The F-cavity, on the other hand, displays a smooth transition of relative intensities across a range of pump powers larger than the stable range in the V-cavity by a factor of 5. The noise is also low at each point. This shift between which wavelength has the higher spectral intensity is expected, as the increased pump power levels will red-shift the gain bandwidth of the VECSEL relative to the supported Fabry-Perot

modes of the etalon enforcing two-color operation. An interesting note from the much wider range of pump powers capable of producing stable two-color emission in the F-cavity is that this likely means the F-cavity is also more robust against pump power or thermal fluctuations that may occur during operation.

The fundamental important difference, relative to multi-wavelength emission, between the V- and F-cavities shown here is the spatial distribution of the intracavity field on the gain structure. For the V-cavity, the field distribution is simply an expected Gaussian mode. By contrast, for the F-cavity, which features what is essentially 4 Gaussian beams interfering at 2 different angles, the spatial distribution is far more jagged. We calculated this interference pattern and a cross section is shown in Fig. 7.10. As can be seen, the transverse location of the intensity spikes in the resulting field pattern is wavelength dependent; this gives rise to the interesting phenomenon that, in the F-cavity, the intracavity fields for the two emission wavelengths are partially decoupled spatially inside the gain. A measure for the partial overlap of these fields is calculated as the mode overlap factor:

$$\frac{\int_{-\infty}^{\infty} I_{\lambda 1}(x) * I_{\lambda 2}(x) dx}{\sqrt{\int_{-\infty}^{\infty} I_{\lambda 1}(x) * I_{\lambda 1}(x) dx * \int_{-\infty}^{\infty} I_{\lambda 2}(x) * I_{\lambda 2}(x) dx}}, \tag{7.1}$$

and plotted in Fig. 7.10b against the wavelength spacing of the two emission wavelengths. The V-cavity features a mode overlap factor of 1, as both wavelengths

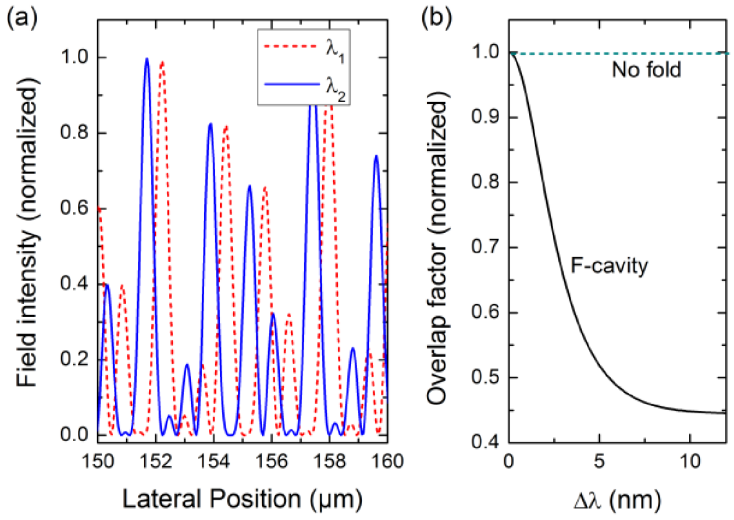

Figure 7.10: (a) Cross section of intracavity field interference pattern at the gain inside an F-cavity. (b)Overlap factor for the intracavity fields of two wavelengths in an F-cavity (solid) and a no-fold V-cavity (dashed), plotted against the wavelength separation. From [41].

are expected to have the same spatial distribution. The F-cavity, however, displays a significant drop-off in overlap at a wavelength separation of more than about 2 nm.

This partial decoupling of the two emission wavelengths' intracavity fields inside the gain is likely responsible for the increased stability of multi-wavelength emission inside the F-cavity. Because these modes can only burn a spatial hole in the gain where their fields strongly exist, they can only compete for gain where their fields overlap spatially. Their fields are not completely decoupled spatially, thus some gain competition does still exist and produces the minor anti-correlated intensity noise between the two emission wavelengths seen in the F-cavity measurements, but their partial decoupling allows for far less gain competition than in the case of the V-cavity, where the spatial distribution of the two modes' intracavity fields have strong overlap. These results are reported in [41].

In some ways, this result is more promising than the two-color modelocking result, as the challenge in attaining stable modelocked operation with two spectral bands of relatively equal power was great. Such modelocking states were reproducible, but incurred significant effort each time to arrive at the proper temperature, pump level, etalon angle, etc. parameters to ensure the stable, temporally overlapped, dual wavelength modelocking that is depicted in the previous section. By contrast, the stability of two-color emission in the F-cavity was easily reproducible, with similar results attainable within minutes, even when briefly tested in F-cavities that were not explicitly designed to minimize the overlap factor discussed above.

Using such multi-wavelength VECSELs to generate THz signals is still an ongoing investigation, but the results reported here, combined with the already established use of VECSELs for DFG of THz signals (e.g. [21]) shows great promise for VECSELs as a future THz source.

7.2 Power Amplification and Spectral Broadening of an Ultrafast VECSEL SEED

As ultrafast VECSELs become more well established near the 100 fs benchmark, they become increasingly attractive for certain frequency comb and spectroscopy applications. Fiber and Titanium:Sapphire modelocked systems are, generally, restricted to repetition rates of hundreds of MHz at the fastest. Because the spacing of the comb teeth in a frequency comb is dependent on the repetition rate of the seed laser, VECSELs offer a potential advantage if one desires specific teeth spacing resulting from rep rates of 1–10 GHz.

An important aspect of any frequency comb is a well-referenced carrier offset frequency. The offset frequency is the frequency difference one would measure if one extended the comb all the way to near-0 Hz and looked at the difference between the comb tooth closest to DC and 0 Hz. This allows for the establishment of the absolute position of the comb teeth, spectrally. Strategies for attaining this referenced

offset frequency generally include attaining a supercontinuum that spans a full octave and beating the lower frequency side of the octave with its corresponding second harmonic components in the higher frequency side (commonly referred to as an "f-to-2f" reference). As even the broadest VECSEL emission spectra are nowhere near an entire octave, additional systems must be used to enhance the bandwidth of the VECSEL seed. A common strategy is to use single-pass, rare-earth-dopant fiber amplifiers (Ytterbium at 1 μm and Erbium at 1.55 μm) to significantly amplify and partially broaden the VECSEL seed before coupling the amplified light into a highly nonlinear photonic crystal fiber (PCF) to exploit nonlinear broadening mechanics and attain a full octave [39, 69, 70].

7.2.1 Broadening Setup

In our lab, I (with Dr. Maik Scheller, initially) have investigated this using two different VECSEL seed lasers: an RPG device operating at 1020 nm and emitting approximately 500 fs pulses in an F-cavity at 390 MHz repetition rate and an MQW device operating at 1040 nm and emitting approximately 250 fs pulses in a V-cavity at 1.7 GHz repetition rate. Both devices were modelocking using surface quantum well SESAMs with a 5 nm cap layer and the MQW structure was also modelocked with a quantum dot SESAM with a 7 nm cap layer after the SQW SESAM was badly degraded (similar pulse durations resulted from both SESAMs). The output powers for these devices were roughly 150 mW for the F-cavity and roughly 70-80mW for the V-cavity when it was run in a stable single-pulse state.

Our amplifier setup was courtesy of a collaboration with the fiber laser group under Dr. R. Jason Jones and features a small, 2.5 m Ytterbium pre-amplifier and a larger 7 m Ytterbium power amplifier. The pre-amplifier is used to ensure proper seeding of the power amplifier to prevent any instabilities from low seeding; this was in some ways a precautionary measure and it bears mention that — particularly for the F-cavity device — it is possible the pre-amp is not strictly necessary and that the VECSEL itself could provide adequate seed power for the power amplifier. Both amplifiers featured angle-cleaved facets to prevent spectral instabilities from arising and the power amplifier possesses an additional safety mechanism at its output facet in the form of an angle-cleaved "endcap" — a 200-500 μm thick piece of bulk glass or multi-mode fiber spliced onto the end of the actual amplifier fiber to prevent coupling back into the amplifier fiber from the fresnel reflection at the glass/air interface of the output facet.

Between each stage of amplification, optical isolators are necessary to prevent instabilities from arising due to back-reflections from various optical components. A repetition rate mixing interlock system is used to protect the amplifiers by cutting power if seed levels drop dramatically (a full description of this interlock system is found in appendix section A.4.3). Two grating-pair compressors are also placed before and after the power amplifier. The first is used to pre-chirp the input pulses to the

power amplifier. This is done because one of the goals is to garner a certain level of spectral broadening in the power amplifier itself and because the fiber of the amplifier will possess positive dispersion properties that will stretch the pulse, it is advantageous to negatively chirp the pulse such that the amplifier actually compresses the pulse duration during initial amplification. Thus, the pulse duration will be shortest after some amplification and more capable of exciting any nonlinear broadening mechanisms in the amplifier fiber. The second grating pair is to control the chirp of the pulse as it enters the PCF. The grating separation for both pairs is determined phenomenologically by optimizing the spectral broadening of the respective fibers which they control the input to. The entire setup is shown in Fig. 7.11

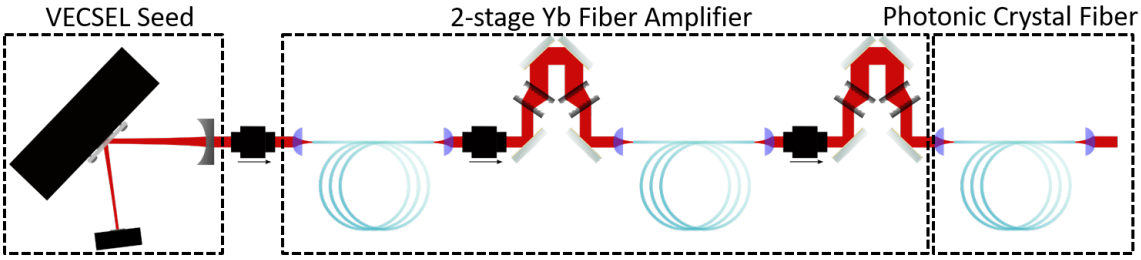

Figure 7.11: Schematic for an amplification and broadening scheme featuring a VECSEL seed, 2-stage Ytterbium amplifier, and highly nonlinear Photonic Crystal Fiber.

I have generated significant supercontinua with two different types of PCF: a 5-micron core fiber with a mode field diameter of 4.4 μm and dispersion 0-point at around 1040 nm (Thorlabs LMA-PM-5) and a 3.7-micron core fiber with a mode field diameter of 3.2 μm and a dispersion 0-point at 975 nm (NKT Photonics SC-3.7-975). The two most important parameters here are the mode field diameter, which can be used to calculate the fiber's nonlinear coefficient using:

$$\gamma = \frac{2 * \pi}{\lambda} \frac{n_2}{A_{eff}}, \tag{7.2}$$

and the 0-point of the dispersion. The 0-point of the dispersion dictates whether broadening happens in a positive or negative dispersion regime, which affects properties of the broadened spectra (see [71]). The nonlinear coefficient for the 1040 nm, 5-micron PCF was around $\gamma_{1040} = 9\frac{1}{W*km}$ and the coefficient of the 975 nm, 3.7-micron PCF was around $\gamma_{975} = 18\frac{1}{W*km}$.

A vital part of utilizing these PCFs was endcapping the input facet. For the result described below using the F-cavity VECSEL seed and the 1040 nm PCF, there was no endcap present on the input facet of the PCF and the input power was limited to around 1-2 W before thermal instability and damage to the input facet caused significant reductions in coupling efficiency (recoverable at lower power in the case of thermal instabilities, unrecoverable when damaged). All subsequent attempts

at generating supercontinua featured PCF with endcapped input facets, raising the thresholds for damage and thermal degradation of the coupling efficiency to as high as 9 W and allowing me to attain supercontinua with as high as 2-3 W of average power. The primary function of the endcap on the input facet is to reduce the intensity at the glass/air interface by first passing said interface while the beam is still relatively wide and replacing the glass/air interface with bulk glass at the point where the light is tightly focused to couple into the core of the PCF. By reducing the intensity in this manner, one can avoid damage thresholds at the glass/air interface, which would otherwise be present where the focused intensity of the input beam is the highest.

7.2.2 Supercontinuum Results

The performance of the amplifier setup was similar for both seed devices. The power after the pre-amp was approximately 500 mW with pulse durations similar to the output of each respective VECSEL. The power after the power amplifier was 17-18 W and the spectral bandwidth was broadened to exceed 50 nm. This was compressible in both cases to pulses of approximately 65 fs duration. A variable beam sampler was then used to reduce this beam to around 400-500 mW for initial coupling and later increase the sampled power to up to 9 W to utilize as much power for broadening inside the PCF as possible.

The first result came from the combination of the 500 fs, 390 MHz F-cavity combined with 2 meters of the 1040 nm, 5-micron PCF. The generated spectra is shown in Fig. 7.12. It spans roughly 800nm to almost 1500 nm, which is 100 nm short of

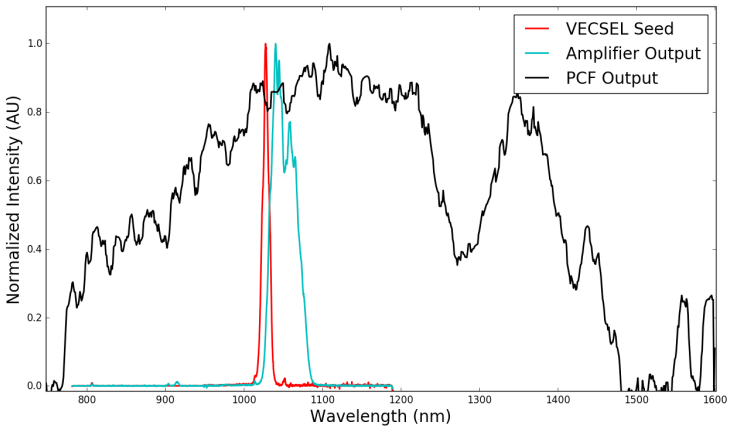

Figure 7.12: 300-400 mW Supercontinuum generated using a 390 MHz F-cavity emitting 500 fs pulses and a 2 m of 5 micron core, 1040 nm PCF. The two peaks near 1600 nm are likely a measurement artefact. The spectra of the VECSEL seed (red) and the high power fiber amplifier output (cyan) are also plotted.

an octave (the two peaks near 1600 nm are likely a measurement artefact discussed shortly as it features prominently in the last result in this section). The output, however, was limited to 300-400 mW as I had not developed a method for endcapping the input facet of the PCF when this measurement was taken, limiting the input to just over 1 W before the coupling degraded due to thermal effects and the threat of potential damage to the PCF arose.

The coupling mechanism had greatly improved by the time the next result was recorded, which is shown in Fig. 7.13. Here, the 1.7 GHz V-cavity seed is utilized with the 1040 nm, 5 micron core fiber. Endcapping procedures for the input facet of the PCF proved to be very important in preventing deleterious thermal effects and allowing me to increase the input power. Unfortunately, the increased repetition rate from 390 MHz to 1.7 GHz lowers the peak power of the pulses by a factor of around 4.5. And as broadening processes are expected to scale nonlinearly with the peak power of the pulse [71], this offsets any improvement to the power handling of the input coupling. Thus, despite the output supercontinua possessing an impressive average power of 3 W, it actually possesses lower spectral width than the previous result, spanning from 800 nm to almost 1400 nm. The input to the PCF in this case had roughly 8 W of average power.

Finally, a full octave was attained using the 1.7 GHz V-cavity with 1.5 meters of the 975 nm, 3.7 micron core fiber. From an input power of around 6W, an average output power of 2 W was attained with a spectrum, shown in Fig. 7.14, that spans 700 nm to 1400 nm. Care must be taken when reading this plot, however, as the signal at the longer side of the spectrum is partially misleading. We found that the OSA used in these measurements produced measurement artefacts by not effectively filtering 2nd-order diffractions from its gratings. Thus, it often displays 2nd-order

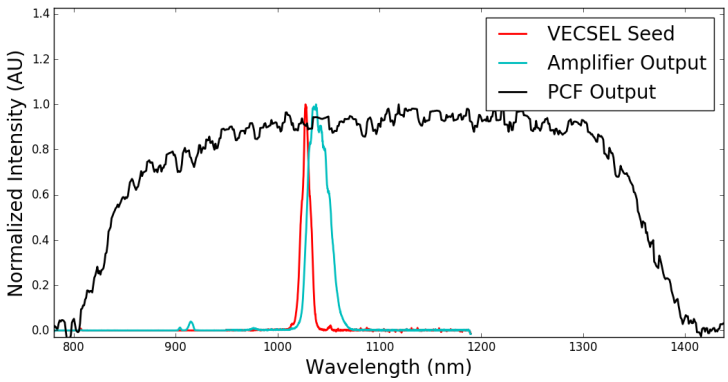

Figure 7.13: 3 W supercontinuum generated using a 1.7 GHz V-cavity emitting 250 fs pulses and 2 m of a 5 micron core, 1040 nm PCF. The spectra of the VECSEL seed (red) and the high power fiber amplifier output (cyan) are also plotted.

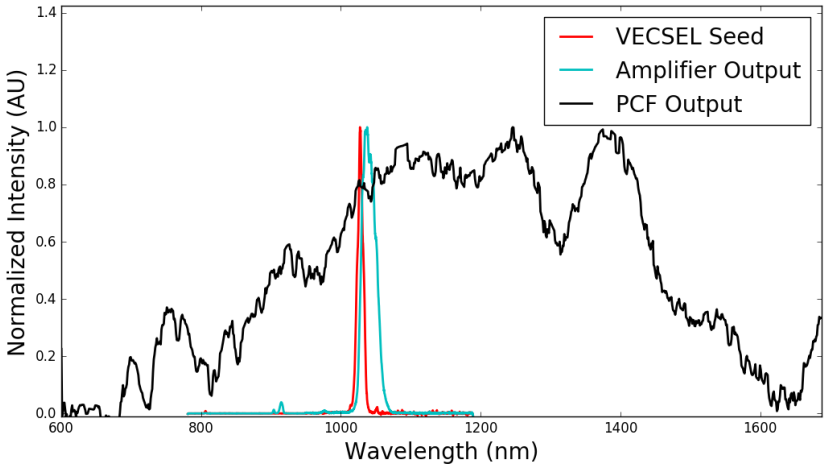

Figure 7.14: 2 W, octave-spanning supercontinuum generated using a 1.7 GHz V-cavity emitting 250 fs pulses and 1.5 m of a 3.7 micron core, 975 nm PCF. All signal to the right of the peak at 1400 nm is likely measurement artefact. The spectra of the VECSEL seed (red) and the high power fiber amplifier output (cyan) are also plotted.

diffractions of the short wavelength side of the spectrum as low-dB signal on the long side of the spectrum. Because of this, the signal in figure 7.14 at 1500-1700 nm is not real, but likely a recorded 2nd order diffraction from the two peaks between 700-800 nm (this same artefact likely features in Fig. 7.12 above). The strong peak at 1400 nm, however, is valid as displayed and thus a full octave was obtained at 2W of average power.

Work is still ongoing to detect the offset frequency using this supercontinuum, but initial results are promising. A beatnote has already been detected and is being characterized on the shorter side of the spectrum. The fact that the peak at 1400 nm, which must be frequency doubled via SHG to perform an f-to-2f reference measurement, is of such high spectral intensity also lends itself to the potential of this system to be used in an f-to-2f scheme to detect and potentially lock the carrier offset frequency.

7.3 Chapter Summary

In this chapter, I have discussed the results of investigations into two specific applications for VECSEL devices. One application featured both modelocked and highly stable CW multi-wavelength emission for potential THz generation, where Dr. Maik Scheller and I established the potential for modelocked VECSELs to emit two temporally overlapped pulses of different bandwidths spaced by about 1 THz in frequency

and we also demonstrated highly stable CW operation in F-cavity designs, vastly improving the robustness of multi-wavelength emission over comparison a comparison V-cavity. Further, I have shown initial attempts at using amplified VECSEL seed lasers to generate supercontinua suitable for potential frequency comb applications, resulting in a 3 W average power supercontinuum that spanned 800-1400 nm and a supercontinuum that spanned a full octave from 700-1400 nm with 2 W of average output power.

The next chapter shall conclude the main body of this thesis by restating some of the main points of each chapter and offering a brief glimpse into the future potential of VECSEL devices.

Chapter 8
Conclusion and Outlook

In this dissertation, we have shown that VECSELs are a highly promising member of the ultrafast class of lasers near 1 μm in wavelength. We have discussed the fundamentals of VECSEL design and modelocking operation, followed by an in-depth discussion of design considerations for ultrafast VECSELs and finally a brief look at two specific applications for VECSEL seed lasers. Here, the key experimental results presented in the dissertation will be summarized and a brief future outlook for VECSEL development will be presented.

8.1 Summary of Experimental Results

The key experimental results featured in this dissertation are the in-situ probing measurements which nicely support the results from the many-body model used to design and optimize the active regions of our devices, and the use of a double fold cavity design to attain record peak powers from modelocked VECSEL operation.

In Chapter 6, the in-situ probing measurement was conducted using an 80 MHz commercial fiber laser producing 120 fs pulses at around 1040nm, which were then spectrally broadened in single-mode fiber to a bandwidth exceeding 100nm and recompressed to 20 fs. By use of an asynchronous scan, the gain dynamics of two VECSEL samples (which emitted 350 fs and 800 fs pulses, respectively) were measured in-situ during modelocked operation with the intracavity, lasing pulse behaving as a strong "pump" and the attenuated fiber laser serving as a weak "probe." The demonstration of this technique itself is a nice result, as it features several advantages over traditional pump/probe measurements: the asynchronous scan confers the ability to measure nanosecond-long windows in milliseconds of real time without a mechanical delay line being required and the in-situ nature of the measurement means that no assumptions are necessary to connect it to a device during operation, as would be necessary with traditional pump/probe measurements. The results themselves give insight into the time dynamics of VECSEL gain during operation (particularly the level of fast recovery during one pass and the overall recovery per round trip) and are shown to agree well with the many-body theory that is used to design our devices.

The record peak power result in Chapter 5 was attained from the use of a double-folded cavity design that allowed for both partial dispersion self-compensation and a repetition rate in the hundreds of MHz without succumbing to multi-pulsing or other undesirable modes of pulsed operation. The multi-pass design of the cavity allowed for no single arm of the cavity to have a travel time on the order of the VECSEL

carrier lifetime, preventing multi-pulsing and allowing fundamental operation at a repetition rate of 390 MHz. The double fold at two different angles allows the blue shifted the dispersion curve of the wide fold to partially compensate for dispersion incurred during the pass at a narrow angle and produce 410 fs pulses. This 390 MHz repetition rate and 410 fs pulse duration combine with an average power of 1.14 W to yield a record-setting 6.3 kW of peak power.

In addition to these results, work was presented in Chapter 7 regarding two-color VECSEL operation for potential DFG of THz frequencies wherein both modelocked operation on two wavelengths simultaneously and highly stable CW two-color operation were demonstrated. In the modelocking case, Dr. Maik Scheller and I demonstrated a VECSEL emitting stable ps-scale pulses comprised of two individual pulses temporally overlapped and spaced spectrally by approximately 1 THz. We also utilized the F-cavity design mentioned above in CW operation with an etalon to enforce two-color operation and found that the partial spatial decoupling between the field distributions of each lasing wavelength produced a two-color emission state that had low gain competition and was far more robust than a comparable single-pass V-cavity design. Both of these results lend themselves to the possibility of using VECSEL devices as tunable sources for difference frequency generation of THz.

Some work has also been presented regarding some of the more specific details in ultrafast VECSEL design which, while unpublished, are nevertheless vital to the day-to-day design and engineering of our devices. This work includes power stability and pump/probe measurement of the relaxation rates of our SESAM devices, compared across different types of absorber design in Chapter 4; different strategies for managing dispersion, including element compensation and, with helpful input from Dr. Alexandre Laurain, precise anti-reflection coating design in Chapter 5; and the practical limitations on strain accumulation inside the gain structure which can cause dislocations and degraded device performance in Chapter 6.

Finally, preliminary results from an ongoing investigation into VECSEL-based frequency combs and spectroscopy — in collaboration with the group of Dr. R. Jason Jones — was presented in Chapter 7. These results include a modelocked VECSEL seed that is amplified in a two-stage Yb-doped fiber amplifier to almost 20 W of power with a spectral bandwidth that exceeds 50 nm and is compressible to approximately 65 fs in pulse duration. This amplified pulse train has been shown to be capable of generating supercontinua with 2-3 W of average power and, in one case, an octave-spanning spectral width.

8.2 Future Outlook

Recent work in the VECSEL community has been essential in establishing the VECSEL as a viable source both in high power CW and modelocked operation. With power levels being demonstrated as high as 100 W CW [3, 4] and pulsed VECSELs

having achieved, in separate instances, the 100 fs benchmark [12] and over 6 kW of peak power [37], much of the foundational proof of concept work has been completed surrounding the viability of VECSELs as a gain medium near 1 micron. Work going forward can now begin to heavily push applications for these VECSEL devices.

The results presented in this dissertation lend themselves directly to the potential for tunable, difference-frequency THz sources and VECSEL-based frequency combs. For the former, DFG has already been established as a viable strategy for intracavity generation of THz frequecies. The F-cavity design could stand to make improvements to the stability of the generated THz along with potentially higher power yield conferred by the four passes through the gain per round trip. If external stabilization mechanisms are employed and additional linewidth filters can be found which do not disrupt the THz frequency separation described in Chapter 7, it may even be conceivable that F-cavity designs could be used as a stable, tunable, single frequency THz source.

There is still much work to be done with regard to VECSEL-based frequency combs. The carrier offset frequency must still be detected by f-to-2f measurement or otherwise and possibly stabilized. This will require much more care be taken to reduce noise in the overall system than was taken to demonstrate a proof-of-concept octave-spanning supercontinuum. Better isolation from mechanical noise and addressing the noise contributions from systems like the diode pump current driver is essential. Further, because the repetition rate of the VECSEL seed is in the GHz regime, the requirements are more strict regarding the response speed of the locking electronics for offset frequency stabilization. It is also interesting to note that the supercontinua presented here meet all the spectral criteria for the biological imaging work detailed in [72], providing another potential spectroscopic application for VECSEL technology.

Appendix A

Practical Lab Guide

Standard academic writing demands a certain level of formality. And while this formality is extremely useful in pursuit of technical correctness and competency, it can often seem daunting and unapproachable to beginners. We also often do not discuss formally some of the more day-to-day practicalities of the laboratory environment, choosing instead to focus primarily on the results and technical details of our experiments to the exclusion of the initial procedures and tricks that helped us get those results. This is understandable, as the refinement of our initial attempts at an experiment into well-documented and characterized procedures is part of what makes scientific results valid and precise. I have found, however, that sometimes the full formality of academic presentation does not make for the most approachable learning tool.

My intention with this appendix is to include in this thesis some of the practical methods, self-built devices, and laboratory tricks that I have learned during my time in graduate school and to do so in a more informal manner. The presentation here is intended mostly for other graduate students earlier in their grad school tenure and I have tried to make it a little more conversational and approachable than the main chapters of the thesis. My hope is that this appendix can serve as a brief reference for certain things that may be hard to find in written form elsewhere and an easily read guide for any students unfamiliar with a laboratory environment.

A.1 Introduction

I have found that there are a lot of very practical things you should know in an experimental laser lab (or optics lab in general) which, despite being incredibly useful and integral to doing timely, effective work in the lab, the majority of scientists consider themselves too dignified to write down in any formal manner. I, however, have no qualms including such enlightening sections to an appendix in my thesis as "advanced mirror mount wiggling techniques."

What follows, then, is a slightly disorganized and haphazard collection of techniques, discussion and instructional vignettes which I hope will serve to assist you in some small way in becoming an effective experimentalist. There will be some practical, general techniques (e.g. laser alignment), some useful jury-rigged devices built with fairly common lab parts, and some long-winded rambling on general lab philosophy. None of what you are about to read should be treated as gospel and all of it is presented with the intention of giving you a basis for you to figure out your own lab tricks and pass them on to grad students after you.

Except the part about mirror mount wiggling. That is gospel and you should treat the technique with the reverence it deserves.

A.2 The Brief Flash of Hope

The spirits of an experimental graduate student live and die on the brief flash of hope. The systems that we are tasked with building are often, procedurally, quite simple to optimize once they produce something measurable, but contain an overwhelming amount of degrees of freedom to consider when the whole thing doesn't work at all. And to make matters worse, before something works in a measurable, optimize-able fashion, it's easy to select the wrong knob to turn and make everything worse for yourself.

The hallmark example of this is aligning a laser cavity. There are many, many variables to consider when getting a laser to work (tip and tilt on both mirrors, location of both mirrors in 3D space, temperature control in gain, pump modematching with cavity mode, pump level, etc.) and they are all mostly orthogonal. That means once the laser works, you can adjust each in turn to optimize the performance in a very straight-forward manner. However, before the laser begins working, you have no real indication from the laser itself which parameter you need to adjust and how you need to adjust it (and adjusting the wrong parameter can misalign the whole thing further!).

In these situations, your greatest ally is the extreme difference between human time scales and optical ones. Because time scales for light are so much faster than for human perception and movement, relatively chaotic motion in human terms can often lead to temporary stability in light terms. The result, if everything else is in working order, is a brief existence of signal, be that a flash on a detection card or a brief spike on a power meter - whichever signal you're expecting to detect. And seeing that brief signal is enough to immediately reduce all your adjustable parameters to a small handful of options that you can methodically work through - a brief flash of hope that maybe your experiment will actually function properly after all.

Again, the most easily understood example is probably laser alignment. You prepare your gain medium, set up your pump, mount your mirrors, and set about aligning everything. There is a completely algorithmic way for doing this and it is painful and time consuming. On the other hand, you can also place one mirror such that you're pretty sure it's sort of aligned and wiggle the other mirror around chaotically. If, in that chaotic movement, you temporarily pass through perfect alignment, the laser will briefly flash on because, again, optical time scales are so much faster than you wiggling the mirror around. And then you immediately know that your gain is functional and your pump is satisfactory and your alignment is close-ish, so you can settle down and work with the tip and tilt screws to permanently align the laser (see the next section for best practices in doing exactly this).

What I think you will find eventually is that while experimentalists love to sell their work as hard science (careful, methodical, etc.) and while that is true for the most part when considering the overall process (planning the experiment, taking publication data, reviewing publications, etc.), the bulk of the experimental construction feels a whole lot more like an art (intuitive, a little haphazard, etc.) than any of us want to admit. This is not so much of a problem, provided that you do at some point stand your work up against a fair standard of scientific rigor, but I think the main take away is to not over focus on acting like some scientist stereotype. Sometimes science is a 8 GB data set covering every possible parameter; sometimes science is going out on an intuitive limb looking for a brief flash of hope.

A.3 Practical Laboratory Techniques

In this section, we're covering some practical techniques for getting things done in the lab. They are generally more procedural or conceptual ideas than the stuff in the practical devices section and if you wind up having a problem with what I call a "technique" and what I call a "device," I don't really care. Why are you being so critical over an arbitrary distinction I made in an appendix to give it some sort of literary structure? Get back to work.

A.3.1 Laser Alignment: Planning the Cavity

This particular part is written at the suggestion of one of my favorite research professors, Dr. Alexandre Laurain, who lamented the inability of "kids these days" to handle the rudimentary skill of planning a laser cavity. The very basics of cavity stability are quite common place in laser textbooks, but I'll quote the stability equations for you and then discuss how to simulate a cavity mode with a relatively simple ABCD matrix approach and extract useful information for constructing a cavity.

The general rule of thumb for cavity stability in simple cavities is that the length of the cavity can't exceed the sum of focal lengths for all the curved mirror interactions that the cavity uses. So if you have one 10cm mirror that you hit once, 10cm is the max cavity length (slightly shorter to account for real life imperfections); if you have two 10cm mirrors that you hit once each, 20cm is the max length; if you have one 10cm that you hit twice, 20cm is the max length again; etc.

The more specific rule is that:

$$-1 \leq \frac{A+D}{2} \leq 1 \tag{A.1}$$

Where A and D are elements in the Ray Transfer Matrix for the entire cavity. Now, if you've never seen Ray Transfer Matrices (often called "ABCD matrices") before, you should really look them up because they're great. Ray Transfer analyses use 2x2 matrices to describe how each part of an optical system affects a ray of

light interacting with it. And the great part about them is that you can describe an entire system with a 2x2 matrix by simply matrix-multiplying its components' ABCD matrices (in order).

We're going to use this to design cavities because we know that cavities have to be self-consistent (i.e that a round-trip through the cavity must return a field that, in spacial distribution, is the same as what you started with). If you multiply ABCD matrices comprising a complete round trip through the cavity, demanding self-consistency allows you to glean important information from the resulting 2x2. Of specific note is your ability to determine the beam width (assuming your cavity supports well-behaved Gaussian beams) and radius of curvature, which allow you to then calculate how the beam evolves at every point inside the cavity.

Here are a list of matrices that will prove useful:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \quad ABCD\ General\ format \tag{A.2}$$

$$\begin{bmatrix} 1 & \frac{d}{n} \\ 0 & 1 \end{bmatrix} \quad Distance\ 'd'\ through\ medium\ with\ index\ 'n' \tag{A.3}$$

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad Flat\ Mirror \tag{A.4}$$

$$\begin{bmatrix} 1 & 0 \\ -\frac{2*cos(\theta)}{R} & 1 \end{bmatrix} \quad Curved\ Mirror\ of\ radius\ 'R'\ at\ angle\ '\theta'(saggital\ axis) \tag{A.5}$$

$$\begin{bmatrix} 1 & 0 \\ -\frac{2}{R*cos(\theta)} & 1 \end{bmatrix} \quad Curved\ Mirror\ of\ radius\ 'R'\ at\ angle\ '\theta'(tangential\ axis) \tag{A.6}$$

$$\begin{bmatrix} cos(\gamma(P)*t) & \frac{1}{n*\gamma(P)}*sin(\gamma(P)*t) \\ -n*\gamma(P)*sin(\gamma(P)*t) & cos(\gamma(P)*t) \end{bmatrix} \quad VECSEL \tag{A.7}$$

The complication in the VECSEL matrix is intended to account for thermal effects in the structure and it need only be used if you require very high precision. The γ function depends on several parameters that need to be measured. The full expression is:

$$\gamma(P) = \frac{1}{w_{Pump} * g} * \sqrt{4 * \frac{\Delta n(P)}{n}} \tag{A.8}$$

where $\Delta n(P) = P * R_{th} * \Delta n_{\Delta T}$, with P being the pump power, R_{th} being the measured thermal impedance of the chip, and $\Delta n_{\Delta T}$ being the shift in material index with temperature ($2.739e^{-4}K^{-1}$ for GaAs).

For most general purpose cavities, you don't need to worry about that complication and can treat the VECSEL as a Flat Mirror (and since Flat Mirrors are just identity matrices, you really don't need to include them in the math).

Now, when you're designing a laser cavity, the first thing you're going to do is pick what reference point to begin with. This is the point in the cavity at which you will start and finish your calculation of a round trip and will be the point you can then derive the beam width and curvature at. Most often, the spot that will make the most sense will be at your gain.

From this point, list the matrices corresponding to each portion of a full round trip through the cavity and multiply them in the correct order. See Fig A.1 and Eq A.9 for an example of this.

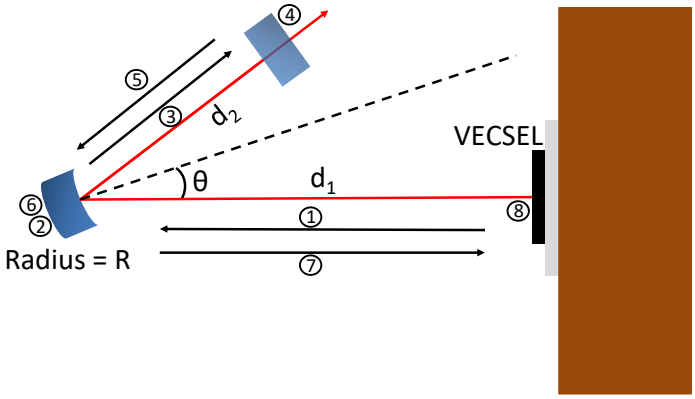

Figure A.1: Example V-cavity showing each step (numbered) that should be included in a ray transfer matrix analysis. See Eq. A.9. Note that since (4) and (8) are flat mirrors, it isn't strictly necessary to include them in the math.

$$Round\,Trip\;=\;\begin{bmatrix}1 & \frac{d_1}{n}\\ 0 & 1\end{bmatrix}_7 * \begin{bmatrix}1 & 0\\ -\frac{2*cos(\theta)}{R} & 1\end{bmatrix}_6 * \begin{bmatrix}1 & \frac{d_2}{n}\\ 0 & 1\end{bmatrix}_5 * \begin{bmatrix}1 & \frac{d_2}{n}\\ 0 & 1\end{bmatrix}_3 * \begin{bmatrix}1 & 0\\ -\frac{2*cos(\theta)}{R} & 1\end{bmatrix}_2 * \begin{bmatrix}1 & \frac{d_1}{n}\\ 0 & 1\end{bmatrix}_1 \tag{A.9}$$

When you've multiplied them together, the resulting matrix should still be in ABCD form, allowing you to check the stability using Eq. A.1 and use Eq A.10 to calculate the beam **radius** at that spot in the cavity. Caution should be given to how you construct your matrices; one can note that in the example, I've folded over a curved mirror, but only calculated a round trip ABCD for the saggital plane of the cavity. To be complete, I should really do the same thing for the tangential plane (swapping elements (2) and (6) for tangential-plane matrices given in Eq. A.6), then calculate the beam width for both projections as the spot should wind up elliptical. I'm lazy, though, so I'm only writing that matrix multiplication out in Tex once.

$$[beam\;radius] \quad w = \sqrt{\frac{\lambda}{\pi}} * \frac{\sqrt{|B|}}{\sqrt[4]{1-(\frac{A+D}{2})^2}} \tag{A.10}$$

The next and final step is to propagate your beam through the rest of your cavity so that you can see how it evolves. A necessary component to this is the beam curvature, obtained similar to the beam width:

$$[beam\;curvature] \quad R = \frac{2*B}{D-A}. \tag{A.11}$$

With both the beam width and curvature, you can use what's called the *complex beam parameter*, usually denoted by "q". The q parameter is written in terms of the curvature and width (Eq. A.12) and can be acted upon by an ABCD matrix (Eq. A.13). The idea is that you start with your width and curvature, calculate your q parameter, act on that q with an ABCD element, and back out your new width and curvature from the new q'. You can see how this can be done procedurally to determine the beam profile for every point in the cavity; it'd be painful by hand, but thankfully not as painful in Python, Matlab, etc. If you're interested, equations A.10 and A.11 come from this sort of procedure. Remember I said that you get them by demanding self consistency? That just means forcing $q' = q$ after a full round trip of the cavity.

$$[complex\;beam,\;or\;\text{“}q,\text{”}\;parameter] \quad \frac{1}{q} = \frac{1}{R} - \frac{i\lambda}{\pi n w^2} \tag{A.12}$$

$$[q\;parameter\;after\;ABCD\;element] \quad \frac{1}{q'} = \frac{C + D\frac{1}{q}}{A + B\frac{1}{q}} \tag{A.13}$$

Once you have mapped out the beam evolution through a full round trip, you're pretty much done with the numerics. You can then look at whether the spot sizes are feasible and how the cavity may shift if you move this mirror or that by a cm or two. After you've decided that your cavity is numerically sound, though, it's time to go align the damn thing...

A.3.2 Laser Alignment: Using Reference Beams

An extremely useful tool for aligning lasers is another laser that's already been aligned. It's pretty obvious, right? Carpenters mark out cuts they're about to make with pencils and rulers; Builders mark out foundations and walls they're about to build with string and chalk; Laser jocks mark out lasers they're about to set up with other lasers!

But using reference beams effectively requires a slight bit of strategy. There are many different types (Diode and HeNe are probably the most common) to pick from and best practices for setting up reference paths are not always immediately obvious.

For example, if you want to swap mirrors in a cavity, you may be tempted to lazily point the reference beam at the mirror, mark the spot it reflects to, switch the mirrors, and point the reference beam at the same spot with the new mirror. And this will work... if you're lucky. If you're not lucky, variations in position (placing a mirror assembly incorrectly; using mirrors with slightly different thicknesses; etc.) will mean that the tip/tilt alignment is now wrong when you match the reference beam's former position. *The best practice is to always overlap the reference beam with your laser's beam as completely as possible.*

The first thing you need to do is pick which reference laser you want. The classic is the HeNe as they're relatively cheap, are very safe and stable, and are very well established. HeNe lasers are great for almost every sort of cavity you could think of aligning. That said, say hello to this little guy:

This is a mW-class green diode that you can buy for $< \$20$ off the internet and is, by most standards, an absolutely terrible laser. It's based on second harmonic generation (SHG), meaning it actually emits converted infrared light. Leftover infrared from conversion inefficiencies would ideally be filtered out, but...uh... you're paying less than twenty dollars for a laser. Expect an eye-unsafe amount of infrared to also be present. Further, it's collimating lens and SHG crystal are temperature sensitive and can sometimes kill the brightness of the green. And you have to solder a power supply onto it yourself.

All that said, these cheap little diodes are incredible reference lasers. Not only are they cheap so you aren't particularly worried about accidentally killing one by dumping too much power into it from the laser you're working on, but as you'll see when I present an example of an alignment procedure, the fact that it has both green and leftover IR can actually be leveraged to your benefit.

Now, the next question is how to use it. You could probably figure this out on

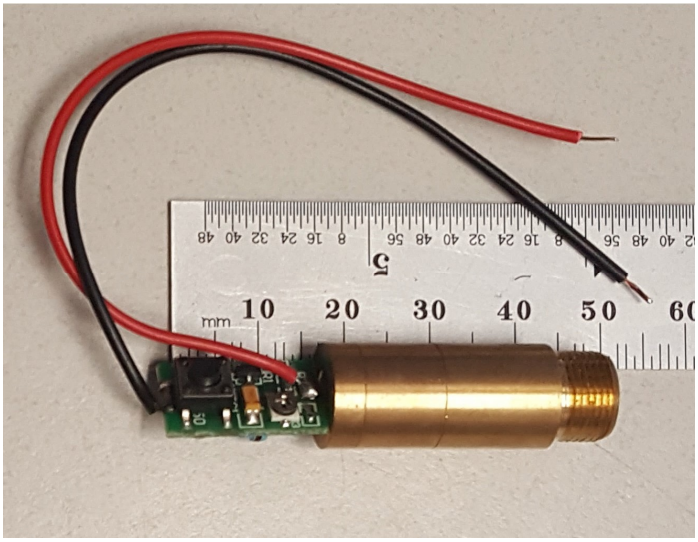

Figure A.2: 30mW Green SHG-based Alignment Diode.

your own, but to align a simple cavity, the easiest method is to ensure your reference beam is parallel to the table, then shoot it through an aperture (The easiest aperture to use? a piece of paper with a hole poked in it), ensure it hits relatively close to the pumped area of your gain, retro-reflect it through the aperture using a flat end mirror, and then move the second end mirror into place. This process is shown in Fig A.3.

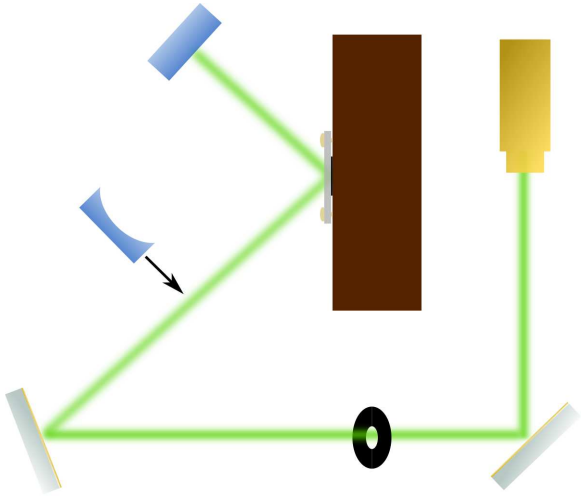

Figure A.3: Aligning a reference for a simple V-cavity. Ensure the beam is level to the table, retro-reflect with the flat mirror, and insert the curved mirror, using the reflection of the back to bring the curved mirror close to alignment.

In most cases, the point of the reference beam is not to attain perfect alignment, it is to get close enough that other techniques (see the mirror mount wiggling section) are virtually guaranteed to get the laser aligned properly. Naturally, the more time spent carefully aligning the reference, the easier it is to get the laser working afterwards, but there is a definite trade off: simple cavities are very easy to align if you use the reference beam for nothing more than centering your curved mirror(s) (You should always be aiming to get your cavity mode as close to a curved mirror's center as possible!). In this case, it's not really worth spending half an hour painstakingly getting the reference aligned. In more complicated cavities, however, a solid effort with the reference laser can save you hours of headache later.

If your aim is to replace a mirror, the procedure is going to be very similar. Only in this case, you're going to shoot the reference through the back of another mirror, match it to the working laser beam as best as you can manage, and use an aperture to track the back reflection from the mirror you're replacing. The easiest way to do this if you're using the diode I mentioned is to put the diode in a tip/tilt mount, shoot the working laser at it (at low power), and use the diode's tip/tilt mount to overlap the beams.

This seems very basic, but I'm mentioning it to bring attention to this: You are relying in this case on your ability to use the working laser to overlap your reference to the proper beam path. It is generally not a good idea to shoot your reference through a mirror without a working laser to overlap with, as refraction will bend the reference going through the mirror and disrupt its ability to act as reference. I've always found it easier to get a laser working by aligning the reference with one mirror out of the cavity (like in Fig A.3), then inserting that mirror, centering it on the reference beam, and wiggling (see next subsection). It may not matter for simple cavities, but that small deviation from refraction will significantly hamper your ability to reference more complicated ones.

Finally, I want to give you an example of how Dr. Maik Scheller and I used references to align a very tricky cavity: the F-cavity. Pictured in Fig A.4, the F-cavity is a relatively complicated multi-fold cavity where most of the difficulty comes from getting both passes through the gain to overlap almost perfectly without the laser being operational before-hand. So in order to align it, we made use of every reference beam trick we could think of, including the fact that our reference diode emitted both green and IR.

The procedure goes as follows:

1. Align the reference beam as shown in Fig A.4(a).

2. Move the flat mirror (2) in; get laser to lase.

3. Overlap reference beam with working laser.

4. Move the flat mirror (2) out again and bring flat mirrors (1) and (3) into F-cavity position, as in Fig A.4(b).

5. Using the IR portion of the reference beam, overlap both gain passes with pump spot by imaging the gain surface onto a camera.

6. Reinsert flat mirror (2), align to reference aperture.

7. Wiggle flat mirror (3) to complete alignment.

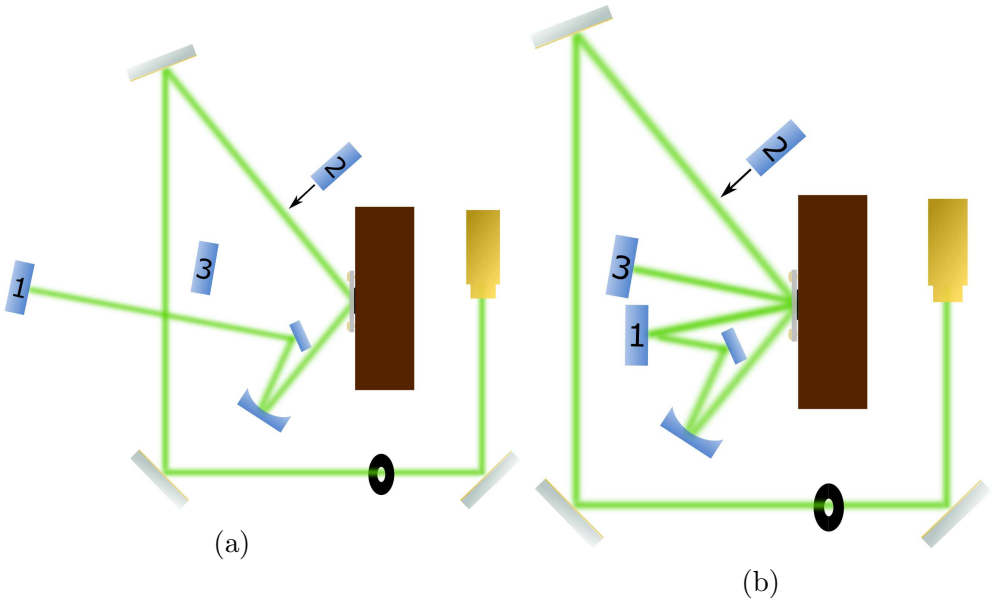

Figure A.4: (a) Alignment of simple pre-F cavity and (b) Alignment of full F-cavity. Not pictured in (b), but integral, is the imaging system to ensure good overlap of the two different passes through the gain and the pump spot

The first two steps (aligning the simpler cavity) are done as a check for the reference alignment. Using the image of the two reference beam spots on the gain surface requires a lot of confidence that it is an effective reference, so we ensure it is by using a cavity with similar set up (you don't have to move the curved or linking flat mirror afterwards) that doesn't have the two passes and is thus easier to make lase. Mixing careful procedures like this with a general understanding of where the "good enough" threshold for which cavity you're building will be very beneficial.

A.3.3 Laser Alignment: Advanced Mirror Mount Wiggling Techniques

If I have made any lasting contribution to science it will most likely turn out to be none of my published work, none of the data I have collected, nor any experiment

that I have ever undertaken. It will most likely be the formal codification of mirror mount wiggling.

As mentioned before in the *Brief Flash of Hope* section, optics happens on a timescale that is ridiculously faster than human time scales. And that means that chaotic motion on human time scales can sometimes lead to temporary stability on optical ones.

For laser alignment, this means that if - in a two-mirror cavity - one mirror is relatively well placed (like, for example, if it were set up with a reference beam), you can sort of wave the other mirror around haphazardly and if you're anywhere near the ballpark of good alignment (and the gain is functional, etc.), odds are you'll see a brief flash from the laser eventually. The following information is sort of a procedure for doing that haphazard waving motion in a manner that gives you the best odds of converting brief flashes into an aligned laser.

So the first thing is, unsurprisingly, to set up one mirror with a reference beam! Take the example of a simple V-cavity in Fig A.3, where you would place the flat mirror, retroreflect the reference beam, and then place the curved mirror in, centered on the reference. The very next step is to wiggle.

Your mirror is likely in a tip/tilt mount with finely threaded screws controlling orthogonal tilt axes, like in Fig A.5. Bring it into relatively close alignment using back reflections of the reference. If this is a curved mirror, be sure to overlap the normal-size spot (off the flat back of the mirror) and divergent spot (off the curved surface) as much as possible - if these aren't overlapped, that means you're off center

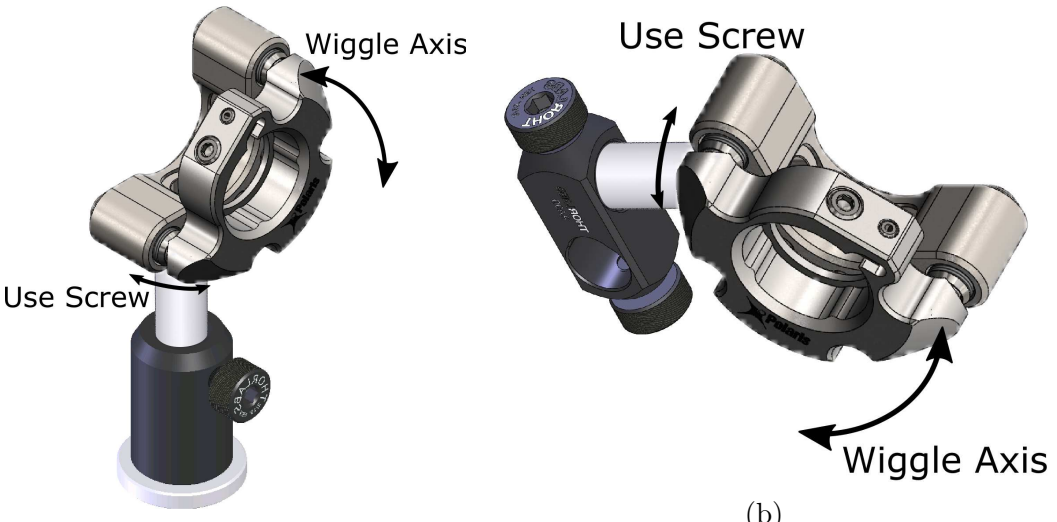

Figure A.5: (a) Proper wiggling form for mirror mount in standard assembly. (b) Wiggling form when mirror is braced horizontally.

of the mirror.

Then, adjust the screws so that the most well-braced axis is tilted farther back than ideal. This is usually the vertical axis tilted upwards, as it's braced against the post assembly that's clamped to the table (wiggling horizontally in this case may unscrew your mount from the post!). Now, quickly tilt that direction back and forth while slowly adjusting the orthogonal direction with the screws. This is depicted (poorly) in Fig A.5(a).

If everything else is working, the laser should flash at some point and since you were adjusting one axis with the actual mount screw, that means it should be close to alignment on that axis. At this point you can simply adjust the axis you were wiggling by using its mount screw until the laser lases. If it won't lase, try backing off, wiggling again, and being extra careful that you're not introducing off-axis tilt while doing so (e.g. slightly wiggling horizontally while trying to wiggle vertically).

For more complicated cavities than the 2-mirror example I just gave, the main trick is simply understanding that any linking mirrors are not that difficult, so long as there isn't another problem at play (such as a mirror clipping a beam). If you use a reference beam to keep yourself honest, the only challenge will be the two end mirrors in most cavities that don't feature some added complication (e.g. the F-cavity's requirement that the wide and narrow folds overlap on the gain), so you can set the rest of the mirrors and one end mirror with your reference, then wiggle the second end mirror.

You can come up with your own wiggling style yourself, but I would caution against methods that leave the potential to disturb more than one degree of freedom. For example, when I was young and foolish I often attempted to wiggle the entire mirror assembly including post holder. This often worked, but the astute reader will notice that moving the entire assembly could also move the mirror around on the table, so it was way less reliable than clamping the post holder down and wiggling like I've described above.

A.3.4 Laser Alignment: So Your Laser Doesn't Work. What Do?

Uh oh. You've gone and done it now, haven't you? You've got the reference beam fairly well aligned, you're wiggling like a pro, and not only does the laser not work but there's no immediately obvious reason why. How could you?! Now you're going to have to fix it.

Don't beat yourself up too badly. There are many, many, *many* reasons that your laser may refuse to work and sometimes it can be quite challenging to pin down what's actually wrong. So to start this section, I'm just going to pretty much leave a bulleted list of possibilities that you can run through as a sort of checklist. If nothing else, it will distract you from the frustration of the laser not working!

- Check your mirrors:

  - Are the mirrors for the correct wavelength?
  - Is the curvature of any curved mirror correct for your cavity design?
  - Check reflectivity; It can be too low *or* too high (sometimes mirrors are so HR you can't actually see flashes through them on a detection card)!

- Check your cavity:
  - Is it the correct length?
  - Are there any mirrors or objects that are clipping the beam?

- Check your pump (I'm assuming optical pumping for this):
  - Is it the right wavelength?
  - Is it working right (same threshold/power scaling that you expect)?
  - Is it focused and/or modematched properly (sometimes your cavity mode and pump can both be elliptical with opposite skew!)
  - If your gain is on a bragg reflector, is the incident angle of your pump correct for the bragg's reflectivity curve?

- Check your gain:
  - Is there any visible damage?
  - Does the design have specific temperature requirements? (Sometimes significantly heating or cooling a device can be required!)
  - Are you just pumping a region that doesn't work or have high enough gain? Can you try a new spot?
  - Did you cool it (or any other element) down too much and let water condense on the surface?

If you checked all of these things and it still refuses to work, there are still some procedures you can try before writing off the laser completely. The go-to method is begging. However, when that fails (it usually does; lasers are quite heartless) the next go-to method is imaging the photoluminescence, or "PL," of your gain onto a photodiode and watching the signal on an oscilloscope. The PL of a device is still sensitive to cavity alignment even if the gain isn't hitting lasing threshold, so you can carefully optimize your cavity's alignment by optimizing the PL signal you record, checking your reference beam occasionally to ensure you aren't optimizing towards a measurement artefact (e.g. you could be catching stray reflections or something with the diode instead of slowly working towards good alignment; weird pointing of the reference can indicate this). If you've fully optimized the PL signal without seeing threshold, try sweeping temperature and pump level parameters again just to be sure.

Another test to try is making the most simple cavity you're comfortable with and design it to be multi-moded. For some, that's a linear cavity; for others that may be a simple V. In general, it is best to try a multi-mode cavity if you are having difficulty getting your device to work. Focusing the pump harder can often also be a good idea as it increases the power densities available to you. $1\frac{kW}{cm^2}$ in the case of an ideal device.

If you've gone through all of the above and the laser *still* won't cooperate, it's probably time to start considering that maybe it's a bad sample. You could try another gain sample in the same cavity; getting one device to lase easily and failing repeatedly with another device under the same conditions is a pretty good indicator that it's the device's fault. Alternatively, this is a good point in the process to go complaining to the nearest post doc or research professor that everything sucks. They'll probably get it working in 5 minutes and you'll feel bad, but such is life in grad school. We've all been through that portion of our careers.

A.3.5 Fiber Coupling: How to start

So your laser works fine, but now your advisor and/or that mean post doc wants you to couple it into some single-mode fiber (or worse, Photonic Crystal Fiber). Worry not, this may at first seem like you're trying to hit a fly with a blow gun from 20 meters while blindfolded, but just like laser alignment, there are ways to make this less difficult.

You've got your laser beam. You've got your fiber. You've got your lens. You are well aware that the core concept is to use the lens to put the laser beam into the fiber, so you carefully line up the laser beam with the fiber and then carefully place the lens one focal length away from the fiber tip, which you have meticulously measured out (because you are an awesome grad student and do things *right*)...

...aaaaaand you've just made everything harder for yourself. Meticulous placement of the lens at the proper focal length away from the fiber is definitely what you eventually want, as that is, of course, optimal for coupling efficiency. However, for first getting light into the fiber, backing the lens off the focal length is going to be way easier; you'll couple less light in, but it will be a lot easier to get an initial signal through the fiber which you can then optimize by slowly moving the lens back into focus (don't go *crazy*, maybe try x1.5f or so).

Be sure to watch the fiber just after its housing at the input side. It's likely that you're going to couple into the cladding at first, from which light scatters out more easily than in the core. By watching for that light scattering out of the cladding, you can keep your alignment honest if you start straying away from the fiber completely. Try to optimize the amount of light you see, even if it's in the cladding at first, and eventually you will probably see a signal through the core at the output. If you're working with infrared, a handheld IR viewer is very helpful for this.

Another technique in a similar vein to the above was suggested by a fellow graduate

student of mine, Sam Nerenberg, who said he always optimizes coupling through a multi-mode fiber (core sizes ranging in the 100s of μm) first, then replaces it with whichever fiber he actually wants to couple into. The core concept is sort of the inverse of moving the lens out of focus: if you're trying to hit a target with darts while blindfolded, you can either make the darts really big (moving the lens) or the target really big (using MM fiber).

A.3.6 Fiber Coupling: Working with Non-Simple Fibers

A simple fiber (single core, single cladding, no microstructure or polarization-maintaining rods) is fairly easy to optimize. If light gets through the whole fiber, it's almost guaranteed to be in the core, since the cladding suffers heavy scattering and bend losses.

When you start to work with more complex fibers, though, things get a little more murky. Consider the fiber profiles shown in Fig. A.6. It is possible that in some of them, parts of the structure that aren't the core (e.g. the stress rods in polarization maintaining fibers) can also support modes with efficiency that may make it difficult to distinguish light propagating through them and light propagating through the core, if all you're looking at is the power at the other side of the fiber.

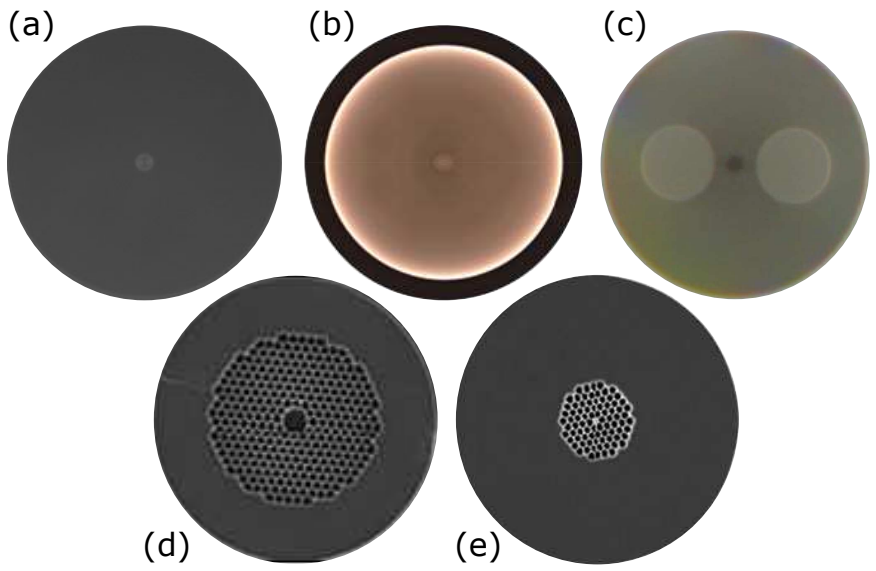

Figure A.6: Fiber profiles for Thorlabs (a) single mode, (b) double clad (the black ring is the second cladding), (c) Polarization maintaining, (d) hollow core photonic crystal, and (e) highly nonlinear, solid core photonic crystal fibers.

So how do you know you're actually optimizing the core mode? Well, there are a handful of checks you can make to convince yourself which part of the fiber you're hitting at the input. The most obvious way to tell is probably imaging the output of

the fiber onto a camera. When you find the proper focal plane, you should be able to determine the modes of the different fiber structures by their image on the camera.

An even simpler method for checking based on the look of the output, though, is to check the divergence of the beam as you translate the lens or input fiber tip. Depending on the fiber, the core will often have a dramatically different NA than the rest of the fiber structure; most often a much higher NA, but not always. Since the divergence of a fiber's output depends on the NA of the structure confining the light, you can often tell if you're in the core or not by watching with a detection card very close to the output tip and raster-scanning the input. For example, if you have a double-clad fiber where the NA of the core is much higher than the NA of the inner cladding, it's easy to mistake light coupled through the inner cladding for light through the core, but if you look near to the output and fiddle with the input coupling, it's likely you'll see a brief flash of light that's much, much more divergent than what you were seeing. That high-divergence mode is likely the core mode you're trying to hit.

A.3.7 Working with Fiber Amplifiers

If you're using a fiber amplifier that does not require free-space coupling at all, you can go ahead and skip this section, honestly. I mean, you're completely cheating and should acknowledge it, but that's okay. If you are coupling into and out of a (likely single-pass) amplifier from free space, there are probably a few things you should know.

The first is that aligning a fiber amplifier is much, much easier than aligning passive fiber. Just pump the amplifier lightly so that you can see the amplified spontaneous emission (ASE) from the un-seeded gain by eye or with a detection card. Then, align the backwards emission of the ASE to your seed laser's beam path. That is generally enough to get an optimizable signal and should take 30 seconds once all the optics are on the table.

The more difficult part of working with amplifiers is handling the power that they're capable of. This includes not blowing up the amplifier itself and not blowing up your seed laser or pre-amplifiers.

First in the backwards direction: if you've managed to couple your seed laser into your amplifier with good efficiency, you've also managed to do two things. First, you've efficiently coupled any backwards emission from the amplifier into any laser or amplifier behind it in the optical line. Second, if your seed laser is not also a fiber laser, you've managed to align the back of your output mirror well enough that it probably reflects right back into the amplifier, producing essentially half of a resonator with the amplifier fiber. Sure, the Q of that resonator may not be as high as a well-designed cavity, but the feedback into the amplifier will likely be well above enough power to destabilize the amplifier's emission.

For this reason, *always* isolate an amplifier from the previous portions of your

optical line. Even though in some cases (e.g. VECSELs) it's unlikely that a fiber amp will pose a threat to the gain or the stability of your seed's operation, in our lab, we have always seen spectral instabilities in the amplifier when we have operated a VECSEL-seeded fiber amp without optical isolation for backwards emission from the amplifier.

Depending on the power level of the amplifier, it can also be a threat to itself without any other element being necessary. There are going to be fresnel reflections at both tips of the fiber, and even they can be enough to mess with operational stability. And at very high power levels, the gain itself can become a physical threat to the fiber if the amp isn't seeded properly.

To address the former issue, the standard solution for low- to mid-power amplifiers is to cleave the input and output tips at an angle. This reduces the coupling efficiency for the reflection off the inside of the tip as it tries to head back through the fiber.

For high-power amplifiers, another additional step can be taken, which is to add a portion of bulk material to the end of the fiber. This way the beam can couple out of the core without any fresnel reflections (as it "feels" no interface) and expand before it is reflected back at the core by the surface of the material. Combined with angle-cleaving, this further reduces the coupling efficiency of any fresnel reflections off the output surface.

Practically, this is called "end capping" and is usually only necessary at the output tip of the amplifier as backwards emission should be minimal if the amp is seeded properly (and if the amp isn't seeded properly, you've got bigger problems). The process is usually to splice multi-mode or coreless fiber to the output tip, and then cleave that additional fiber to within 100-500 μm. This thickness requirement is to ensure clean beam expansion so that interference from the walls of the end cap does not distort the beam's profile later. You can also angle-cleave an end cap for additional protection. The principle of end capping is visualized in Fig. A.7

Finally, and most dramatically, you want to make sure your amplifier is seeded and remains seeded for the entire time it is being run at high power (which also means ensuring that your seed is in stable, pulsed operation, if the amp is intended to amplify pulses). This means using some form of interlock (a feedback device that shuts off the amp's pump if there's a problem with the seed). See the practical devices section for my favorite interlock design for pulsed operation.

The reason maintaining the seeding is so important is that the gain is so high in some amplifiers that without a seed exhausting all the excited carriers, ASE will build up inside the fiber. And at high enough powers, that ASE will begin to get a little unstable and actually pulse itself. Then, that Q-switch instability will build up and quickly reach damage thresholds for the fiber tips. One Q-switch ASE pulse then blows a portion of glass off the tip of the fiber and a little bit later the next pulse comes through to blow up more. It's actually pretty awesome. It's called "fiber fusing" and I would highly recommend you watch a video of this on the internet (but,

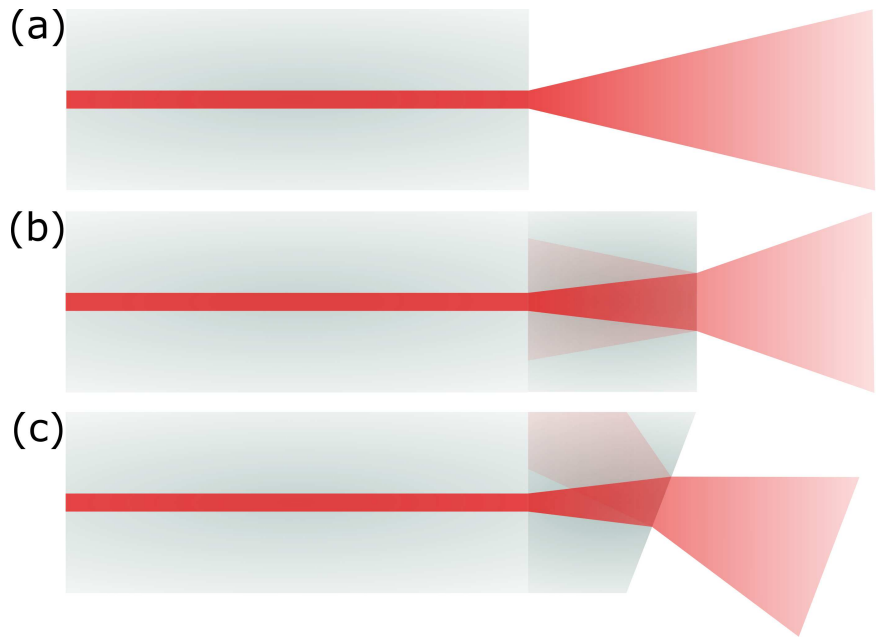

Figure A.7: Schematic drawing for output and back-reflection at (a) a standard fiber tip (note that fresnel reflection couples directly back into the core), (b) a fiber with an end cap, and (c) a fiber with an angle-cleaved end cap.

you know, maybe don't try it with your advisor's $10,000 amplifier...).

A.3.8 Checking Your M^2 With a CCD Camera

So you have your laser. It works; there's light coming out of it; nothing is on fire... And then someone (probably your advisor) asks you "what's the beam quality?". And your answer is probably "it's... I mean, it's *working*, what more do you want from me?!"

The correct answer to this question, though, is some reference to the M^2 value of the beam, which is a parameter that describes the deviation from an ideal Gaussian, TEM_{00} beam profile. The lower limit on M^2 values is 1, which represents a perfect TEM_{00} beam. As M^2 values increase, beam qualities get progressively more deviant. This affects the minimum waist that the beam can be focused to and also the divergence angle of the beam during propagation.

There are devices, aptly named Beam Profilers, that can directly give you the M^2 value for your laser with fairly good precision. Commercial versions of these devices are also fairly expensive, so you may be tasked with determining your laser's beam quality without one to back you up, in which case you're going to need a lens with a long-ish focal length (probably $\geq 20cm$), a webcam or other CCD camera, and a python script (or MATLAB, if you're partaken with demon worship).

The general procedure is to collimate the laser, focus it with the longer-length lens (which is used so that the waist is large enough for the camera to be useful), and then record five spots near the waist and five spots at least one Rayleigh Length, $z_R = \frac{\pi * w_0^2}{\lambda}$, away from the waist (and carefully record the positions of all those measurements along the propagation axis).

Next, you calculate the second moment width (often written as the "$D4\sigma$" width) of each spot you measured. The analytic form for this calculation based on the intensity profile, $I(x, y)$, is

$$D4\sigma = 4 * \sqrt{\frac{\int\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})^2 dxdy}{\int\int_{-\infty}^{\infty} I(x,y) dxdy}}, \tag{A.14}$$

where

$$\bar{x} = \frac{\int\int_{-\infty}^{\infty} I(x,y)\, x\, dxdy}{\int\int_{-\infty}^{\infty} I(x,y) dxdy}. \tag{A.15}$$

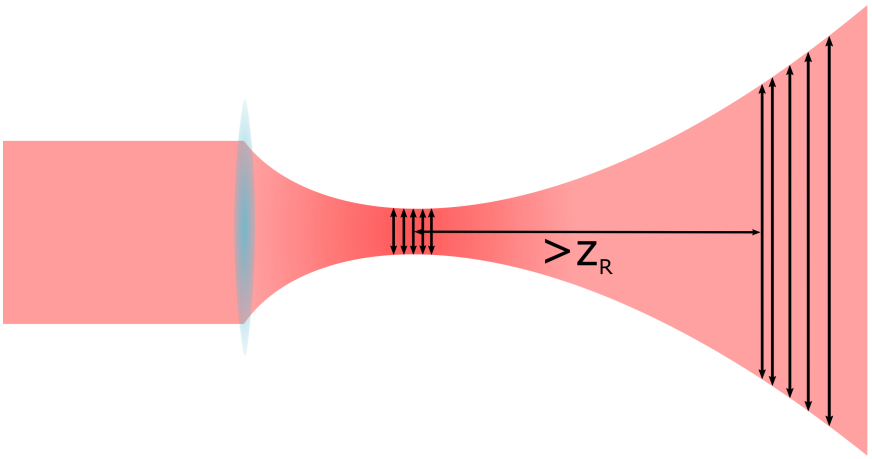

Figure A.8: Schematic for M^2 measurement.

In terms you can use in a computer program, the integrals are essentially directional sums over the intensity array measured by your CCD camera. For example, to calculate $\int\int_{-\infty}^{\infty} I(x,y)\, x\, dxdy$, you would get your MxN data array from the CCD, create a second array that uses the sensor's width and pixel dimensions to calibrate each pixel to a real-space location (each element of the data array = a pixel), sum the data array by column, and finally sum that summed array with each element multiplied by it's corresponding real location. This process is shown for the integral above in Eqs. A.16 - A.19:

$$[Data \quad array \quad from \quad CCD] \quad \begin{bmatrix} d_{00} & d_{01} & \dots & d_{0n} \\ d_{10} & d_{11} & \dots & d_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ d_{M0} & d_{M1} & \dots & d_{MN} \end{bmatrix} \tag{A.16}$$

$$[Data \quad after \quad y \quad \text{“}integral\text{”}] \quad \begin{bmatrix} \sum_{m=0}^{M} d_{m0} & \sum_{m=0}^{M} d_{m1} & \dots & \sum_{m=0}^{M} d_{mN} \end{bmatrix} \tag{A.17}$$

$$[Pixel \quad position] \quad X = \begin{bmatrix} 0 & 1 & 2 & \dots & N \end{bmatrix} * \frac{sensor \quad width}{\#\quad of \quad pixels \quad wide} \tag{A.18}$$

$$\int \int_{-\infty}^{\infty} I(x,y)\, x\, dxdy = \sum_{n=0}^{N} (X[n] * \sum_{m=0}^{M} d_{mn}) \tag{A.19}$$

When you follow this process for each integral in the second moment equation (Eq. A.14), you can obtain that spot's $D4\sigma$ width. Once you have $D4\sigma$ values for every measurement you took (that should be 10 of them; they're $\sigma(z)$ below), you can use those values and your carefully-measured positions for each measurement to fit them to

$$\sigma^2(z) = \sigma_0^2 + M^4 \left(\frac{\lambda}{\pi \sigma_0} \right)^2 (z - z_0)^2, \tag{A.20}$$

where z_0 and σ_0 are the waist location and second moment width of the waist, respectively. You can obtain M^2, z_0, and σ_0 from this fit, but you're mostly interested in your M^2 value, which you can finally go give to your advisor (or hide, if it's really bad).

Whew.

A.4 Practical Lab Devices

In this section, we're moving onto practical devices you can make in the lab. I'm sure there are commercial products for each of these devices, but the idea here is that we can maybe use common, cheaper parts which, because they're cheap, you may just have laying around in the lab to begin with. I would caution against getting too excited about the whole Do-It-Yourself thing. If you let the mentality run away with you, you'll wind up convincing your advisor that they can save tons of money by just putting pressure on you to build equipment instead. You should avoid giving this impression *at all costs.*

A.4.1 Variable Beam Sampler

This is one of the first things that Dr. Maik Scheller taught me to do in our lab and the genesis of my idea for this particular practical devices section. It will seem super obvious immediately after you see it, but it might not have occurred to you without any prompting.

The reflection coefficient off a piece of glass at an angle is polarization dependent. That is easy to see from the fresnel reflection coefficients for s- and p-polarized light:

$$R_s = \left| \frac{n_1 cos(\theta_i) - n_2 \sqrt{1 - \left(\frac{n_1}{n_2} sin(\theta_i)\right)^2}}{n_1 cos(\theta_i) + n_2 \sqrt{1 - \left(\frac{n_1}{n_2} sin(\theta_i)\right)^2}} \right|^2 , \tag{A.21}$$

$$R_p = \left| \frac{n_1 \sqrt{1 - \left(\frac{n_1}{n_2} sin(\theta_i)\right)^2} - n_2 cos(\theta_i)}{n_1 \sqrt{1 - \left(\frac{n_1}{n_2} sin(\theta_i)\right)^2} + n_2 cos(\theta_i)} \right|^2 . \tag{A.22}$$

If you plot these for glass ($n \simeq 1.5$), it looks like Fig. A.9. One can see easily that for above maybe 30° there exists a significant difference between reflection coefficients for s- and p-polarizations. Particularly around 60-70°, the difference is quite dramatic.

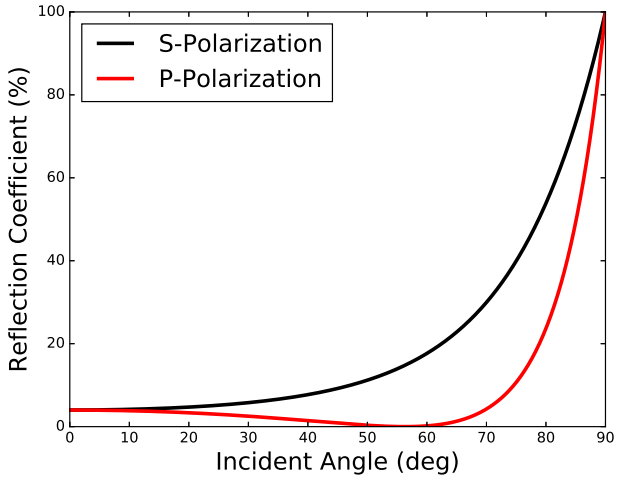

Figure A.9: Reflection coefficients for S- and P-polarized light

Microscope coverslips are extremely cheap, extremely common, easily-mounted pieces of glass. $\frac{\lambda}{2}$ waveplates rotate the polarization of light passing through them and are also fairly cheap and common. I'm sure you've probably got the idea right

about now: take your cheap, common piece of glass and tape/glue it to a post. Then, put it in your beam at ∼60°, after a $\frac{\lambda}{2}$ waveplate, and if your beam has one linear polarization, you now have a beam sampler that you can tune with your waveplate from close to 0% to upwards of 20% .

A.4.2 Poor Man's Spectrometer

Depending on the funding level of your lab, spectrometers can be either easily accessible or so coveted a resource that the graduate students have fractured into warring factions waging endless campaigns against each other for control of them. If you find yourself in a *Mad Max* style wasteland without a spectrometer, though, you may not be without recourse.

Find yourself any dispersive element and a webcam. Blaze gratings tend to have high dispersion so they tend to work best (most common spectrometers use gratings), but other dispersive elements such as prisms also work if you don't have access to a decent grating.

The idea is fairly simple in principle: a grating produces wavelength-dependent diffraction angles. Therefore, collimate the light from whichever source you're interested in and align a grating to it so that the first order diffraction is maximized. Now you have a wavelength-resolved spatial distribution which you can image with your webcam. Depending on your signal level and desired resolution the webcam can be used directly or can collect light scattered off a target (e.g. a piece of white paper).

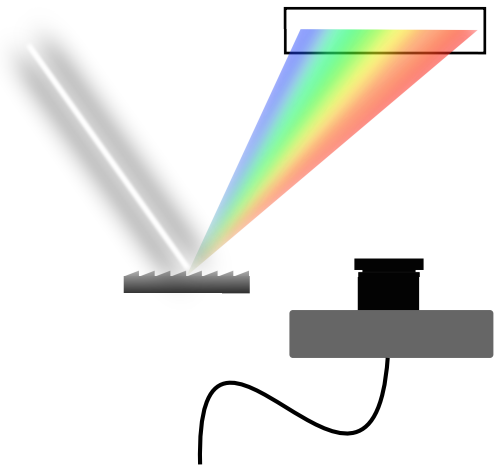

Figure A.10: Schematic for a simple grating spectrometer with target screen. Note that the projection angle onto the screen *must* be accounted for when calibrating spectra in this configuration.

There are limits to this; You'll need to take several calibration points by comparing with a calibrated spectrometer to ensure relative accuracy and you probably

shouldn't use grating/webcam spectrometer data for publications unless you *really* sink some time into its calibration (unless you've invested enough time to say with confidence exactly what parameters such as your spectral resolution are, I would recommend using a commercial spectrometer for publication data). However, for day to day operation such as knowing within 1-2nm accuracy what your wavelength is or discriminating between narrow CW and wide-band modelocking spectra, this sort of cheap, easily-made spectrometer is very useful.

A.4.3 Repetition Rate Mixing Interlock

Recall back to the section regarding fiber amplifiers and its mention of an interlock system. The purpose of an interlock system is to prevent your big, bad, high-power fiber amplifier from blowing itself up by ensuring that it has a steady stream of ~~sacrifi~~ seed pulses saturating its gain. An interlock must therefore be capable of doing at least two things: monitoring the pulse state of the seed laser and ensuring that pulses are still making it into the amplifier (e.g. that there's nothing blocking the optical path).

To give a little context: our first version of an interlock consisted of a second harmonic signal check of a 1 μm laser. We clipped a portion of the output beam of the amplifier, focused into an LB_3O_5 crystal, filtered out the infrared, and focused the green SHG light onto a photodiode. Because SHG is a nonlinear process, the power of the green light should drop dramatically if the seed stops pulsing (or gets blocked, but that's sort of obvious, isn't it?). We thus used a voltage comparison circuit (the same as in Fig. A.12) designed to trigger and cut power to the amplifiers when the green SHG signal dropped significantly.

This setup worked, but was extremely temperamental. It was sensitive to room

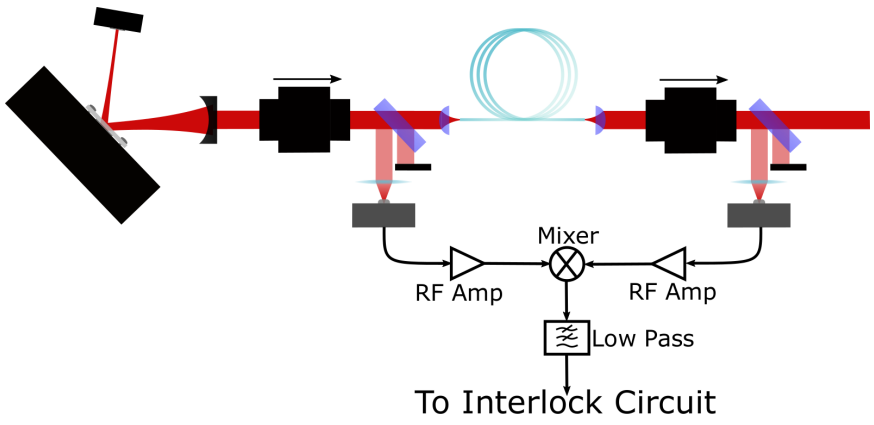

Figure A.11: Setup for a repetition rate mixing interlock. Exercise caution in selecting parts (e.g. the mixer *can not* have a DC block) and consider putting attenuators before RF amps and after the Mixer output.

light and the photodiode, being cheap, had a pronounced thermal drift, producing all sorts of problems with false readings if we ran the system for a long time.

Collaborators in Dr. R. Jason Jones' lab then suggested the improvement of switching to a repetition rate mixing interlock. Shown in Fig. A.11, this interlock scheme measures the laser's pulse train twice (and at least once after the amplifier to make sure pulses are getting through the whole system) and mixes their repetition rate signals (often called their "RF" signal since repetition rates are often in the MHz or GHz range of "radio frequencies"). Because we're always measuring twice from the same pulse train, we should be guaranteed that the two measured RF signals are always equal and thus will at least partially rectify in the mixer, giving us a DC voltage to put into a voltage comparison circuit (Fig. A.12). If either diode is blocked or if the seed laser stops pulsing, there will be no DC signal as the RF amplifiers (which should not amplify DC voltages) after both diodes should ensure that the CW beams do not produce enough DC signal for a false positive (it can also be a good idea to put 3-5 dB attenuators before the amplifiers and after the mixer to quiet down any noise in the system).

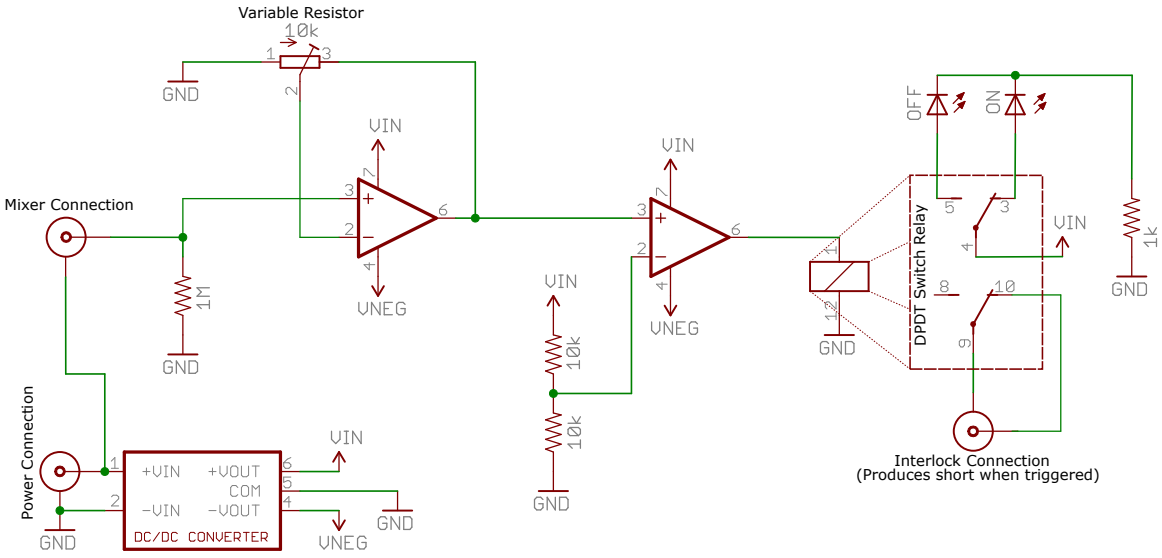

Figure A.12: Circuit diagram for voltage comparison intended to be used in a repetition rate mixing interlock. Part list in Table A.1. When triggered, produces an electrical short, which is expected by many laser current drivers.

When selecting parts for this, it is important to consider the repetition rate of the laser you're working with. I have included at the end of this section a table of parts for the circuit; because this only expects a DC input, the same circuit should work for pretty much any pulsed laser. The other components, shown in Fig. A.11, must be chosen specifically to work at your laser's RF regime. For example, the low-pass

Qty	Type	Specs
1	DC converter	24 V to ±15 V; 1W; 2 O/P
1	DPDT Relay	6 V DC; 1 A
2	Op Amp	5 MHz; 10-36 V
1	Potentiometer	10 kΩ; 250 mW; ±10%
1	LED (green)	2.1 V; 20 mA
1	LED (red)	2.1 V; 20 mA
2	Resistor	10 kΩ; 125 mW; ± 1%
1	Resistor	1 kΩ; 125 mW; ± 1%
1	Resistor	1 MΩ; 125 mW; ± 1%
3	Connector	generic SMA or BNC, your choice

Table A.1: Table of parts for interlock circuit. Specs are for the parts we used but are not strict requirements.

filter should be chosen with a low enough limit as to effectively block the system's RF signal and only allow the rectified DC signal from the mixer.

There are a couple of notes regarding the use of this setup. First, the DC signal that comes out of the mixer can have a variable sign and amplitude based on the phase between the two diode's signal. This phase difference is caused by a difference in the timing between the pulses hitting each diode, which is caused by the effective optical path difference between the two branches. This means that sometimes the interlock will seem to not work after adjusting things in the primary optical path of the system. Usually, this is simply a matter of connecting the output of the mixer to an oscilloscope, monitoring the DC signal, and slightly changing the distance to whichever diode is more convenient to move.

Another thing is that the DPDT relay (that stands for "double pull, double throw," a type of switch) may have a manufacturing defect that causes the switches to throw in the opposite direction. When I built this interlock the first time, it happened on the switch that controls the indicator LEDs (pins 3&5 in Fig. A.12). I checked the circuit more times than I could count and could only conclude the relay was throwing to pin 5 when it triggered instead of pin 3. It's a minor detail if the indicators are backwards, but if it happens to the interlock switch (pins 8&10 in Fig. A.12), you'll have to reverse those connections. Keep that in mind as a trouble shooting step if something doesn't work and the rest of the circuit checks out.

A.5 Inspirational Conclusion

This is where I write some deeply impactful words about mentality and perseverance and then tie it all together with a commentary about the indefatigable resilience of the human spirit. Then you're like "thanks coach" and go out and win graduate

school.

But let's be honest, that's a more than a little trite and you're pretty smart so instead I'll just kind of awkwardly end this appendix without any real sense of closure.

References

[1] Charles Townes. *A Century of Nature: Twenty-One Discoveries that Changed Science and the World*, pages 107–112. University of Chicago Press, 2003.

[2] Mark Kuznetsov. Vecsel semiconductor lasers: a path to high-power, quality beam, and uv to ir wavelength by design. *Semiconductor Disk Lasers: Physics and Technology*, pages 1–71, 2010.

[3] B. Heinen, T.-L. Wang, M. Sparenberg, a. Weber, B. Kunert, J. Hader, S.W. Koch, J.V. Moloney, M. Koch, and W. Stolz. 106 W continuous-wave output power from vertical-external-cavity surface-emitting laser. *Electronics Letters*, 48(9):516, 2012.

[4] Tsuei Lian Wang, Bernd Heinen, Joerg Hader, Colm Dineen, Mino Sparenberg, Antje Weber, Bernardette Kunert, Stephan W. Koch, Jerome V. Moloney, Martin Koch, and Wolfgang Stolz. Quantum design strategy pushes high-power vertical-external-cavity surface-emitting lasers beyond 100 W. *Laser and Photonics Reviews*, 6(5):12–14, 2012.

[5] Tsuei-Lian Wang. *High-Power Optically Pumped Semiconductor Lasers for Near Infrared Wavelengths*. PhD thesis, The University of Arizona., 2012.

[6] Marc Schmid, Sarah Benchabane, Firuz Torabi-Goudarzi, Richard Abram, Allister I Ferguson, and Erling Riis. Optical in-well pumping of a vertical-external-cavity surface-emitting laser. *Applied Physics Letters*, 84(24):4860–4862, 2004.

[7] F Gires and P Tournois. An active interferometer permitting translation or modulation of laser impulsion frequency. *COMPTES RENDUS HEBDOMADAIRES DES SEANCES DE L ACADEMIE DES SCIENCES SERIE B*, 268(4):313, 1969.

[8] Keith G Wilcox, Anne C Tropper, Harvey E Beere, David A Ritchie, B Kunert, Bernd Heinen, and Wolfgang Stolz. 4.35 kW peak power femtosecond pulse mode-locked VECSEL for supercontinuum generation. *Optics express*, 21(2):1599–1605, 2013.

[9] I.R. Kilen, S.W. Koch, and J.V. Moloney. Nonequilibrium pulsed femtosecond semiconductor disk laser, May 11 2017. US Patent App. 15/264,335.

[10] Peter Klopp, Florian Saas, Martin Zorn, Markus Weyers, and Uwe Griebner. 290-Fs Pulses From a Semiconductor Disk Laser. *Optics express*, 16(8):5770–5775, 2008.

[11] Alexandre Laurain, Declan Marah, Robert Rockmore, John McInerney, Jorg Hader, Antje Ruiz Perez, Wolfgang Stolz, and Jerome V Moloney. Colliding pulse mode locking of vertical-external-cavity surface-emitting laser. *Optica*, 3(7):781–784, 2016.

[12] Dominik Waldburger, Sandro M Link, Mario Mangold, Cesare GE Alfieri, Emilio Gini, Matthias Golling, Bauke W Tilma, and Ursula Keller. High-power 100 fs semiconductor disk lasers. *Optica*, 3(8):844–852, 2016.

[13] J.V. Moloney, I.R. Kilen, and S.W. Koch. Nonequilibrium pulsed femtosecond semiconductor disk laser, October 11 2016. US Patent 9,466,948.

[14] Alexej Chernikov, Jens Herrmann, Martin Koch, Bernardette Kunert, Wolfgang Stolz, Sangam Chatterjee, Stephan W Koch, Tsuei-Lian Wang, Yushi Kaneda, J Michael Yarborough, et al. Heat management in high-power vertical-external-cavity surface-emitting lasers. *IEEE Journal of Selected Topics in Quantum Electronics*, 17(6):1772–1778, 2011.

[15] K Gbele, A Laurain, Jorg Hader, W Stolz, and Jerome V Moloney. Design and fabrication of hybrid metal semiconductor mirror for high-power vecsel. *IEEE Photonics Technology Letters*, 28(7):732–735, 2016.

[16] Jorg Hader, T-L Wang, Jerome V Moloney, B Heinen, M Koch, Stephan W Koch, B Kunert, and W Stolz. On the measurement of the thermal impedance in vertical-external-cavity surface-emitting lasers. *Journal of Applied Physics*, 113(15):153102, 2013.

[17] Mario Mangold, C A Zaugg, Sandro M. Link, Matthias Golling, Bauke W. Tilma, and Ursula Keller. Pulse repetition rate scaling from 5 to 100 GHz with a high-power semiconductor disk laser. *Optics Express*, 22(5):6099, 2014.

[18] M Butkus, E A Viktorov, T Erneux, C J Hamilton, G Maker, G P A Malcolm, and E U Rafailov. 85.7 MHz repetition rate mode-locked semiconductor disk laser: fundamental and soliton bound states. *Optics express*, 21(21):25526–25531, 2013.

[19] Alexandre Laurain, Cody Mart, Jorg Hader, Jerome V. Moloney, Bernardette Kunert, and Wolfgang Stolz. 15 W single frequency optically pumped semiconductor laser with sub-megahertz linewidth. *IEEE Photonics Technology Letters*, 26(2):131–133, 2014.

[20] Lukas E Hunziker, Chris Ihli, and Daniel S Steingrube. Miniaturization and power scaling of fundamental mode optically pumped semiconductor lasers. *IEEE Journal of Selected Topics in Quantum Electronics*, 13(3):610–618, 2007.

[21] Maik Scheller, Joe M. Yarborough, Jerome V. Moloney, Mahmoud Fallahi, Martin Koch, and Stephan W. Koch. Room temperature continuous wave milliwatt terahertz source. *Opt. Express*, 18(26):27112–27117, Dec 2010.

[22] Pankaj Ahirwar, Thomas J. Rotter, Darryl Shima, Nahid A. Jahan, Stephen P R Clark, Sadhvikas J. Addamane, Ganesh Balakrishnan, Alexandre Laurain, Jörg Hader, Yi Ying Lai, Jerome V. Moloney, Ikuo Suemune, and Robert G. Bedford. Growth and optimization of 2-micron InGaSb/AlGaSb quantum-well-based VECSELs on GaAs/AlGaAs DBRs. *IEEE Journal on Selected Topics in Quantum Electronics*, 19(4), 2013.

[23] Stefan Baumgärtner, Hermann Kahle, Roman Bek, Thomas Schwarzbäck, Michael Jetter, and Peter Michler. Comparison of AlGaInP-VECSEL gain structures. *Journal of Crystal Growth*, 414:219–222, 2015.

[24] Jennifer E. Hastie, Lynne G. Morton, Alan J. Kemp, Martin D. Dawson, Andrey B. Krysa, and John S. Roberts. Tunable ultraviolet output from an intracavity frequency-doubled red vertical-external-cavity surface-emitting laser. *Applied Physics Letters*, 89(6):1–4, 2006.

[25] Jennifer Hastie, Stephane Calvez, Martin Dawson, Tomi Leinonen, Antti Laakso, Jari Lyytikäinen, and Markus Pessa. High power CW red VECSEL with linearly polarized TEM00 output beam. *Optics express*, 13(1):77–81, 2005.

[26] S.T. Keller, A. Sirbu, V. Iakovlev, A. Caliman, A. Mereuta, and E. Kapon. 85 W VECSEL output at 1270 nm with conversion efficiency of 59 %. *Optics Express*, 23(13):17437, 2015.

[27] H. Lindberg, M. Strassner, E. Gerster, and A. Larsson. 0.8 W optically pumped vertical external cavity surface emitting laser operating CW at 1550 nm. *Electronics Letters*, 40(10), 2004.

[28] Cherry May N. Mateo, Uwe Brauch, Hermann Kahle, Thomas Schwarzbäck, Michael Jetter, Marwan Abdou Ahmed, Peter Michler, and Thomas Graf. 25 W continuous wave output at 665 nm from a multipass and quantum-well-pumped AlGaInP vertical-external-cavity surface-emitting laser. *Optics Letters*, 41(6):1245, 2016.

[29] M. Rahim, A. Khiar, F. Felder, M. Fill, H. Zogg, and M. W. Sigrist. 5-micron Vertical External-Cavity Surface-Emitting Laser (VECSEL) for spectroscopic applications. *Applied Physics B: Lasers and Optics*, 100(2):261–264, 2010.

[30] S Ranta, M Tavast, T Leinonen, N Van Lieu, G Fetzer, and M Guina. 1180 nm VECSEL with output power beyond 20 W. *Electronics Letters*, 49(1):59–60, 2013.

[31] J. M. Rey, M. Fill, F. Felder, and M. W. Sigrist. Broadly tunable mid-infrared VECSEL for multiple components hydrocarbon gas sensing. *Applied Physics B: Lasers and Optics*, 117(3):935–939, 2014.

[32] Alexandre Laurain, Laurent Cerutti, Mikhael Myara, and Arnaud Garnache. 2.7-μm single-frequency TEM 00 low-threshold Sb-based diode-pumped external-cavity VCSEL. *IEEE Photonics Technology Letters*, 24(4):246–248, 2012.

[33] Emmi Kantola, Tomi Leinonen, Jussi-Pekka Penttinen, Ville-Markus Korpijärvi, and Mircea Guina. 615 nm GaInNAs VECSEL with output power above 10 W. *Optics Express*, 23(16):20280, 2015.

[34] Emmi Kantola, Tomi Leinonen, Sanna Ranta, Miki Tavast, and Mircea Guina. High-efficiency 20 W yellow VECSEL. *Optics Express*, 22(6):77–81, 2014.

[35] Yushi Kaneda, J M Yarborough, Li Li, N Peyghambarian, Li Fan, Chris Hessenius, Mahmoud Fallahi, Jörg Hader, Jerome V Moloney, Yoshiyuki Honda, Masato Nishioka, Youhei Shimizu, Kenshi Miyazono, Hiroya Shimatani, Masashi Yoshimura, Yusuke Mori, Yasuo Kitaoka, and Takatomo Sasaki. Continuous-wave all-solid-state 244 nm deep-ultraviolet laser source by fourth-harmonic generation of an optically pumped semiconductor laser using CsLiB_6O_10 in an external resonator. *Optics Letters*, 33(15):1705, 2008.

[36] Bauke W Tilma, Mario Mangold, Christian A Zaugg, Sandro M Link, Dominik Waldburger, Alexander Klenner, Aline S Mayer, Emilio Gini, Matthias Golling, and Ursula Keller. Recent advances in ultrafast semiconductor disk lasers. *Light: Science & Applications*, 4(7):e310, 2015.

[37] Caleb W Baker, Maik Scheller, Alexandre Laurain, Antje Ruiz-Perez, Wolfgang Stolz, Sadhvikas Addamane, Ganesh Balakrishnan, Stephan W Koch, R Jason Jones, and Jerome V Moloney. Multi-angle vecsel cavities for dispersion control and peak-power scaling. *IEEE PHOTONICS TECHNOLOGY LETTERS*, 29(3), 2017.

[38] M Fallahi, Li Fan, Y Kaneda, C Hessenius, J Hader, Li Hongbo, J V Moloney, B Kunert, W Stolz, S W Koch, J Murray, and R Bedford. 5-W Yellow Laser by Intracavity Frequency Doubling of High-Power Vertical-External-Cavity Surface-Emitting Laser. *IEEE Photon. Technol. Lett.*, 20(20):1700–1702, 2008.

[39] C A Zaugg, Alexander Klenner, Mario Mangold, A S Mayer, S M Link, F Emaury, M Golling, E Gini, Clara J Saraceno, Bauke W Tilma, and U Keller. Gigahertz self-referenceable frequency comb from a semiconductor disk laser. *Optics Express*, 22(13):16445–16455, 2014.

[40] Maik Scheller, Caleb W Baker, Stephan W Koch, and Jerome V Moloney. Dual-wavelength passively mode-locked semiconductor disk laser. *IEEE Photonics Technology Letters*, 28(12):1325–1327, 2016.

[41] Maik Scheller, Caleb W. Baker, Stephan W. Koch, Jerome V. Moloney, and R. Jason Jones. High power dual-wavelength vecsel based on a multiple folded cavity. *IEEE Photonics Technology Letters*, 2017.

[42] Chris Hessenius, Michal Lukowski, and Mahmoud Fallahi. High-power tunable two-wavelength generation in a two chip co-linear t-cavity vertical external-cavity surface-emitting laser. *Applied Physics Letters*, 101(12):121110, 2012.

[43] AH Quarterman and KG Wilcox. Design of a solar-pumped semiconductor laser. *Optica*, 2(1):56–61, 2015.

[44] Peter W Milonni and Joseph H Eberly. Laser physics. 2010.

[45] Herman A Haus. Mode-locking of lasers. *IEEE Journal of Selected Topics in Quantum Electronics*, 6(6):1173–1185, 2000.

[46] Hans J Weber and George B Arfken. *Essential Mathematical Methods for Physicists, ISE*, pages 689–707. Academic Press, 2003.

[47] Hermann Haus. Theory of mode locking with a slow saturable absorber. *IEEE Journal of Quantum Electronics*, 11(9):736–746, 1975.

[48] FX Kartner, ID Jung, and U Keller. Soliton mode-locking with saturable absorbers. *IEEE Journal of Selected Topics in Quantum Electronics*, 2(3):540–556, 1996.

[49] FX Kärtner and U Keller. Stabilization of solitonlike pulses with a slow saturable absorber. *Optics Letters*, 20(1):16–18, 1995.

[50] Oliver D Sieber, Martin Hoffmann, Valentin J Wittwer, Mario Mangold, Matthias Golling, Bauke W Tilma, and Thomas Su. Experimentally verified pulse formation model for high-power femtosecond VECSELs. *Applied Physics B: Lasers and Optics*, (113):133–145, 2013.

[51] I Kilen, Jorg Hader, Jerome V Moloney, and Stephan W Koch. Ultrafast nonequilibrium carrier dynamics in semiconductor laser mode locking. *Optica*, 1(4):192–197, 2014.

[52] J Hader, M Scheller, A Laurain, I Kilen, C Baker, JV Moloney, and SW Koch. Ultrafast non-equilibrium carrier dynamics in semiconductor laser mode-locking. *Semiconductor Science and Technology*, 32(1):013002, 2016.

[53] Charles Henry. Theory of the linewidth of semiconductor lasers. *IEEE Journal of Quantum Electronics*, 18(2):259–264, 1982.

[54] J Hader, SW Koch, and JV Moloney. Microscopic theory of gain and spontaneous emission in gainnas laser material. *Solid-State Electronics*, 47(3):513–521, 2003.

[55] A Girndt, F Jahnke, A Knorr, Stephan W Koch, and WW Chow. Multi-band bloch equations and gain spectra of highly excited ii–vi semiconductor quantum wells. *physica status solidi (b)*, 202(2):725–739, 1997.

[56] Alexandre Laurain, Robert Rockmore, Hsiu-Ting Chan, Jorg Hader, Stephan W Koch, Antje Ruiz Perez, Wolfgang Stolz, and Jerome V Moloney. Pulse interactions in a colliding pulse mode-locked vertical external cavity surface emitting laser. *JOSA B*, 34(2):329–337, 2017.

[57] Anthony J Demaria, William H Glenn, Michael J Brienza, and Michael E Mack. Picosecond laser pulses. *Proceedings of the IEEE*, 57(1):2–25, 1969.

[58] CA Zaugg, Z Sun, VJ Wittwer, D Popa, S Milana, TS Kulmala, RS Sundaram, M Mangold, OD Sieber, M Golling, et al. Ultrafast and widely tuneable vertical-external-cavity surface-emitting laser, mode-locked by a graphene-integrated distributed bragg reflector. *Optics express*, 21(25):31548–31559, 2013.

[59] Kai Seger, Niels Meiser, Sun Young Choi, Boo He Jung, D-I Yeom, Fabian Rotermund, Olof Okhotnikov, Fredrik Laurell, and Valdas Pasiskevicius. Carbon nanotube mode-locked optically-pumped semiconductor disk laser. *Optics express*, 21(15):17806–17813, 2013.

[60] U Siegner, R Fluck, G Zhang, and U Keller. Ultrafast high-intensity nonlinear absorption dynamics in low-temperature grown gallium arsenide. *Applied physics letters*, 69(17):2566–2568, 1996.

[61] ZY Zhang, AEH Oehler, B Resan, S Kurmulis, KJ Zhou, Q Wang, Mario Mangold, T Süedmeyer, Ursula Keller, KJ Weingarten, et al. 1.55 μm inas/gaas quantum dots and high repetition rate quantum dot sesam mode-locked laser. *Scientific reports*, 2:477, 2012.

[62] John Venables. *Introduction to surface and thin film processes*. Cambridge University Press, 2000.

[63] Edmond Treacy. Optical pulse compression with diffraction gratings. *IEEE Journal of quantum Electronics*, 5(9):454–458, 1969.

[64] Caleb Baker, Maik Scheller, Stephan W Koch, Antje R Perez, Wolfgang Stolz, R Jason Jones, and Jerome V Moloney. In situ probing of mode-locked vertical-external-cavity-surface-emitting lasers. *Optics letters*, 40(23):5459–5462, 2015.

[65] Ada Bäumner, Stephan W Koch, and Jerome V Moloney. Non-equilibrium analysis of the two-color operation in semiconductor quantum-well lasers. *physica status solidi (b)*, 248(4):843–846, 2011.

[66] A Chernikov, M Wichmann, MK Shakfa, M Scheller, JV Moloney, SW Koch, and M Koch. Time-dynamics of the two-color emission from vertical-external-cavity surface-emitting lasers. *Applied Physics Letters*, 100(4):041114, 2012.

[67] Fan Zhang, Mahmoud Gaafar, Christoph Möller, Wolfgang Stolz, Martin Koch, and Arash Rahimi-Iman. Dual-wavelength emission from a serially connected two-chip vecsel. *IEEE Photonics Technology Letters*, 28(8):927–929, 2016.

[68] Maik Scheller, Abram G Young, Joe M Yarborough, Jerome V Moloney, Stephan W Koch, Christian Y Drouet d'Aubigny, and Christopher K Walker. Heterodyne detection of intracavity generated terahertz radiation. *IEEE Transactions on Terahertz Science and Technology*, 2(3):271–277, 2012.

[69] Pascal Dupriez, Christophe Finot, Andy Malinowski, Jayanta K Sahu, Johan Nilsson, David J Richardson, Keith G Wilcox, Hannah D Foreman, and Anne C Tropper. High-power, high repetition rate picosecond and femtosecond sources based on yb-doped fiber amplification of vecsels. *Optics express*, 14(21):9611–9616, 2006.

[70] Adrian H Quarterman, Lucy E Hooper, Peter J Mosley, and Keith G Wilcox. Gigahertz pulse source by compression of mode-locked vecsel pulses coherently broadened in the normal dispersion regime. *Optics express*, 22(10):12096–12101, 2014.

[71] John M Dudley, Goëry Genty, and Stéphane Coen. Supercontinuum generation in photonic crystal fiber. *Reviews of modern physics*, 78(4):1135, 2006.

[72] Haohua Tu, Yuan Liu, Marina Marjanovic, Eric J Chaney, Sixian You, Youbo Zhao, and Stephen A Boppart. Concurrence of extracellular vesicle enrichment and metabolic switch visualized label-free in the tumor microenvironment. *Science Advances*, 3(1):e1600675, 2017.